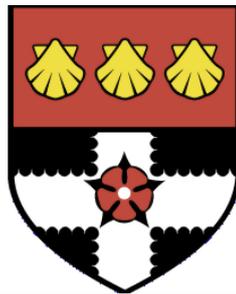


UNIVERSITY OF READING

Department of Meteorology



**Improving local corrections for
the radar vertical reflectivity
profile using the linear
depolarisation ratio**

Caroline Georgina Sandford

A thesis submitted for the degree of Doctor of Philosophy

July 2018

Declaration

I confirm that this is my own work and the use of all material from other sources has been properly and fully acknowledged.

Caroline Georgina Sandford

Acknowledgements

I would like to thank my supervisors Anthony Illingworth, Robert Thompson and Dawn Harrison for their creativity, patience and support throughout this PhD. I am very grateful to Malcolm Kitchen for his expertise and engagement with this project, which were a great source of encouragement to me. I would also like to thank Donal Adams from the Met Office radar hardware team, for his input on the radar network upgrades and reflectivity calibration scheme; and Nawal Husnoo, for his unfailingly prompt responses and invaluable insights during the writing of this thesis.

Abstract

A major source of errors in radar-derived quantitative precipitation estimates (QPEs) is the vertical reflectivity profile (VPR). A feature of particular importance is the radar “bright band”: a reflectivity enhancement due to the melting of large snowflakes that occurs in the majority of high latitude rainfall, and which if misrepresented can cause order of magnitude errors in surface QPEs. Recent upgrades of several national weather radar networks to dual polarisation provide opportunities to refine the identification of bright band in operational radar measurements, and to improve subsequent determination and corrections for VPR.

This thesis applies information from the linear depolarisation ratio (LDR) to improve classification of VPRs at the pixel scale. Using a unique dataset of high resolution vertical profiles, values of LDR in the melting layer are shown to provide a seven-fold increase in probability of detection of non-bright band reflectivity profiles over the current UK operational criterion. In the context of the Met Office local VPR correction scheme, an LDR-based classification step alone is shown to produce improvements in bias and RMSE of more than 1 mm h^{-1} for high rain rates in non-bright band conditions.

The high resolution vertical profile dataset is then further used to improve the parameterisation of VPR shapes for correction at the local scale. Using a simulation method adapted from previous literature, three possible non-bright band VPR shapes are defined and their performance objectively compared to a control (no correction) profile. The most skilful profile is applied to a high intensity non-bright band case study in combination with LDR-based VPR classification, yielding further improvements to QPE bias and RMSE. Improvements to the stratiform profile are also investigated. The conclusions of this thesis indirectly support the use of local over global VPR corrections to maximise the accuracy of radar QPEs at the sub-kilometre scale.

Contents

1	Introduction	7
1.1	Background	7
1.2	Principles of meteorological radar	9
1.2.1	Rayleigh scattering, reflectivity and the radar equation	9
1.2.2	Reflectivity and rain rate	11
1.3	Operational radar configurations	13
1.3.1	Scanning meteorological radar	13
1.3.2	Properties of the radar beam	14
1.3.3	Radar networks	16
1.4	Dual polarisation radar	17
1.4.1	Dual polarisation transmission modes	18
1.4.2	Interpretation and applications of SHV mode parameters	18
1.4.3	Single polarisation transmission: the linear depolarisation ratio	22
1.5	Quantitative precipitation estimation	24
1.5.1	Identifying and removing non-meteorological echoes	24
1.5.2	Power loss corrections for blockage and attenuation	26
1.5.3	Estimating reflectivities at ground level	27
1.5.4	Calculating surface rain rates	29
1.6	Impacts of dual polarisation for radar QPE and motivation for this study	31
2	Existing approaches to VPR: classification, determination and correction	34
2.1	Introduction	34
2.2	Microphysical basis of the VPR	35
2.2.1	Defining “stratiform” and “convective” rain	36
2.2.2	Stratiform VPRs	36
2.2.3	Convective VPRs	41
2.2.4	Other types of VPR	42
2.3	VPR determination and correction in radar PPIs	43
2.3.1	Introduction to terminology	43

2.3.2	Mean apparent VPRs: the ratio method	45
2.3.3	Mean field background VPRs	45
2.3.4	Localised VPR methods	49
2.3.5	Combining apparent and parameterised VPRs	51
2.3.6	Time-averaged precipitation estimates	52
2.3.7	The question of scale	53
2.4	Identifying non-bright band VPRs	54
2.4.1	Convective diagnosis in the SHY framework	55
2.4.2	Alternative approaches using volume reflectivities	56
2.4.3	Surface drop size distributions	58
2.4.4	Bright band identification	59
2.4.5	Combining classification algorithms	61
2.5	Operational VPR correction schemes	61
2.5.1	Météo France	61
2.5.2	Finnish Meteorological Institute (FMI)	63
2.5.3	Met Office (UK)	64
2.6	Summary	65
3	A high resolution VPR dataset for testing and verification	68
3.1	The Wardon Hill research radar	68
3.2	Data collection strategy	70
3.2.1	Scan details	70
3.2.2	Quality of the LDR measurement	71
3.3	Calibration	71
3.3.1	Reflectivity	72
3.3.2	LDR	74
3.4	Extracting meteorological profiles	76
3.4.1	Removal of non-meteorological echoes	76
3.4.2	Regridding to vertical profiles	77
4	Using the linear depolarisation ratio to identify non-bright band VPRs	79
4.1	Introduction	79
4.2	Defining VPR types within the high resolution dataset	82
4.2.1	Locating the melting layer	82
4.2.2	Observed profile types	83
4.2.3	Properties of the classified dataset	87
4.2.4	Quantitative peak size thresholds for “true” VPR types	87
4.3	Results and discussion	89
4.3.1	Comparing LDR skill with a high level reflectivity criterion	89
4.3.2	Bright band intensity	91

4.3.3	Sensitivity to VPR type definitions	92
4.4	Conclusions	95
5	Applying an LDR-based classification criterion to operational data	96
5.1	Introduction	96
5.2	Classification of VPRs using simulated LDR measurements	99
5.2.1	Radar simulator	100
5.2.2	Sampling regimes: defining the melting layer region	100
5.2.3	Overall skill of LDR measurements	102
5.2.4	Range dependent LDR thresholds	103
5.2.5	Results by sampling regime	105
5.3	Benefits for QPE	106
5.4	Quality and requirements for operational LDR scans	108
5.4.1	Hardware quality requirements	109
5.4.2	Scan strategy	110
5.4.3	Quality control	113
5.4.4	Pre-processing to mitigate noise issues	114
5.5	Evaluation of Radarnet implementation	116
5.5.1	Case study evaluation	116
5.5.2	Gauge-radar statistical trials	119
5.6	Conclusions	119
6	Introducing a VPR correction for non-bright band precipitation	122
6.1	Introduction	122
6.2	Constructing profile shapes	125
6.2.1	Existing literature	125
6.2.2	RHI profile observations	127
6.3	Simulation studies using high resolution VPRs	130
6.3.1	Profile selection	130
6.3.2	Quality of fit: an illustrative comparison	133
6.3.3	Further simulation statistics	133
6.3.4	Rain rate scatterplots	136
6.3.5	Distributions of error with range	139
6.4	Case study evaluation with LDR-based classification	139
6.4.1	Description of case	140
6.4.2	Results and discussion	140
6.5	Conclusions	143
7	Improving operational correction for bright band VPRs	146
7.1	Introduction	146
7.2	Proposed change to the Kitchen et al. (1994) stratiform profile	147

7.3	Evidence of residual bias in stratiform QPEs	149
7.3.1	Selecting cases from the radar archive	149
7.3.2	Choosing suitable radars	150
7.3.3	Discussion	151
7.4	Investigating offset values	153
7.5	Impact of a 2 dB offset for real time rainfall estimation	154
7.5.1	Reducing long range biases in stratiform QPEs	154
7.5.2	Real time trialling over March 2017	157
7.6	Conclusions	163
8	Summary and outlook	165
8.1	Improving local classification and correction for the vertical reflectivity profile	165
8.2	Future work	167
8.2.1	Correcting for the VPR in non-bright band conditions	167
8.2.2	Extending LDR findings to operational systems	168
8.2.3	The question of scale: revisited	169
	Appendices	181
A	UK dual polarisation network parameters	182
B	Definition of statistics	183
C	LDR calibration factors for the Wardon Hill VPR dataset	184
D	NCIC temperature data for the winter of 2013-14	187

Chapter 1

Introduction

1.1 Background

The use of radar for monitoring and quantification of meteorological phenomena has become globally established over the past several decades. The term “radar”, originally an acronym RADAR, stands for “radio detection and ranging”, and dates from the use of radio frequency installations for military surveillance in the 1940s. Since then the details of radar hardware and frequency have diversified, and have come to fulfil a range of operational and research functions within the meteorological community.

The Met Office weather radar network began its development in 1976 with the installation of a 5.4 GHz (5.6 cm wavelength) Plessey radar at Hameldon Hill, in Lancashire (Kitchen and Illingworth, 2011). The initial aim of the installation was to monitor real time rainfall to inform preparations to mitigate flood impacts, such as occurred in 1952-1953 in the South and West of England. By 1985 the network had expanded to four radars, to 12 within the next decade, and currently stands at 15. UK radar products now also incorporate data from two non-Met Office radars in Ireland and one in the Channel Islands.

In the 33 years since the official launch of the network in 1985, substantial developments have been made both to the science of rainfall estimation, and to the uses of radar analyses. Algorithms to identify and remove non-meteorological echoes, and processing of meteorological artefacts such as the radar bright band (section 1.5.3), have led to steadily increasing reliability and corresponding demands on the accuracy of radar precipitation estimates. Met Office radar products are currently used by operational meteorologists to inform and advise, and by the Met Office / Environment Agency collaborative Flood Forecasting Centre in high impact situations. Rain rate composites are also assimilated into the Met Office numerical weather prediction (NWP) model via a latent heat nudging

scheme (described in Jones and Macpherson, 1997), as well as providing a strong component of the Short Term Ensemble Prediction System (STEPS) rainfall extrapolation nowcast (Bowler et al., 2006). As well as their scientific applications, real time radar products are highly visible to the general public, through publication on the external website and availability as part of the Met Office mobile app.

Radar quantitative precipitation estimates (QPEs) suffer from substantial uncertainties from a range of different sources (Villarini and Krajewski, 2010). Common problems in rainfall estimation include remote sensing issues, such as attenuation and contamination from non-meteorological echoes, as well as more complex inaccuracies arising from the microphysical relationships between the radar measurement and liquid water content. For hydrological users, these errors can have significant detrimental impacts to downstream products, such as runoff, pluvial and fluvial flood forecasting (Berne and Krajewski, 2013). The impacts are severe enough that some modellers prefer to ingest point rain gauge accumulations instead of a gridded radar product, despite the benefits in spatial coverage and grid representativity that radar provides. There is therefore strong motivation for continuing research and development towards improving the accuracy of radar QPEs.

With the proliferation of radar networks around the world over recent decades, radar systems have also undergone significant developments. One of the most important of these breakthroughs has been the introduction of dual polarisation: the ability to transmit and receive independently in the horizontal and vertical polarisations. This technology provides the capability to measure the shape of hydrometeors rather than their size, which allows more complex inferences to be made as to the type and quantity of precipitation falling at any given time.

While research into the uses of dual polarisation data has been active since the mid-1970s, over the past 10-15 years dual polarisation hardware has developed to a high enough quality and reliability to be used operationally. This has prompted many national meteorological services to upgrade their radar networks to dual polarisation capability (eg Kumjian, 2013; Figueras I Ventura and Tabary, 2013; Helmert et al., 2014; Gabella et al., 2016). The recent culmination of several upgrade projects has shifted the focus of dual polarisation research into the operational sphere, prompting the development of a range of algorithms for real time radar quality control, corrections and QPE.

The UK national radar network, including the Channel Islands radar, has recently undergone a complete upgrade to dual polarisation. The Met Office Weather Radar Network Renewal project (WRNR), which began in 2011 and was completed in December 2017, replaced and renovated existing hardware at these 16 UK radars (but not the two in Ireland). The new radars were designed and built in-house using commercial off-the-shelf components (Darlington et al., 2004, 2016), and are capable of high quality dual

polarisation measurements comparable to those of leading research facilities (appendix A). This upgrade, combined with the quality of the new data available, provides unique opportunities for research and development in the context of the UK climate and existing radar processing algorithms.

This thesis investigates improvements to radar rainfall estimation in the UK through the use of a specific dual polarisation parameter, the linear depolarisation ratio (LDR), to inform corrections for the vertical profile of radar reflectivity (VPR). Understanding the variability in radar-measured reflectivity with height, and how this compares to intrinsic reflectivity at the surface, is a crucial step in obtaining accurate QPEs (section 1.5). In this introductory chapter the principles of meteorological radar are described, from the initial measurement to the estimation of surface rain rates. The nature of the radar measurement and its meteorological interpretation are covered in section 1.2, followed by an overview of operational networks and scan strategies in section 1.3. Section 1.4 introduces the concept of dual polarisation, which in recent years has become the new standard for operational radar networks. The sequence of processes by which UK radar reflectivities are converted into rainfall estimates, including determination and correction for the VPR, is described in section 1.5. Having established this context, the full aims and an outline of this thesis are presented in section 1.6.

1.2 Principles of meteorological radar

Weather radar works on the principle of echo location. A transmitter first sends out a pulse of electromagnetic energy, followed by a passive phase in which the receiver “listens” for echoes. These microwave pulses are transmitted via a parabolic antenna whose properties determine the radar beam pattern and gain. In “receive” mode the antenna assembly acts in reverse: power returning is focused into the antenna and measured by the receiver. Radar receivers must be robust in detecting power returns across the 8-9 orders of magnitude spanned by meteorological echoes, and must also be extremely sensitive, given the significantly decreased power of reflected echoes in comparison to the transmitted pulse. A physical “T/R” switch operates to transition the radar between “transmit” and “receive” modes, which protects the receiver assembly from the higher energy transmission.

1.2.1 Rayleigh scattering, reflectivity and the radar equation

The basic meteorological measurement derived by weather radar is known as “reflectivity”. This is calculated directly from the received power, and is a function of properties of meteorological targets within the radar pulse volume.

The power received at the radar from a population of hydrometeors is given by:

$$p_r = \frac{A \sum \sigma_i}{r^2} \quad (1.1)$$

$$A = \frac{p_t g^2 \lambda^2 \Theta^2 l}{1024 \pi^2 \ln(2)} \quad (1.2)$$

Here p_t is transmitted power, g is the antenna gain, l the pulse length, λ the wavelength, and Θ the angular beam width; which makes A a constant for a given radar (Bringi and Chandrasekar, 2001). These fixed beam properties are discussed in section 1.3.2. Variable factors in this equation are range from the radar, r , and the scattering cross-section of each hydrometeor, σ_i .

The scattering cross-section of a radar target depends on both the size of the target and the radar wavelength, which determines the scattering regime. The wavelength of meteorological radars is chosen so that the hydrometeors of interest are around an order of magnitude smaller than the incident wavelength, and are therefore sampled within the Rayleigh regime. In this regime, the electromagnetic pulse \vec{E}_i transmitted by the radar induces an oscillating dipole in the target hydrometeors, and the associated movement of electrons generates its own field \vec{E}_s , which is scattered both forward and backward along the transmission path. The amplitude of the backscattered field from a dielectric sphere relates to properties of the incident radiation and the target as:

$$\left| \vec{E}_s \right| = \frac{k_0^2}{4\pi r} \frac{m^2 - 1}{m^2 + 2} \times 4\pi D^3 \left| \vec{E}_i \right| \quad (1.3)$$

where k_0 is the wavenumber of the incident radiation (which is inversely proportional to its wavelength λ), m is the complex refractive index and D is the diameter of the spherical target (Bringi and Chandrasekar, 2001, equations 1.32a and b). The backscattered power ($\propto \left| \vec{E}_s \right|^2$) from a point target in the Rayleigh regime is therefore proportional to λ^{-4} and the sixth power of the target diameter D^6 .

Following from equation 1.3, the scattering cross-section of a Rayleigh point target is defined as:

$$\sigma_i = \frac{\pi^5 |\kappa|^2 D_i^6}{\lambda^4} \quad (1.4)$$

where dielectric factor κ relates to the complex refractive index m as:

$$|\kappa| = \left| \frac{m^2 - 1}{m^2 + 2} \right| \quad (1.5)$$

Radar reflectivity, which is the sum of diameters to the sixth power, can then be defined in terms of the sum of scattering cross-sections from hydrometeors i of different diameters

per unit volume:

$$Z = \sum D_i^6 = \frac{\lambda^4}{\pi^5 |\kappa|^2} \sum \sigma_i \quad (1.6)$$

The definition of reflectivity as a sum of diameters is valid only for liquid rain drops with sphere-equivalent diameters D_i ; however the definition in terms of scattering cross section is general. Combining equations 1.1 and 1.6 leads trivially to the radar equation: $Z = Cp_r r^2$. The conversion from power to reflectivity is typically performed at the radar site, using the precalculated radar constant C and the range r calculated from the echo arrival time.

The relationship between reflectivity and the volume of water sampled by a radar pulse (through the D^6 term) depends on the dielectric factor κ , which varies with precipitation phase. Reflectivity (equation 1.6) is defined with respect to the value for liquid water: $|\kappa_w|^2 = 0.93$. However dielectric factor values for frozen hydrometeor species are significantly lower, varying in proportion to snow density up to a maximum of $|\kappa_i|^2 = 0.176$ for solid ice (Sauvageot, 1992, chapter 2, pg 97-98). This means that the reflectivity of snow is at least 5 times lower than the reflectivity of the equivalent volume of liquid rain drops, which impacts the drop diameters and water volume that can be inferred from a reflectivity measurement. The behaviour of κ and its impact on radar reflectivities is discussed in more detail in section 2.2.2.

For meteorological purposes reflectivity is usually expressed as an integral rather than a sum. Defining the drop size distribution $N(D)dD$ as the number of hydrometeors per unit volume with diameters between D and $D + dD$ yields the most widely used form of the definition:

$$Z = |\kappa_w|^{-2} \int_0^\infty N(D) |\kappa|^2 D^6 dD \quad (1.7)$$

This linear Z is in units of $\text{mm}^6 \text{m}^{-3}$, with D in mm and $N(D)$ in $\text{mm}^{-1} \text{m}^{-3}$.

Given that meteorological reflectivities span several orders of magnitude, reflectivity values are often given in logarithmic dBZ units:

$$Z_{dBZ} = 10 \log(Z) \quad (1.8)$$

Reflectivities in rain typically fall within the range 18-55 dBZ, with reflectivities greater than 55 dBZ often indicating the presence of hail in convective situations (Bringi and Chandrasekar, 2001).

1.2.2 Reflectivity and rain rate

The main function of meteorological radar in an operational environment is to provide real time estimates of rain rate at the ground surface. Rain rate R is the volume of

liquid water reaching the ground per unit time, and like reflectivity (equation 1.7) can be defined in terms of an integral over rain drops of different diameter:

$$R = \frac{\pi}{6} \int_0^{\infty} N(D)v_t(D)D^3 dD \quad (1.9)$$

where v_t is the terminal fall velocity of each drop. The fall velocity is generally approximated by a power law $v_t(D) = \alpha D^\beta$, where β lies in the range 0.6-0.67 for liquid rain drops (Bringi and Chandrasekar, 2001). Radar rain rates are usually expressed in millimetres per hour (mm h^{-1}).

The relationship between reflectivity and rain rate (ZR relation) depends on the form of the drop size distribution. Marshall and Palmer (1948) use a rain rate dependent exponential to describe a measured drop size distribution. The slope of this exponential can be expressed as a function of the median drop diameter, D_0 (Bringi and Chandrasekar, 2001), so that:

$$N(D) = N_0 \exp\left(-3.67 \frac{D}{D_0}\right) \equiv N_0 e^{-x} \quad (1.10)$$

Assuming a rain drop fall velocity proportional to $D^{0.67}$, it follows that:

$$Z \propto \int_0^{\infty} N_0 D^6 \exp\left(-3.67 \frac{D}{D_0}\right) dD \propto N_0 D_0^7 \int_0^{\infty} x^6 e^{-x} dx \quad (1.11)$$

$$R \propto \int_0^{\infty} N_0 D^{3.67} \exp\left(-3.67 \frac{D}{D_0}\right) dD \propto N_0 D_0^{4.67} \int_0^{\infty} x^{3.67} e^{-x} dx \quad (1.12)$$

$$Z \propto N_0^{(1-7/4.67)} R^{7/4.67} \approx N_0^{-0.5} R^{1.5} \quad (1.13)$$

This specific example illustrates the more general form of the ZR relation, which is expressed as a power law:

$$Z = aR^b \quad (1.14)$$

where a is inversely proportional to the square root of the intercept parameter N_0 . Such fixed ZR relationships are used throughout the operational community to estimate rain rates from radar networks.

In reality $N(D)$ is not always exponential, nor is N_0 a constant, but varies with rain type and climatology. Optimal coefficients for equation 1.14 can be determined empirically, using disdrometer or other independent measurements in conjunction with radar. These methods typically find the exponent b in the range 1.4-1.6, but estimates of a vary much more widely due to its dependence on drop concentration (Battan, 1973). The form suggested by Marshall and Palmer (1948) for stratiform precipitation is $Z = 220R^{1.6}$, which is frequently approximated as $Z = 200R^{1.6}$ (the ‘‘Marshall-Palmer relation’’).

1.3 Operational radar configurations

1.3.1 Scanning meteorological radar

Operational radars collect the majority of their data in scanning mode. For scanning meteorological radar the antenna rotates in azimuth while transmitting and receiving, building up a circular “plan position indicator” (PPI) image of targets surrounding the radar (eg figure 1.1). A PPI scan is typically taken at a fixed non-zero elevation angle above the horizontal, and has a maximum range of between 60 and 300 km. Due to Earth curvature and the refractive properties of the atmosphere, the height of the radar beam above the ground increases with range according to the four-thirds Earth approximation:

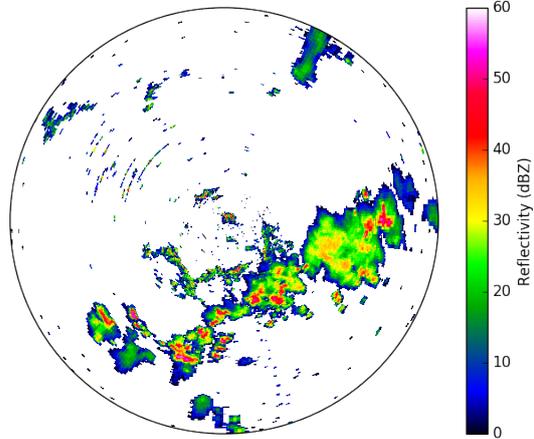


Figure 1.1: Example of a reflectivity PPI at 0.5° elevation from the Channel Islands radar, 14:59 UTC on 28th May 2018. Maximum range of 255 km.

$$H = H_{rad} + \sqrt{r^2 + R'^2 + 2rR'\sin\theta} \quad (1.15)$$

$$\text{where } R' = \frac{4}{3}R_E \text{ and } R_E = 6374 \text{ km (Earth radius)} \quad (1.16)$$

where H_{rad} is the antenna height, r the range (in km) from the radar, and θ the scan elevation angle (Doviak and Zrnić, 1993). This means for example that at a range of 250 km, a radar PPI measurement at 0.5° elevation would be centred around 6 km above the ground (figure 1.2). The radar beam also broadens with range, with a typical operational beam width of 1° (eg Harrison et al., 2012; Figueras I Ventura and Tabary, 2013; Helmert et al., 2014; Koistinen and Pohjola, 2014) spanning an azimuthal and vertical distance of 1 km at 50 km range, 2 km at 100 km, and more than 5 km at the maximum 255 km range of the Met Office radars (as illustrated in figure 1.2).

Operational weather radars typically operate a scan strategy including several PPIs at different elevations above the horizontal. Radar “volumes”, consisting of a series of PPIs taken within a fixed time period, are repeated at frequencies of 5 to 15 minutes depending on the update frequency required of radar precipitation estimates. The use of volumes provides additional vertical coverage in PPI scan mode, and allows higher elevation scans to be used to replace lower elevation data contaminated by non-meteorological echoes (section 1.5.1). Radar volumes can also be used to generate 3D products such as hydrometeor type, for vertical reflectivity profile estimation, and for the construction of 2D

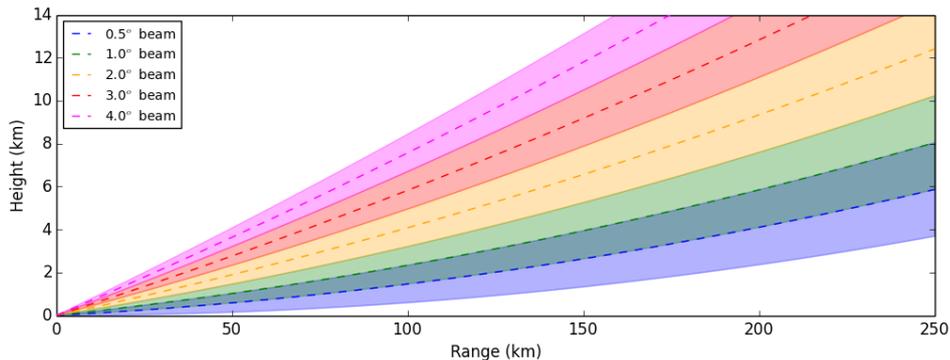


Figure 1.2: Schematic showing beam heights and broadening of the five PPI elevations used in the Met Office QPE scan strategy.

gridded products at fixed altitudes (constant altitude PPIs, or CAPPIs).

The five beam QPE scan strategy used by Met Office radars is shown in figure 1.2. This is a compromise well suited to the precipitation systems most frequently observed in the UK. In lower latitude climates where convective systems can reach greater depths, and at centres where volume reflectivities may be used in numerical weather prediction (NWP), higher elevation scans may also be included. There is a trade-off between including more scans to observe the whole precipitation column and completing the radar volume scan within the time available.

1.3.2 Properties of the radar beam

The spatial resolution of a radar and the limits of its sensitivity to precipitation depend on properties of the hardware and the radar beam. The radar constant C , which relates the radar reflectivity to the received power, is influenced by several factors including the wavelength λ , pulse length l , and the angular width of the beam Θ .

Wavelength and dish diameter

To maximise both sensitivity and spatial resolution (in azimuth), it is desirable for the power transmitted by a radar to be focused into a narrow beam. Assuming a Gaussian beam power profile, the half power beam width Θ_{hp} achieved by a given radar is a function of the ratio between the radar wavelength λ and the diameter of the parabolic reflector dish d_A :

$$\Theta_{hp} = \frac{1.27\lambda}{d_A} \quad (1.17)$$

where Θ_{hp} in radians is twice the off-axis angle at which the radar beam power reduces to half of the peak power (Doviak and Zrnić, 1993).

Bandwidth	Frequency (GHz)	Wavelength (cm)	Dish diameter (m) for a 1° beam	Application
S-band	2-4	7.5-15	5.5-11.0	Rainfall
C-band	4-8	3.75-7.5	2.7-5.5	
X-band	8-12	2.5-3.75	1.8-2.7	
Ku-band	12-18	1.67-2.5	1.2-1.8	Cloud
K-band	18-24	1.11-1.67	0.8-1.2	
Ka-band	24-40	0.75-1.11	0.5-0.8	

Table 1.1: The main frequency bandwidths used by rainfall (S-X) and cloud (Ku-Ka) radars, for operational and research applications.

Operational weather radars within the UK and Europe typically operate with an angular beam width of around 1° (eg Kitchen et al., 1994; Tabary, 2007), which provides the kilometre-scale spatial resolution with which rainfall estimates have historically been required by the modelling and verification communities. The main radar wavelength bands used for this purpose are listed in table 1.1. A 1° beam width at a C-band wavelength of 5.3 cm (as used by the UK radar network) requires a parabolic reflector dish 3.86 m in diameter. At S-band, due to the longer wavelength, a 7-8 m dish is required. Conversely at X-band, this spatial resolution can be achieved with a dish 2 m in diameter. For these reasons, X-band radars are a popular choice for field studies, as higher spatial resolution can be achieved with a smaller and more portable instrument. Longer wavelengths tend to be chosen for ground based operational radars, due to the wider spatial coverage and lower sensitivity to meteorological attenuation (section 1.5.2).

Pulse length, averaging and pulse repetition frequency

Each pulse transmitted by a weather radar has a duration T , which relates to the spatial pulse length l by the speed of light ($l = cT$). The shorter the pulse duration, the finer the radial resolution of the radar image. For two targets along the same azimuth to be resolved, the end of the transmitted pulse must reach the first target before the start of the echo from the second target arrives back at the first. The minimum target separation at which this is possible defines the spatial resolution, which is equal to half of the spatial pulse length.

Shorter pulse length also allows for a higher pulse repetition frequency (PRF), defined as the number of pulses transmitted per second (in Hz). A higher sampling rate can reduce the error on derived reflectivities. However, since an echo can only be attributed with certainty to a given pulse if it is received before the next pulse is transmitted, a higher PRF reduces the maximum unambiguous range of the radar. Given a time between successive pulse initiations (in seconds) of $\tau = 1/\text{PRF}$, the maximum unambiguous range is $c\tau/2$.

The reflectivities returned from a radar site are not calculated from a single pulse, but from the average received power $\overline{p_r}$ from several transmitted pulses. The phase of the echo from a meteorological target is random, meaning that the standard deviation of the pulse-to-pulse power received from a target is equal to the mean power, assuming that the pulses are independent. Averaging over M independent pulses reduces the standard deviation on the derived measurement by a factor of \sqrt{M} , which improves the accuracy of the measurement.

The number of independent pulses is not the same as the total number of pulses averaged to calculate reflectivity, due to the temporal correlation between successive measurements. The correlation between pulses separated by lag m varies with wavelength λ and PRF:

$$\rho(m, \tau) = \exp \left[- \left(\frac{4\pi m \sigma_v \tau}{\lambda} \right) \right] \quad (1.18)$$

where σ_v is the velocity spectrum width (Doviak and Zrnić, 1993). This can be used to calculate the number of independent pulses M contributing to the average of a larger, correlated sample M_t .

For Met Office radars, standard QPE reflectivity measurements are collected at a PRF of 300 Hz with a spatial pulse length of 300 m. Two radial bins are averaged for each measurement, meaning that the data are returned at a radial resolution of 600 m. Given an azimuthal resolution of 1° and a scan rotation rate of 8.4°s^{-1} , this means that each reflectivity measurement is calculated using the average power from $M_t = 37$ pulses. With a Doppler spectrum width of 1 m s^{-1} this is equivalent to roughly $M = 16$ independent pulses, or a four-fold reduction in noise on the averaged reflectivity measurement. The PRF of 300 Hz corresponds to a maximum unambiguous range of 500 km; but due to the height of the beam (equation 1.15) only reflectivities within 255 km of the radar are used operationally.

1.3.3 Radar networks

For monitoring rainfall on a national level, most meteorological services operate a network of radars to maximise spatial coverage and service reliability. Data from these networks are typically combined to synthesise a “composite” product, which is composed of maximum, highest quality, or weighted average rain rate estimates from all contributing radars. The number of radars in a network is related to land area: the smallest European networks consist of only two or three radars, with networks of 15 to 20 radars being typical of the larger European nations.

The Met Office currently operates 15 radars across mainland UK, and receives data

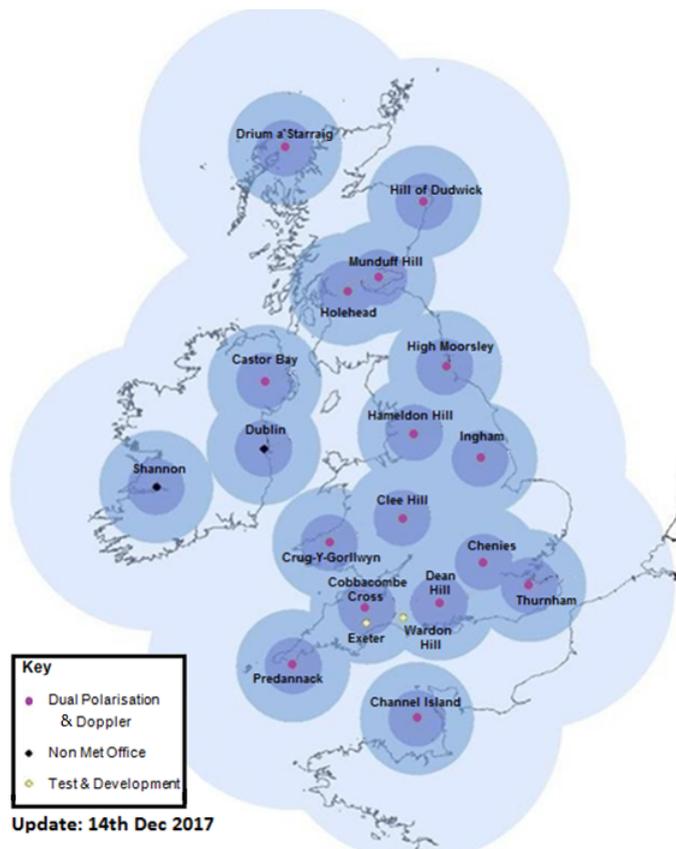


Figure 1.3: Status of the UK national radar network after completion of the Weather Radar Network Renewal project (WRNR). Different shading (dark to light) indicates coverage within 50 and 100 km respectively of the nearest radar.

from two more in Ireland and one in the Channel Islands (figure 1.3). This network provides multiple coverage of most land areas, increasing the resilience of the network to outages and maintenance schedules. The Met Office QPE composite is generated using the highest quality rain rate estimate at each point, taking account of both geometrical factors (range and height of the measurement) and meteorological conditions. The 1 km gridded rain rate composite underpins the majority of downstream applications in the UK.

1.4 Dual polarisation radar

Section 1.2 introduced the general relationships between rainfall rate, reflectivity, and raw received power for meteorological radar. A simple reflectivity measurement for rain rate estimation can be made using a microwave pulse transmitted in any polarisation. However, developments in radar hardware over the past several decades have made it possible to transmit and receive power independently in two orthogonal polarisation

channels. Dual polarisation parameters, derived by comparing both the power and relative phase of the polarised returns, can provide significant additional information on the properties of meteorological echoes.

1.4.1 Dual polarisation transmission modes

Most modern dual polarisation systems operate using a linear polarisation basis with horizontal (H) and vertical (V) channels. The preferred method for operational radars is to transmit equal power simultaneously in both the H and V channels: a 45° polarised or slant elliptical transmission. This is referred to in the literature as “SHV mode” or “hybrid transmission”. In addition to reflectivity (usually Z_h), the following SHV mode dual polarisation parameters can be calculated:

- Differential reflectivity (Z_{DR})
- Copolar correlation coefficient (ρ_{hv})
- Differential phase shift (Φ_{dp})

The information provided by dual polarisation parameters, in general, relates to the shape of hydrometeors or other reflectors in the radar pulse volume. Section 1.4.2 gives an overview of the dual polarisation parameters listed above, with definitions and a brief description of the information they contain. This information is included for context, and is not used further in this thesis. The main focus for this thesis is section 1.4.3, which describes in detail the measurement and interpretation of the linear depolarisation ratio.

1.4.2 Interpretation and applications of SHV mode parameters

Differential reflectivity

Differential reflectivity is defined as:

$$Z_{DR} = 10 \log \left(\frac{Z_h}{Z_v} \right) \quad (1.19)$$

where Z_h and Z_v are the reflectivities calculated respectively from the H and V received power. For meteorological echoes, Z_{DR} provides a measure of the average axis ratio. For liquid drops axis ratio is linked to diameter, with larger diameters producing more oblate spheroids. This means that Z_{DR} can be used in combination with reflectivity Z_h to estimate parameters of the drop size distribution $N(D)$ (Brandes et al., 2004), which in turn determine the optimal ZR coefficients for rainfall estimation (section 1.2.2). This is discussed further in section 1.5.4.

Since the axis ratio of rain drops is known to increase with increasing size, Z_{DR} measurements in rain should increase monotonically with reflectivity Z_h . Meteorological values of Z_{DR} range from close to zero for light rainfall (low Z_h), to up to 2 dB for very heavy rain (high Z_h), and can be in excess of 4 dB for wet hail and melting ice, where liquid water forms a torus around the melting particle that increases its axis ratio (Rasmussen and Heymsfield, 1987). Z_{DR} responds to axis ratio differently for rain and snow, due to the difference in dielectric factor. A rain drop with an axis ratio of 2 would have a very high Z_{DR} of 6 dB; but a dry ice pellet or graupel with the same axis ratio has Z_{DR} of 2-3 dB, and a dry snowflake has Z_{DR} of less than 1 dB (Doviak and Zrnić, 1993). High reflectivities colocated with regions of near-zero Z_{DR} correspond to tumbling hail or graupel, where hydrometeors are irregular but the average axis ratio due to non-preferential orientation is equal to 1.

It is possible for measured Z_{DR} to be lower than expected behind strong reflectivity echoes. This is due to “differential attenuation”, usually through heavy rain, where the horizontal component of the radar pulse is more strongly attenuated than the vertical. In extreme cases this can result in negative Z_{DR} measurements, despite the positive or neutral aspect ratio of hydrometeors in the radar pulse volume (eg figure 1.4b).

Copolar correlation coefficient

Radar parameters for a single range bin are not derived from a single pulse, but typically are averaged from up to 30 to 40 pulses (section 1.3.2). The relationship between power returns from individual pulses can be used to derive measures of signal variability in a single polarisation (eg the clutter indicator, section 1.5.1), as well as providing information as to the range of dual polarisation properties within the radar pulse volume.

The copolar correlation coefficient ρ_{hv} is the correlation between the Z_h and Z_v time-series for separate pulses within the same range bin. The back-scattering matrix for a hydrometeor in the radar pulse volume is defined:

$$\begin{bmatrix} S_{hh} & S_{hv} \\ S_{vh} & S_{vv} \end{bmatrix} = e^{j\delta_{hh}} \begin{bmatrix} |S_{hh}| & |S_{hv}|e^{j(\delta_{hv}-\delta_{hh})} \\ |S_{vh}|e^{j(\delta_{vh}-\delta_{hh})} & |S_{vv}|e^{j(\delta_{vv}-\delta_{hh})} \end{bmatrix} \quad (1.20)$$

where δ_x elements describe the phase of the backscattered signal (relative to the incident phase δ_{hh}) and the magnitude elements $|S_x|$ are proportional to the square root of the back-scattered power (and thus the square root of reflectivity) (Bringi and Chandrasekar, 2001). Using this notation, ρ_{hv} is defined:

$$\rho_{hv} = \frac{|\langle S_{hh}S_{vv}^* \rangle|}{\sqrt{\langle |S_{hh}|^2 \rangle \langle |S_{vv}|^2 \rangle}} \quad (1.21)$$

The behaviour of ρ_{hv} is dependent on both the quality of polarisation separation and the target hydrometeors. If the sampling volume contains a population of hydrometeors of similar size and shape, then the backscattered phase will tend to be consistent from pulse to pulse, so the correlation between timeseries will be high. These conditions are typical of rain. Where the quality of polarisation separation in the radar hardware is good, ρ_{hv} measurements in rain are typically greater than 0.99.

Volumes containing a mixture of hydrometeor shapes and types give much less correlated returns. ρ_{hv} values of around 0.8 at low levels can indicate the presence of hail in heavy convective rainfall. Melting snowflakes are also associated with a reduction in ρ_{hv} due to their variety of axis ratios and canting angles. This property is often used to locate the melting layer in radar PPIs (eg Tabary et al., 2006; Boodoo et al., 2010; Giangrande et al., 2008; Kalogiros et al., 2013), the significance of which is introduced in section 1.5.3 and discussed in detail in chapter 2. ρ_{hv} can also be used to identify non-meteorological echoes, which have correlation values typically below 0.5 with a distribution that is well-separated from that of precipitation.

Differential phase shift and specific differential phase

Differential phase shift (Φ_{dp}) measures the phase difference in degrees between echoes returned in the H and V channels. The value of Φ_{dp} quantifies the difference in the extent to which radar pulses are slowed by the change in refractive index (from air to precipitation) in the two polarisation channels.

The amount of phase shift is related to the amount of precipitation encountered, in terms of scattering cross-section. When a radar pulse passes through a region of oblate rain drops, the horizontal component is phase shifted more than the vertical component, which is defined as positive Φ_{dp} (eg figure 1.4d). Similar to Z_{DR} , drops with more oblate shapes produce a higher positive Φ_{dp} . However unlike Z_{DR} , Φ_{dp} is a difference not a ratio, and is therefore also sensitive to the total drop concentration in the radar pulse volume.

Since Φ_{dp} is cumulative an additional useful parameter, “specific differential phase” (KDP), can be defined as the gradient of Φ_{dp} along a radial. KDP in rain provides an indirect measure of drop shape, but is also sensitive to the quantity of water in the radar pulse volume. This means that under certain conditions KDP can be used as an alternative to reflectivity in rainfall estimation (eg Brandes et al., 2003, see also section 1.5.4).

Φ_{dp} is a noisy measurement, with estimated random errors of at least $\pm 3^\circ$ even after pulse averaging (section 1.3.2). A major problem for KDP calculation is that due to

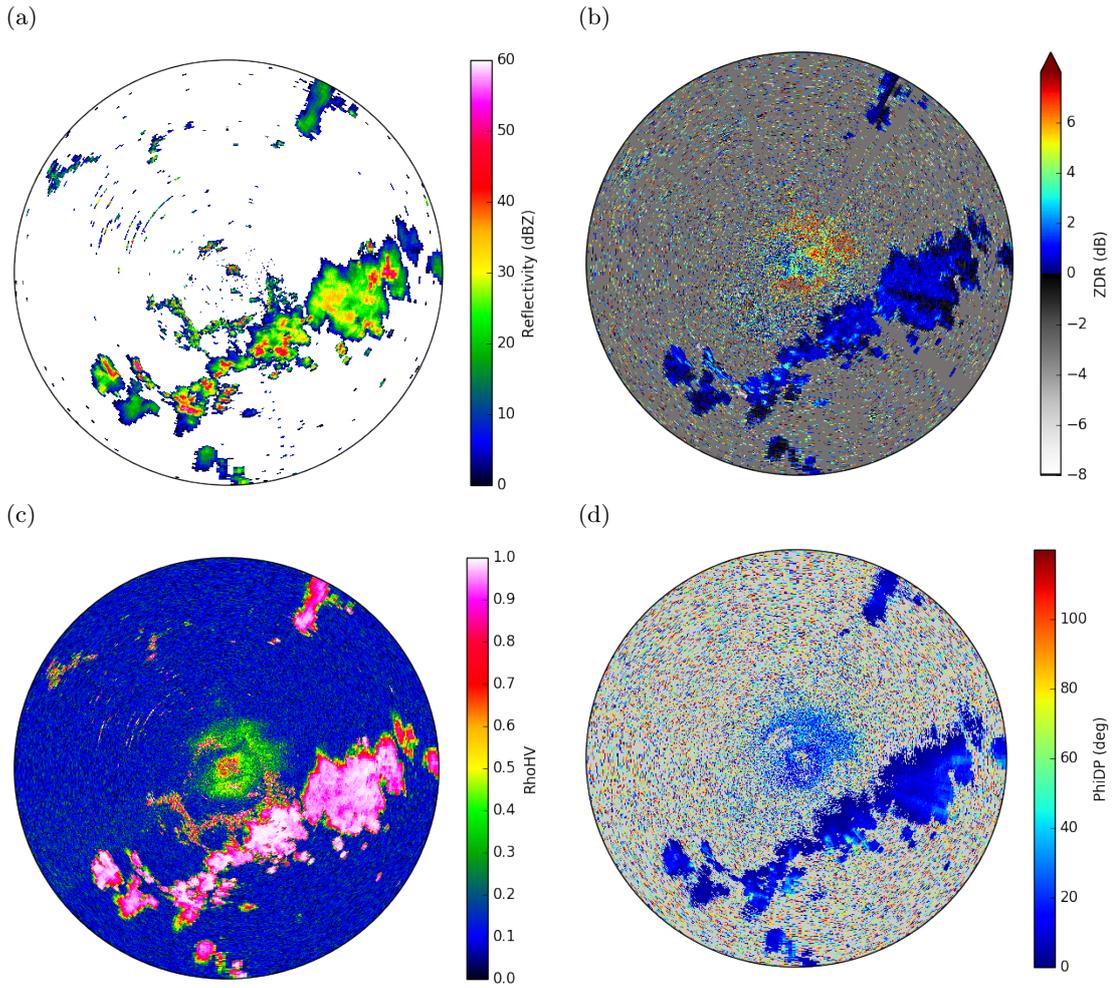


Figure 1.4: Example of different dual polarisation measurements from a rainfall event observed by the Channel Islands radar at 14:59 UTC on 28th May 2018. Top left: reflectivity Z (copied from figure 1.1); top right: differential reflectivity Z_{DR} ; bottom left: ρ_{hv} ; bottom right: Φ_{dp} .

these effects, Φ_{dp} values do not always increase monotonically with range. The noisiness of the Φ_{dp} range profile means that it is not always possible to extract the underlying monotonically increasing function, which means that accurate KDP values cannot be derived for every range gate or in all conditions. At the Met Office, KDP is calculated only in regions of rapidly increasing phase shift where the resulting value would be greater than 16° km^{-1} (corresponding to rain rates in excess of 10 mm h^{-1}).

Combining parameters

In addition to their independent interpretation, the information provided by SHV mode parameters can be used in tandem to improve and extend beyond radar precipitation analyses. The main combined applications of dual polarisation are in the identification

of non-meteorological echoes and the classification of different hydrometeor types (eg Rico-Ramirez and Cluckie, 2008; Park et al., 2009; Wen et al., 2016). Classification of rain is based on the relationship between Z_h and Z_{DR} in regions of high ρ_{hv} ; while lower correlations (0.8-0.9), high reflectivities, and zero or negative Z_{DR} at low levels can indicate the presence of hail. The melting layer in stratiform precipitation is identifiable by reduced ρ_{hv} , high reflectivity and high Z_{DR} (eg Tabary et al., 2006). Non-meteorological echoes can be identified by their rough “texture” (high spatial variability) in Z_{DR} and Φ_{dp} , and by very low correlation coefficients (Rico-Ramirez and Cluckie, 2008).

The dual polarisation characteristics of meteorological and non-meteorological echoes can be seen in the example in figure 1.4, which shows a rainfall event observed from the Channel Islands radar in May 2018. The rain in the radar image is characterised by high ρ_{hv} and by smooth Z_{DR} and ϕ_{dp} , whereas regions of high reflectivity and ρ_{hv} to the North-West of the radar are identifiable as non-meteorological by their noisy signatures in Z_{DR} and ϕ_{dp} . Areas of negative Z_{DR} and high ϕ_{dp} to the South and East of the radar indicate strong attenuation, which is consistent with the high reflectivities observed (panel (a)). The lower (≈ 0.9) values of ρ_{hv} in the region to the East of the radar suggest the presence of strong convection and hail.

Although quality control can be applied throughout the radar domain, more quantitative applications of dual polarisation data, such as corrections for attenuation and rainfall estimation, are as yet limited to radar measurements in rain, where the preferential orientation of liquid drops leads to quantifiable relationships between drop shape and volume that provide useful additional information on liquid water content (Herzogh and Jameson, 1992).

1.4.3 Single polarisation transmission: the linear depolarisation ratio

In addition to simultaneous transmission, further information on drop properties can be derived from transmission in a single polarisation. The linear depolarisation ratio (LDR) is measured by transmitting plane polarised horizontal pulses and receiving in both polarisations. LDR is defined as the fraction of the plane polarised signal that is returned in the opposite polarisation:

$$\text{LDR} = 10\log\left(\frac{Z_{vh}}{Z_{hh}}\right) \quad (1.22)$$

The mechanism by which a hydrometeorological target may depolarise the incident radar pulse is illustrated in figure 1.5. Figure 1.5a shows scattering from a horizontally oriented drop which does not depolarise. Figure 1.5b illustrates how when a drop is canted, the dipole induced along the major and minor axes generates a scattered electric field which

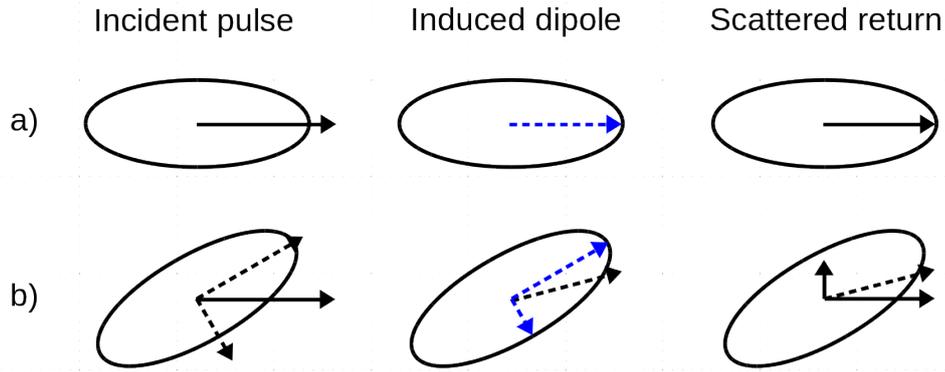


Figure 1.5: Depolarisation of a horizontally polarised incident radar pulse by a hydrometeorological target. Solid lines indicate the incident and scattered returns in the horizontal and vertical polarisations, while dashed lines show how these resolve along the axes of the target hydrometeor. a) An oblate, horizontally oriented drop which does not depolarise. b) An oblate, canted drop which generates a small vertical component from the horizontally polarised transmission.

is not aligned with the incident polarisation. It is clear from this schematic that both aspect ratio and orientation contribute to LDR, with higher axis ratios causing more depolarisation for a given canting angle.

The unique advantage of LDR as a parameter is this responsiveness to the orientation of hydrometeors within the radar pulse volume (Illingworth, 2004). Rain and snow consist of populations of largely horizontally oriented particles, and are therefore not depolarising (figure 1.5a). The depolarised component in this case is not identically zero, but is limited by the cross-polar isolation of the dual polarisation system, so that the minimum LDR measured in light rain (where drops are near-spherical) is typically between -30 and -40 dB. However in the melting layer, the oscillations of large, partially melted snowflakes cause significant depolarisation (figure 1.5b). The sum of depolarised reflectivity components across the range of canting angles generates a characteristic melting layer peak of around -18 dB in LDR PPIs (Smyth and Illingworth, 1998; Illingworth and Thompson, 2011).

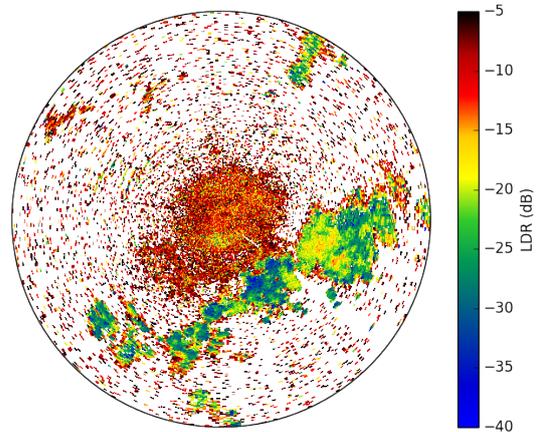


Figure 1.6: LDR PPI at 0.5° elevation from the Channel Islands radar at 14:58 UTC on 28th May 2018. The high, localised values to the East of the radar are consistent with hail.

Depolarisation is not unique to melting snowflakes. In convective conditions, tumbling wet hail and graupel also generate high values of LDR, due to their essentially random canting angle. Figure 1.6 shows the LDR scan through the event shown in figure 1.4. Low

LDR of around -35 dB to the South-East of the radar corresponds to areas of rain. High values to the East are consistent with the signatures of attenuation in Z_{DR} and ϕ_{dp} , and support the interpretation of reduced ρ_{hv} in this region as hail. The high and noisy LDR field surrounding the radar at short range is attributable to non-meteorological echoes.

LDR cannot be calculated from SHV mode measurements, nor other dual polarisation parameters from LDR mode data. Due to competition for time in a real time scan strategy, operational measurements of LDR are therefore uncommon. Unlike many commercial radar installations, Met Office radars are designed to allow the collection and interpretation of LDR mode data, in addition to the more typically collected SHV mode scans.

1.5 Quantitative precipitation estimation

The process of obtaining quantitative precipitation estimates (QPEs) from radar data varies between National Meteorological Services (NMSs), but broadly speaking can be described in four main stages. Initially, the raw reflectivity volumes are filtered to remove echoes from non-meteorological targets such as buildings, trees, and complex terrain. Measurements are then corrected for power loss effects, such as partial beam blockages and attenuation through heavy precipitation (section 1.5.2). This may be followed by an adjustment to account for the height of the beam above ground level. The final stage is the conversion of corrected reflectivities, and occasionally dual polarisation parameters, into estimates of precipitation rate at the ground.

This section gives an overview of the main QPE steps applied by the Met Office’s operational radar data processing system (Radarnet). This provides a context for the work of this thesis, and a reference for the additional processing performed as part of the analyses of chapters 5-7. While a full review of different NMS’s processing steps is not provided, reference is made where relevant to the main types of dual polarisation algorithms used within the operational community.

1.5.1 Identifying and removing non-meteorological echoes

Although weather radar is designed to observe and quantify precipitation, in reality echoes can be received from any number of sources. Close to the radar, a beam transmitted at low elevation may intercept targets such as buildings, trees, or complex terrain. Such stationary non-meteorological targets are known as “ground clutter”. Ground clutter echoes are often associated with very high reflectivities (at least 40-50 dBZ), which if misidentified can cause extreme overestimation or misleading detections where no precip-

itation is present (“false alarms”). Strong echoes may also arise from airborne targets, such as aircraft, which often cause contamination close to airports. Aircraft echoes are strong but extremely localised, tending to be restricted to one range bin across one or two radar azimuths. Weak distributed echoes may be caused by “biological clutter”: flocks of birds or swarming insects. These echoes tend to be detected only at the high sensitivities and relatively low sampling heights of measurements very close to the radar.

As well as aircraft, ground and biological clutter, radars near the coast may receive additional returns from the sea surface. This “sea clutter” has very different characteristics from ground clutter, since the beam is scattered from liquid water, so the reflectivity values and “texture” can be very similar to that of a precipitation field. Coastal radars may also detect large container ships, which have similar properties to aircraft echoes.

Finally, non-meteorological echoes can arise from the ground itself. Anomalous propagation, or “anaprop”, occurs in situations where the atmospheric refractivity profile differs from that expected from climatology. Equation 1.15 for beam height assumes a “standard atmosphere”: a climatological profile of temperature and relative humidity that defines refractivity and dictates the expected propagation of the radar beam. In reality, variations in these physical profiles can cause the beam to deviate from its expected path. This is primarily an issue in “super-refraction” conditions, where the beam is deflected more towards the ground than would be expected from the four-thirds Earth model. In anaprop cases, radar echoes from the ground and sea surface can be incorrectly interpreted as meteorological echoes at some height in the atmosphere.

The identification of non-meteorological reflectivities in Radarnet is performed on a scan-by-scan basis using one of two different methods, depending on whether or not dual polarisation parameters are available for that particular scan. Reasons for dual polarisation data being unavailable may be temporary, for example as a result of scheduled maintenance, or can be due to standing issues such as limited bandwidth or non-upgraded radar hardware. For these reasons Radarnet remains flexible to inputs from both single and dual polarisation scans.

For single polarisation data, a number of successive filters are applied to identify and remove non-meteorological signals. Echoes are classified hierarchically, starting with ground and biological clutter, “speckle” (strong echoes over a handful of radar pixels due to ships and aircraft), radio frequency interference, and anaprop. The single polarisation filters in Radarnet rely heavily on a “clutter indicator” (CI) measure of signal variability (Sugier et al., 2002). Since ground clutter is stationary, echoes from clutter fluctuate very little from pulse to pulse, whereas precipitation echoes are constantly fluctuating. The initial CI-based clutter filter is followed by filters based on satellite cloud masks and precipitation climatologies built up from individual radars (Harrison et al., 2000). Sea clutter, which is difficult to distinguish from precipitation, is masked using a fixed

“clutter map” for affected regions of the lowest elevation scans.

For dual polarisation data, a naïve Bayesian classifier based on Rico-Ramirez and Cluckie (2008) is used to identify non-meteorological echoes. This includes parameters based on the copolar correlation coefficient (ρ_{hv}), textures of differential reflectivity (Z_{DR}) and differential phase shift (Φ_{dp}) (section 1.4.2), and clutter phase alignment (CPA) (Hubbert et al., 2009), which replaces the clutter indicator. CPA is used throughout this thesis for quality control, and is described in more detail in chapter 3. The main benefits of the dual polarisation algorithm are in its ability to distinguish sea clutter from precipitation, retaining valid data which would have been removed by a fixed clutter map, and in retaining some low reflectivities associated with drizzle rather than misclassifying them as “noise”. This allows the upgraded radars’ high sensitivities to be fully exploited at short range.

1.5.2 Power loss corrections for blockage and attenuation

Once non-meteorological echoes have been identified and removed from the raw radar data, the remaining reflectivities must be corrected for power loss effects. Physical obstructions such as complex terrain can block all or part of the radar beam, causing a reduction in transmitted power which must be accounted for in order to derive accurate meteorological properties from measurements at longer range.

In the Radarnet system, beam occlusions are mapped to the radar sampling volume using a high resolution digital terrain map and the standard beam propagation model (equation 1.15). Where partial beam blockages affect less than half of the radar beam the data are flagged “usable”, but a power correction is applied to the measured reflectivity as part of correction for the vertical reflectivity profile (section 1.5.3). If a blockage affects more than 50% of the beam the data are flagged as unusable. Single ray occlusions are interpolated from reflectivities in adjacent rays, but for larger occluded sectors, QPE data is taken from higher elevation scans.

“Attenuation” refers to the weakening of the transmitted signal propagating behind any target in proportion to the amount of backscattered power. This reduction in transmitted power means that the beam is less sensitive to precipitation at ranges beyond strong echoes. Attenuation at C-band can reduce the intensity of long range precipitation, but complete extinction of the radar beam is rare, and corrections are usually effective in improving the quality of the attenuated data.

Correction for attenuation is performed for both single and dual polarisation data from the Met Office network. The single polarisation method is iterative, based on the relationships between rain rate and attenuation derived by Hitschfeld and Bordan (1954)

and Gunn and East (1954), with:

$$A_{dB} = 0.0044R^{1.17} \quad (1.23)$$

for the two way attenuation A_{dB} at each range gate. An estimate of attenuation from the first non-clutter range gate is used to adjust the reflectivity at the next gate, before calculating the attenuation from that next gate, adjusting the next, and so on. Such iterative correction methods are known to be unstable, so the Met Office correction is capped operationally at a factor of 3 in linear reflectivity (approximately 5 dB).

The dual polarisation attenuation correction method follows a two step procedure using reflectivity, Φ_{dp} and long range radiometric noise emissions (Thompson et al., 2011). Although Φ_{dp} can be used directly in dual polarisation attenuation correction schemes (Bringi et al., 1990), the coefficient of the $A(\Phi_{dp})$ relationship has been found to vary significantly with both weather and climatology (Carey et al., 2000; Park et al., 2005; Vulpiani et al., 2008). The Met Office method uses a combination of the iterative calculation, $A(\Phi_{dp})$ with an *a priori* coefficient, and emissions to constrain the total path integrated attenuation at each radar azimuth. The median of these three estimates dictates the most suitable coefficient of Φ_{dp} , which is then used to calculate and correct for attenuation at each range gate. This method has been shown to reduce over-correction resulting from the instability of the iterative reflectivity-based correction, and to improve the performance over a fixed coefficient $A(\Phi_{dp})$ relation (Husnoo et al., 2018, additional details Husnoo 2017, personal communication).

1.5.3 Estimating reflectivities at ground level

Radar reflectivities corrected for attenuation and beam blockage can be assumed to give a reliable representation of the meteorological properties of the radar pulse volume at the measurement height. However, for QPE, the desired measurement is of rainfall at the ground. Chapter 2 presents a more detailed review of the microphysics underlying the vertical profile of reflectivity (VPR), and the literature pertaining to its classification and correction in radar PPIs. An introductory overview is provided in this section.

The radar reflectivity measurement is influenced by both the diameter D and dielectric constant κ of hydrometeors within the radar pulse volume (equation 1.7). Changes in both of these parameters with phase, which is determined by temperature gradient in the lower atmosphere, mean that the atmospheric reflectivity varies significantly with height. This has particular impacts in areas sampled around or above the height of the 0°C isotherm (or “freezing level”).

In typical conditions, hydrometeors just above the freezing level are snow aggregates,

comprised of a mixture of ice and air inclusions. These aggregates are much larger than their liquid counterparts, and their $|\kappa|^2$ is at least a factor of 5 smaller than that of liquid water (section 1.2.1). When these large, aerated snowflakes begin to melt, the accumulation of liquid water in and around the melting flakes causes a sudden increase in κ , and a correspondingly sharp increase in reflectivity of order 5-7 dB. As melting continues, the decreasing diameter and increasing terminal velocity of the target hydrometeors causes the reflectivity to fall again, reaching its final “rain” value typically a few hundred metres below the onset of melting (eg Kitchen, 1997; Kirstetter et al., 2013). This characteristic peak in reflectivity in the melting layer is known as the “radar bright band” (figure 1.7).

Identifying the bright band in radar PPIs is not straightforward. Figure 1.8 shows how the height and vertical extent of the radar beam increase with range. With increasing range the beam passes through rain, melting snow, snow and ice aggregates, and finally overshoots the precipitation top. The geometry of the measurement means that it is not possible to observe variations in reflectivity with height directly with any degree of precision. Additional information is therefore needed in order to quantify and correct for the effects of bright band.

The bright band can pose significant problems for radar QPE. If uncorrected, the enhancement can cause severe (even order of magnitude) overestimation in surface precipitation estimates. Despite the conceptual simplicity of the bright band, however, in reality the details are extremely difficult to model, due to its strong dependence on the specifics of how each snowflake melts (see section 2.2 for a detailed review). Correcting for this vertical profile of reflectivity (VPR) is therefore extremely complex; and yet it is essential to obtaining accurate estimates of surface rain rate.

A large number of studies, both practical and theoretical, focus on situations such as those described above, where there is little vertical motion and the layers of frozen, melting and liquid precipitation are well separated. These conditions are described as “stratiform”. The Met Office corrects for stratiform VPR using a climatological profile shape, which is constrained using gridded NWP forecasts of 0°C isotherm height and satellite cloud top measurements. The idealised shape is fitted iteratively to the radar

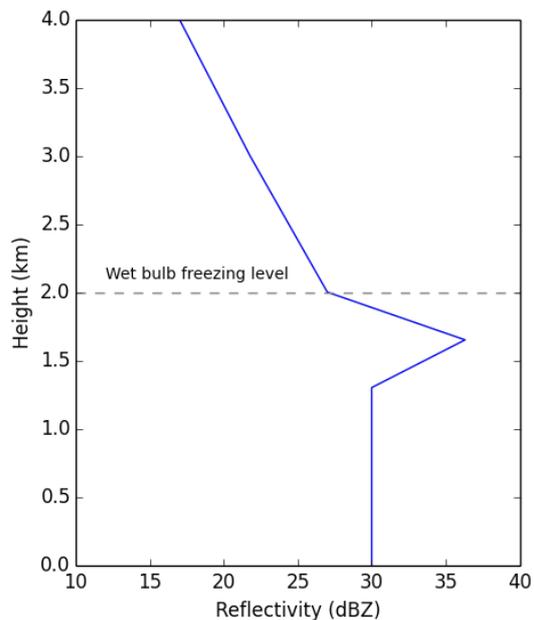


Figure 1.7: Example of the variation in reflectivity with height in a typical bright band case. The quantitative axes are for illustration only.

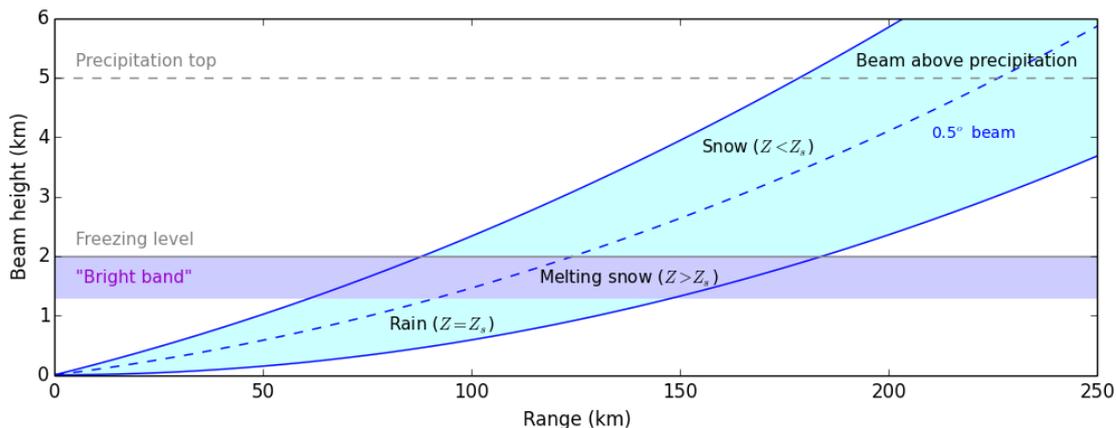


Figure 1.8: Illustration of how and where the 0.5° elevation radar scan can sample liquid, melting and frozen precipitation at different levels of the atmosphere. The height of the freezing level and precipitation top are for illustration only - in reality these have a range of values that vary spatially across the radar domain.

reflectivity measurement at each pixel using a known off-axis beam power profile (Kitchen et al., 1994; Kitchen, 1997). Bright band intensity is empirically related to rain-level reflectivity (Kitchen et al., 1994), and the integration over beam power is truncated to account for partial beam blockages in the vertical. A detailed discussion of this method and the associated profile shape is included in section 2.5.3.

An additional challenge to be overcome in correcting operationally for VPR is that different types of precipitation generate differently shaped profiles. In particular, “convective” situations characterised by strong vertical motion (Steiner et al., 1995) produce distinctly different profile shapes, which usually do not display a bright band (section 2.2.3). The shape of convective profiles is much more variable and difficult to constrain than the standard “stratiform bright band” (figure 1.7), and significantly fewer attempts have been made to develop convective corrections. As part of the Met Office VPR correction scheme, convection is diagnosed using a minimum reflectivity threshold 1 km above the 0°C isotherm. In these cases the reflectivity is assumed to be constant with height, and no VPR correction is applied.

1.5.4 Calculating surface rain rates

Having achieved a reflectivity estimate representative of surface conditions, the final conversion from reflectivity to rain rate can be applied. As discussed in section 1.2.2, the relationship between Z and R is not generally analytical, but can be approximated by a power law of the form $Z = aR^b$. The coefficients of this power law are dependent on details of the drop size distribution.

It is not possible to measure $N(D)$ in real time for operational rainfall estimation. Operational ZR relations are usually climatological, using fixed coefficient values determined from previous studies. The Met Office uses the Marshall-Palmer relation $Z = 200R^{1.6}$ (Marshall and Palmer, 1948), which performs well in the light to moderate stratiform rainfall responsible for the majority of precipitation in the UK.

For radar measurements in the rain layer, dual polarisation parameters can sometimes be used to improve QPEs. The ability to constrain the rain drop size distribution using Z_{DR} (section 1.4.2) is the principle underlying alternative $R(Z, Z_{DR})$ rain rate estimators. Although such estimators arise frequently in the literature (eg Ryzhkov and Zrnić, 1995; Zhang et al., 2001; Brandes et al., 2002, 2003; Ryzhkov et al., 2005), the requirements on Z_{DR} calibration accuracy for QPE are stringent (± 0.1 dB, Thompson (2007)), and require a high quality system isolation not often achieved by operational radar networks (Hubbert et al., 2010).

In intense rainfall KDP can also provide an independent measurement of rain intensity (Brandes et al., 2003). The differential phase shift Φ_{dp} relates both to the average axis ratio and the total water content along the radar beam propagation path (section 1.4.2), and its gradient provides an estimate of water content per km range. As the difference in phase shift between two channels, KDP has the significant advantage in heavy precipitation of being less sensitive than reflectivity to attenuation. However, the inherent noisiness of the Φ_{dp} measurement means that KDP values are highly suspect at low intensities. In the UK, KDP measurements greater than 16° km^{-1} (approximately 10 mm h^{-1} in rain rate) are used for rainfall estimation where the beam is below the melting layer, and have been shown to reduce QPE underestimation significantly at high intensities.

Dual polarisation rain rate estimators have been shown to deliver significant improvements the accuracy of QPEs calculated from rain level measurements (eg Brandes et al., 2003). However in high latitude climates, and particularly in winter months, the majority of the radar composite is derived from measurements at altitudes within or above the melting layer. For 0°C isotherm at 2 km, which represents an average for the UK climate (Kitchen et al., 1994), the domain in which the radar beam is fully within the rain extends only 60 km from the radar (figure 1.8). In conditions such as these, reflectivities in and above the bright band remain a vital source of QPE information, ensuring the continued relevance of research into improving corrections for the VPR.

1.6 Impacts of dual polarisation for radar QPE and motivation for this study

The advent of dual polarisation radar has radically changed the nature of operational radar QPE. Use of multi-parameter classifiers for quality control and dual polarisation rain rate estimators (Brandes et al., 2002, 2003; Giangrande and Ryzhkov, 2008) are feeding improvements in the quality and reliability of radar products worldwide (Ryzhkov et al., 2005; Figueras I Ventura and Tabary, 2013; Helmert et al., 2014). However, the applicability of dual polarisation measurements directly to QPEs is still limited to the rain level. The assumptions underpinning the relationships between Z_{DR} , KDP and rain rate are based on the strongly preferential orientation behaviour of liquid rain drops, but equivalent constraints for more randomly oriented ice crystals, aggregates and melting particle mixtures have not yet been possible to derive. At high latitudes, this severely limits the applicability of dual polarisation rain rate estimators to what is often a only a small proportion of the radar sampling domain (eg figure 1.8). The use of single polarisation reflectivity measurements in and above the bright band for surface QPE, and the methods required to adjust these for VPR, remains an important area for future research.

There has already been significant investigation into the use of dual polarisation information for improving determination of stratiform VPRs (reviewed in chapter 2). Many papers have established skill in ρ_{hv} , Z_{DR} and LDR for identifying and locating the radar bright band (eg Tabary et al., 2006; Boodoo et al., 2010; Hall et al., 2015). For systems without access to freezing level information, using dual polarisation parameters to determine a domain-averaged melting layer height is a significant refinement of the climatological bright band corrections proposed in much of the literature, and provides for the first time the prospect of a VPR correction accurate enough for real time application (eg Tabary, 2007). However, little attention has been given to the use of microphysical information contained in dual polarisation measurements to determine profile characteristics on the more local scale. The demonstrable skill of dual polarisation parameters in hydrometeor classification raises the question as to whether this information could be combined with knowledge of microphysics underlying the VPR to improve the accuracy of local surface reflectivity estimates.

The pixel-by-pixel Kitchen et al. (1994) scheme applied in the UK provides a unique testbed for investigating the potential of dual polarisation to improve VPR classification and correction at the local scale. While stratiform VPRs are well-treated by this scheme, the identification and correction of non-bright band VPRs could benefit significantly from further research. Use of a reflectivity-based criterion for convective diagnosis is known to underdiagnose non-bright band conditions, meaning that potentially a sig-

nificant proportion of non-bright band cases are being corrected inappropriately. This implies widespread underestimation of rain rates in the very high impact situations for which accurate QPEs are most urgently required. Dual polarisation measurements have the potential to provide more reliable methods of identifying bright band (Smyth and Illingworth, 1998; Illingworth and Thompson, 2011), which could reduce the occurrence of inappropriate bright band corrections and the associated rain rate underestimation.

Beyond identifying non-bright band conditions, recent observational literature (eg Delrieu et al., 2009; Kirstetter et al., 2013; Matrosov et al., 2016) suggests that the assumption that reflectivity is constant with height in all cases without bright band may be inaccurate. Once the issue of classification has been addressed, there is scope for improving the characterisation of non-bright band profiles within the Kitchen et al. (1994) framework, by developing new idealised local profiles for different types of VPR.

This thesis aims to apply new information from dual polarisation parameters to improving correction for VPR in the Met Office operational radar processing chain. High quality measurements from the upgraded radar network will be used to distinguish intelligently between different types of vertical profile, which are characterised using a large new dataset from the C-band research radar at Wardon Hill. The investigation focuses on the linear depolarisation ratio (LDR), which has been shown to respond particularly to the large melting snowflakes responsible for stratiform reflectivity bright bands (Smyth and Illingworth, 1998; Illingworth and Thompson, 2011).

This thesis is organised as follows. Chapter 2 reviews the current literature on VPR classification and correction schemes, and includes a detailed description of the Kitchen et al. (1994) approach on which this thesis builds. Chapter 3 provides details of the Wardon Hill radar, and describes the high resolution dual polarisation dataset collected to support this investigation. In chapter 4, the quantitative skill of LDR in distinguishing between different VPR types is investigated, and is compared to the skill of the current UK operational convective diagnosis criterion. Chapter 5 develops this result into an operationally feasible algorithm, and demonstrates the impact of LDR-based classification on QPE accuracy in a real time environment.

Having improved the real time classification of VPRs using dual polarisation measurements, the remainder of this thesis investigates refinements to the different profile shapes available for correction in different meteorological conditions. Chapters 6 and 7 exploit the reflectivity information from the high resolution profile dataset to suggest improvements to the idealised shapes used for VPR correction in the UK. In chapter 6 a new non-bright band VPR shape is proposed, with support from previous literature, and evaluated through both simulations and a real time implementation. Chapter 7 uses observations from stratiform VPRs in the Wardon Hill dataset to improve residual long range QPE bias through small changes to the Kitchen et al. (1994) idealised profile. Fi-

nally, chapter 8 summarises the outcomes of these investigations and suggests areas for future research and development.

Chapter 2

Existing approaches to VPR: classification, determination and correction

2.1 Introduction

Radar quantitative precipitation estimation (QPE) is achieved through the conversion of a reflectivity measurement aloft into a rain rate estimate at the ground (section 1.5). An important step in this process is the adjustment of a meteorological reflectivity from its value measured at height to an estimate of the value near the ground. This is known as correcting for the vertical profile of reflectivity, or VPR.

The VPR defines the variation in atmospheric reflectivity as a function of height above the ground surface. While there is little variation at low levels other than that caused by partial beam blocking, in areas where the radar beam samples above or close to the 0°C isotherm the difference between measured and surface reflectivities can exceed an order of magnitude. In high latitude climates, the ranges at which radar PPI measurements intersect the melting layer combined with broadening of the radar beam cause these effects to influence a significant proportion of the radar domain. VPR is therefore a significant, if not the most significant, source of error in high latitude radar QPEs.

In this chapter the existing literature pertaining to VPR classification, determination and correction is reviewed. Section 2.2 describes the microphysical processes that occur in different types of precipitation, with reference to both observational and modelling studies, and how they lead to the general shape of the resulting profiles. This provides context for the many approaches to correcting for the VPR in PPIs which are reviewed

in section 2.3. These range from purely empirical ratio-based methods to highly complex linear parameterisations and probabilistic schemes, on a range of different spatial and temporal scales. Section 2.4 presents approaches to the classification of different precipitation types, such as convection, which are needed as input to certain bright band and VPR correction schemes. Three operational correction schemes and their supporting classification methods are discussed in detail and compared in section 2.5. Finally, section 2.6 summarises the main points of the review and revisits the motivations for the remainder of this thesis.

2.2 Microphysical basis of the VPR

In stratiform conditions, which dominate at high latitudes, the VPR has a characteristic structure that includes the radar “bright band”: a region of enhanced reflectivity below the freezing level attributable to large melting snowflakes. The spatial influence and magnitude of bright band enhancement in radar PPIs (figure 1.8) has made this an area of significant interest for observational, microphysical and modelling studies over the past several decades.

The manifestation of the VPR in radar measurements is a function of both the underlying microphysics and the radar sampling characteristics. The reflectivity of a beam-filling meteorological radar target is defined in equation 1.7 and depends on drop diameters D , dielectric factors κ , and the drop size distribution $N(D)$ per unit volume. Each of these variables depends strongly on phase, and by extension varies with height in the atmosphere.

In this section, the microphysics underlying different types of precipitation (specifically “stratiform” and “convective” rain) is described, and is used to explain the shapes of the resulting reflectivity profiles. Section 2.2.1 first presents a definition of “convection” and outlines the major distinguishing features of the two precipitation types. Section 2.2.2 then proceeds in much more detail through the well defined vertical structures that occur in stratiform rain, considering height levels above, within and below the melting layer, and with a strong focus on the radar bright band. This is followed by a much shorter section on convective VPRs (section 2.2.3), which reflects not so much the scientific understanding of convection itself, but the limited relevance of that science for interpreting reflectivity profiles. The section concludes with an overview of other types of VPR which are less frequently mentioned in the literature.

2.2.1 Defining “stratiform” and “convective” rain

Precipitating systems can be classified broadly in terms of their vertical structure as “stratiform” or “convective”. Steiner et al. (1995) define stratiform precipitation as existing when vertical wind speeds are lower than the fall velocity of ice-phase hydrometeors (typically $\approx 1 \text{ m s}^{-1}$). This means that precipitation formation and growth processes in stratiform conditions occur while hydrometeors are falling.

The dominant ice growth process in cold stratiform precipitation is initially vapour diffusion (via the Bergeron-Findeisen process), in which water vapour is directly sublimated onto existing ice crystals. At lower levels, as the temperature rises towards freezing, aggregation becomes more significant, producing lower density snowflakes. When these snowflakes begin to melt, the abrupt increase in dielectric factor κ (section 1.2.1) causes a strong reflectivity enhancement. A decrease in reflectivity with continued melting, through decreasing drop diameter and number density, completes the classic “bright band” peak that is characteristic of stratiform precipitation. Reflectivities at the base of the bright band tend to be slightly higher than those immediately above it (Fabry and Zawadzki, 1995), and undergo little variation between this “reference level” and the surface.

Convective precipitation is defined by the presence of significant vertical air motion. During convective initiation, updrafts of $1\text{-}10 \text{ ms}^{-1}$ lift water vapour from lower levels of the atmosphere, increasing the relative humidity (RH) at high levels. When the RH reaches 100% the air is no longer sub-saturated with respect to liquid water, so liquid drops do not evaporate to feed vapour diffusion. Instead the high relative concentration of liquid water droplets leads to “riming”, where liquid water freezes directly onto the surface of ice crystals. Growth dominated by riming produces solid hydrometeors in the form of high density graupel and hail, rather than the low density snow aggregates characteristic of stratiform precipitation. In well developed convection, entrainment into updrafts from low levels can expose hydrometeors to several cycles of partial melting, refreezing and riming, resulting in the growth of large hail which can cause damaging impacts at the ground.

2.2.2 Stratiform VPRs

Vertical reflectivity profiles are usually described in terms of a ratio or difference from the reflectivity at the “reference level” (section 2.2.1). In stratiform precipitation, where there is little vertical motion, the VPR shows a characteristic layered structure on three levels: the melting layer “bright band”, frozen precipitation above it, and the rain at the “reference level” and below.

The key features of stratiform VPRs are conceptually simple, and have been described both here and in section 1.5.3. The following subsections present further detail on the variations in κ , D and $N(D)$ within the three layers of stratiform rain, and how these affect the atmospheric reflectivity profile. Notation wise, the properties of liquid drops are indicated by the subscript w , frozen by s (snow) or i (ice), and melting hydrometeors by the subscript m .

Hydrometeor formation and growth

In stratiform precipitation ice crystals form, often by deposition or contact nucleation, at a generating level, and grow by vapour deposition. Since κ_i is constant, the change in reflectivity with height near the top of the profile is a function of the increasing number of ice particles, N , and the growth in diameter D .

As the atmospheric temperature rises at lower levels, ice crystals begin to aggregate and form snowflakes. Aggregation causes a steady increase in reflectivity, as the decrease in number concentration with aggregation is outweighed by the impact of increasing diameter. Once aggregation begins the dielectric factor κ_s is no longer constant, but varies in proportion to snow density (Sauvageot, 1992, chapter 2, pg 97-98). Since $Z \propto |\kappa|^2 D^6$ (equation 1.7), it follows that the reflectivity of each individual snowflake is proportional to the square of its mass, and that the total reflectivity is related to the overall mass of precipitation per unit volume.

Given the majority of VPR modelling literature is focused on the bright band, there are few models which include a fully parameterised reflectivity profile shape in the ice layer. The microphysically-based correction scheme of Kirstetter et al. (2013) uses a parameterisation for ice-level reflectivity in which N and D increase linearly from the precipitation top until the onset of melting, and κ_s is defined by the matrix-inclusion model of Boudevillain and Andrieu (2003) (discussed further below). This model has the benefit of simplicity, but neglects the difference between the vapour deposition and aggregation layers, and has a tendency to overestimate the reflectivity gradient above the melting layer.

A more common element of modelling studies is the inverse relationship between snowflake density and diameter, which has the general empirical form $\rho_s = \gamma_s D_s^{y_s}$ (where $y_s < 0$). The chosen form can be combined with mass conservation and a standard assumption of no aggregation or breakup during melting to derive direct analytical relationships between reflectivities immediately above and below the melting layer (Hardaker et al., 1995; Szyrmer and Zawadzki, 1999; Wood et al., 2015). These studies typically aim toward developing detailed, predictive models of the bright band reflectivity profile.

Modelling the bright band

The most prominent feature of radar VPRs is the reflectivity enhancement associated with the stratiform melting layer, or “bright band”. The impacts of uncorrected bright band in radar QPEs include visible ring-shaped artefacts containing order of magnitude overestimation errors, which through a combination of freezing level height and beam broadening can affect the majority of the radar’s domain. A desire to mitigate these impacts has led to the large number of studies aiming to characterise bright band behaviour and correct for the effects in radar precipitation estimates.

The origin of the bright band can be easily understood with reference to the definition of reflectivity and its three contributing variables. Immediately above the 0°C isotherm, stratiform precipitation is made up of a population of aggregate snowflakes. The presence of air inclusions means that snow aggregates at this level are much larger than their melted counterparts ($D_s \gg D_w$). At the onset of melting, liquid water begins to collect in air inclusions and on the surface of aggregate snowflakes. This leads to a rapid increase in κ towards that of liquid water. Since $\kappa_w \approx 5 \times \kappa_i$ and $\kappa_i > \kappa_s$ due to the density relation (section 1.2.1), the immediate increase in reflectivity can be extremely large. Matrosov et al. (2007) observe typical bright band enhancements in PPI scans of 5-7 dB, which is likely to be smaller than the true enhancement due to the vertical smoothing caused by broadening of the radar beam. As melting continues, the increase in κ is balanced by decreases in both drop diameter (through melting), and the decrease in drop concentration N through increasing fall velocity (from 1 m s⁻¹ in snow to 5 m s⁻¹ in rain, Mittermaier et al., 2004). This causes a decrease in reflectivity between the initial peak and completion of melting, where precipitation is recognisable as rain and which defines the standard “reference level”.

Despite its conceptual simplicity, the details of the melting layer bright band are extremely difficult to model. The size and depth of the observed bright band peak depend strongly on how each snowflake melts: specifically where the liquid water collects, as inclusions or on the surface of the melting snowflake (Fabry and Szyrmer, 1999, discussed below). The modelled reflectivity is therefore extremely sensitive to the chosen melting model, which may or may not reflect the realities of melting in any given bright band.

The difficulties of accurately modelling stratiform melting are exemplified by the early attempt of Hardaker et al. (1995) to characterise the bright band shape in detail for VPR correction. This paper introduces a physically-based one-dimensional bright band simulator, which uses a model of melting snowflakes as ice with air inclusions and an adiabatic lapse rate of 6°C km⁻¹ above the melting layer to calculate a full analytical profile of reflectivity with height. Modelled VPRs are then compared with radar observations. The peak bright band reflectivities generated by the Hardaker et al. (1995) model

were found to be correlated with surface rainfall rate over the range 0-7 mm h⁻¹, which is consistent with both the vertically-pointing radar observations of Fabry and Zawadzki (1995) and the bright band area relation of Kitchen et al. (1994). However in general, despite the careful consideration given to every detail of this model, the features of the simulated bright band did not agree well with observed VPRs.

Although their outcomes differ, there are certain consistent assumptions which underpin the majority of microphysical bright band studies, and lead to certain standard results. Typical models of the melting layer are one-dimensional, and assume no aggregation or breakup of partially melted droplets (Hardaker et al., 1995; Szyrmer and Zawadzki, 1999; Heyraud et al., 2008; Kirstetter et al., 2013). The requirement for mass flux conservation then applies on an individual particle basis, which implies the constraint:

$$D_s^3 = \frac{\rho_w D_w^3}{\rho_s} \quad (2.1)$$

on the diameter of any given snowflake. Combining this with a snow density-diameter relation makes it possible to model the diameter of a snowflake immediately above the melting layer in terms of the diameter of its melted counterpart. Similarly, the behaviour of D_m with melted mass fraction f (as a function of height) is governed by mass conservation:

$$D_m(D_w, f) = D_w \left[\left(\frac{\rho_w}{\rho_s} \right) (1 - f) + f \right]^{\frac{1}{3}} \quad (2.2)$$

The initial assumption of no aggregation or breakup implies that knowledge of the rain drop size distribution $N(D_w)$ provides constraints on both $N(D_m)$ at all heights within the melting layer, and on $N(D_s)$ immediately above it. A further standard assumption that all hydrometeors fall at their terminal velocity V_t , which varies with diameter and phase (Heyraud et al., 2008), renders $N(D_w)$ sufficient to describe fully the changes in particle size distribution with height up to the top of the melting layer (Hardaker et al., 1995). In practise however, $N(D_w)$ itself is a parameterised approximation, and the combination of many empirical relations and assumptions means such models cannot yet reproduce with reliable accuracy the shape of observed VPRs.

Apart from changes in diameter, reflectivity behaviour in the bright band is also influenced by changes in the dielectric factor. The dielectric factor of melting hydrometeors, κ_m , can be calculated using a modified form of the equations of Boudevillain and Andrieu (2003) for snow aggregates. Modelling a snowflake as an ice matrix with air inclusions,

they define:

$$m_s^2 = \frac{1 - f_a}{1 - (1 - \beta)f_a} m_i^2 + \frac{\beta f_a}{1 - (1 - \beta)f_a} m_a^2 \quad (2.3)$$

$$\text{where } \beta = \frac{2m_i^2}{m_a^2 - m_i^2} \left[\frac{m_a^2}{m_a^2 - m_i^2} \log \left(\frac{m_a^2}{m_i^2} \right) - 1 \right] \quad (2.4)$$

where f_a is the volume fraction of air inclusions within the snowflake and κ_s relates to m_s as defined in equation 1.5. Replacing m_i with the complex refractive index of the chosen matrix and m_a by that of the inclusions, these formulae can be extended to describe the mixtures of water, ice and air that comprise melting snow.

Since equation 2.3 is asymmetrical, the exact value of κ_m depends on the nature of the three-level matrix-inclusion model used to describe melting snowflakes. Fabry and Szyrmer (1999) evaluate the following six different microphysical models for melting snow:

1. A snow core (ice inclusions in an air matrix) with a water shell
2. Snow (ice inclusions in air) inclusions in a water matrix
3. Air inclusions in a melting snow (ice inclusions in water) matrix
4. Melting snow (ice inclusions in water) inclusions in an air matrix
5. As model 4, but with a density discontinuity (higher density core surrounded by lower density shell)
6. Inner core represented by model 3, with an outer shell represented by model 4

Each of these is used to simulate the melting layer reflectivity profile, and the results compared with observations to infer the model that best represents reality. Simulated bright band reflectivities were found to be extremely sensitive to the melting model. The final, most complex model best reproduced the observations available to this study; models 1-3 overestimated the magnitude of the bright band, whilst models 4 and 5 underestimated the peak. On this basis model 6 is adopted by Kirstetter et al. (2013) in developing their microphysically-based VPR determination scheme.

A key conclusion arising from Fabry and Szyrmer (1999) echoes that of earlier work (Hardaker et al., 1995, and others): that modelling the bright band is extremely complex and sensitive to small details of the underlying microphysical situation. It follows that detailed microphysical simulations are unlikely ever to provide a practical solution to VPR determination in real time. Approximations and parameterisations remain necessary in estimating the VPR from real time data, in order to make the corrections required for operational QPE.

Reflectivities below the reference level

The major microphysical processes governing cold stratiform rain occur above the melting layer. Low-level growth, aggregation and breakup of liquid drops are generally assumed to be negligible, implying no change in D_w or $N(D_w)$ with height, and since κ_w is constant reflectivity is not expected to change significantly below the reference level (at the base of the melting layer). The exceptions to this are warm rain, which is not dealt with here, and contributions from orographic processes such as the feeder-seeder mechanism (Carruthers and Choularton, 1983). Kirstetter et al. (2013) parameterise low-level processes with a linear reflectivity gradient from the base of the melting layer, whilst Kitchen (1997) account for contributions from the feeder-seeder mechanism using a model estimate of orographic enhancement, which is applied to the lowest 1.5 km of the VPR in complex terrain. Changes in reflectivity with height in the rain layer are not considered further as part of this thesis.

2.2.3 Convective VPRs

Convective precipitation is associated with relatively strong vertical motion, which can be in the region of 10 m s^{-1} (Steiner et al., 1995). These conditions favour the formation of small, high density ice phase hydrometeors through riming, rather than the large snow aggregates characteristic of stratiform precipitation.

Whilst the reflectivity characteristics of solid ice are similar to those of snow, the melting profile of convective precipitation differs from that of stratiform bright bands. Unlike snow, there are only two possible matrix-inclusion models of melting ice (noting the asymmetry of equation 2.3). Solid ice particles have smaller diameters and higher fall speeds (Kumjian et al., 2016) than snowflakes of the same mass (and hence reflectivity), and also melt more slowly due to their lower ventilation coefficient (Szyrmer and Zawadzki, 1999; Fabry and Szyrmer, 1999). This means that convective melting layers tend to be both deeper and less enhanced than those of stratiform precipitation. The depth of convective melting can also be extended by updrafts, causing wet hail and graupel to be lifted to heights well above the freezing level. This disrupts the clear vertical structures observed in stratiform conditions, and makes it more difficult to describe and characterise a “typical” convective VPR.

In strong convective conditions, rimed hydrometeors can become large enough to move from the Rayleigh scattering regime ($D \ll \lambda$ and $Z \propto D^6$) into Mie scattering (where $D \approx \lambda$). This occurs at smaller diameters for shorter wavelength radars (see also section 1.3). In the Mie scattering regime the relationship between reflectivity and diameter is not monotonic, but oscillating, and remains so at diameters between the Rayleigh and

geometric limits (where $D \gg \lambda$). When rimed hydrometeors become large enough to cause Mie scattering, therefore, reflectivity values are no longer useful for estimating precipitation rates.

2.2.4 Other types of VPR

Despite the focus of the majority of the literature on “stratiform” and “convective” profiles, there are a handful of existing papers which consider different specific subsets of precipitation type. Fabry and Zawadzki (1995) present a detailed observational study of VPRs observed by a vertically pointing X-band radar. Rather than the traditional two, the authors propose a total of five profile classifications, which are defined based on temperature and vertical velocity characteristics.

The most common precipitation profile observed by Fabry and Zawadzki (1995) is stratiform “rain with bright band” (section 2.2.2). A difference of 1-2 dB is typically observed in reflectivities immediately above and below the bright band, with an average initial decrease in ice reflectivity with height of 6-7 dB km⁻¹. A second stratiform VPR type is defined as “low-level rain”, where the top of the precipitation profile is below the 0°C isotherm and hydrometeor growth processes occur in the liquid phase. In shallow “warm rain” cases, the dominant growth processes of collision and coalescence produce fewer and smaller rain drops, resulting in very light rain or drizzle and a reflectivity profile that decreases monotonically with height. Fabry and Zawadzki (1995) define their third “shower” class as applying to warm rain in convective conditions, where reflectivity is constant with height up to a level approximately 1 km below the precipitation top. These shallow warm profiles rarely cause significant precipitation at higher latitudes.

The fourth and fifth profile types are subdivisions of events typically classed (in the observational, if not the microphysical literature) as convective. “Rain from compact ice” displays a VPR similar to that of stratiform rain, with $Z(h)$ relatively constant below the melting layer and decreasing with height above it, but without a bright band. This can be explained by the similarity in diameter and terminal velocity between compact ice and liquid hydrometeors, which means there is little decrease in reflectivity below the bright band peak (section 2.2.2). The inference of Fabry and Zawadzki (1995) is therefore that these profiles occur in cases when the ice particle size distribution contains a high proportion of small, dense hydrometeors. This is supported by the more recent observations of Matrosov et al. (2016), who find that rain from “compact ice” VPRs contains a much larger proportion of small drops at the surface than bright band or convective rainfall. This profile is also observed as a distinct and separate “type” later in this thesis (chapter 4). The final category described by Fabry and Zawadzki (1995) is the traditional “deep convection”, associated with strong updrafts and downdrafts, which is

observed in structures such as thunderstorms and squall lines.

2.3 VPR determination and correction in radar PPIs

The VPR as sampled by radar PPIs, and particularly the melting layer bright band, can have a significant impact on the quality and accuracy of QPEs. The aim of VPR correction is to estimate the surface reflectivity from an in-situ radar measurement, which is both taken aloft and subject to the effects of beam broadening. This is usually done via an intermediate step, in which the shape of the full VPR or bright band is estimated using the available measurements.

Due to the limited vertical sampling of radar PPIs, it is not possible to observe or determine local VPRs using only reflectivity volume measurements; and as established in section 2.2.2, neither is it possible to model VPRs in detail using microphysics alone. Approaches to the determination problem include large scale spatial or temporal averaging of radar observations, simple parameterisations, and the use of independent information to constrain the profile.

This section presents a review of the range of approaches to VPR determination and correction available in the current theoretical literature. Section 2.3.1 begins by introducing the terminology of “mean apparent” and “background” VPRs, which are presented in terms of their defining equations. These will be referred to throughout the subsequent discussion. The early method of ratios to determine mean apparent VPRs is followed by a discussion of the different approaches to estimating mean field background VPRs: by far the most common approach in the existing literature. More localised and flexible methods are covered in sections 2.3.4 and 2.3.5, which describe the use of ancillary data to constrain underdetermined profiles. Finally, the overarching questions of scale, accuracy and representativity are discussed in section 2.3.7.

2.3.1 Introduction to terminology

In the existing literature the term “VPR” is used interchangeably in referring to what are actually two different functions. The distinction between the atmospheric “background” and measured “apparent” VPRs can be understood by considering the manner in which the radar samples a vertical profile.

The power transmitted by a radar pulse is a function of the angular distance from the centre of the beam. For example in the UK, the fraction of beam power transmitted

between off-axis angles ϕ and $\phi + d\phi$ is:

$$f(\phi)d\phi = P_t^{-1} \left[\frac{\sin(k\phi)}{k\phi} \right]^4 d\phi \quad (2.5)$$

where $k = 159.46$ for ϕ in radians, and P_t is the integrated power over the total beam width (Kitchen et al., 1994). Given the half-power angular beam width for an operational radar is typically of order 1° , the vertical extent of the beam increases significantly with range, causing the detail of the vertical profile to become smoothed. Data from a low elevation PPI samples more than a kilometre in the vertical beyond 50 km range (figure 1.8) - while features of interest in the VPR, such as the bright band, may span depths of only a few hundred metres. There is substantial difference, therefore, between the underlying “background” VPR and what the radar can resolve.

The vertical profile of reflectivity $z(h)$ at location \mathbf{x} is defined as a dimensionless ratio of linear reflectivities at the surface Z_s and aloft:

$$Z(h, \mathbf{x}) = z(h, \mathbf{x})Z_s(\mathbf{x}) \quad (2.6)$$

(from Andrieu and Creutin, 1995a, equation 1). The reflectivity Z_a observed by a radar with an angular beam width of 2α is then described by:

$$Z_a(h, \mathbf{x}) = \int_{-\alpha}^{\alpha} f^2(\phi)Z_b(h, \mathbf{x})d\phi \quad (2.7)$$

$$\text{where } Z_b(h, \mathbf{x}) = z_b(h, \mathbf{x})Z_s(\mathbf{x}) \quad (2.8)$$

$$\text{and } Z_a(h, \mathbf{x}) = z_a(h, \mathbf{x})Z_s(\mathbf{x}) \quad (2.9)$$

Equation 2.8 defines the true atmospheric vertical reflectivity profile $z_b(h, \mathbf{x})$: the “background VPR”. Equation 2.9 defines the profile directly observed by the radar, $z_a(h, \mathbf{x})$, which is a convolution of the background VPR with the radar beam power profile. This is termed the “apparent VPR”. Whilst both the background and apparent VPRs can be spatially varying, $z_a(h, \mathbf{x})$ has superimposed on this an additional range-dependent behaviour associated with broadening of the radar beam.

The majority of the research literature assumes a “global” VPR shape that does not vary spatially across the radar domain. This can be expressed as $z_b(h, \mathbf{x}) \equiv z_b(h)$. Since the radar beam height is a function of range from the radar, both the apparent VPR and the required bias correction vary spatially, but the underlying profile remains constant. Some more recent papers allow for a profile with limited spatial variability, for example by determining bright band parameters independently at different azimuths (Giangrande et al., 2008; Boodoo et al., 2010). Then $z_b(h, \mathbf{x}) \equiv z_b(h, \theta)$, where θ is the azimuth angle, usually measured clockwise from grid North. Only a few papers allow for full spatial

variability of the background VPR (Kitchen, 1997), where $z_b(h, \mathbf{x}) \equiv z_b(h, \theta, r)$. These “local” schemes rely on information additional to the radar reflectivity, such as model freezing height level fields.

2.3.2 Mean apparent VPRs: the ratio method

An early approach to VPR determination, the method of ratios was developed by Andrieu and Creutin (1995a) with the aim of deriving a mean field background reflectivity profile from volume reflectivity measurements. The generalised definition of the VPR (equation 2.6) implies that the ratio of colocated reflectivities from different PPIs is independent of reflectivity at the reference level: it depends only on the form of the VPR. (This early assumption of independence contrasts with later findings, for example in Kitchen et al. (1994) and Rico-Ramirez et al. (2005), that Z_s is in fact correlated with bright band intensity.) By extension, given enough measured ratios from scans at different elevation angles, a complete VPR as a function of height $z(h)$ can be inferred. Andrieu and Creutin (1995a) formulate this as a matrix inversion problem, using ratios from short range reflectivity measurements (where $z_a(h) \approx z_b(h)$) to constrain one of two piecewise idealised VPRs (with and without bright band). Bias corrections based on this short range profile are then applied to the whole radar domain. Case study evaluation (Andrieu and Creutin, 1995b) finds that this method outperforms a reference scheme based on event-mean intensity ratios.

2.3.3 Mean field background VPRs

Vignal et al. (1999) observe that it is impossible to obtain accurate vertical profiles from an individual radar PPI in the absence of additional information. This is particularly true for measurements beyond 60 km range, where it is generally agreed that even apparent VPRs derived from multiple scans are no longer useful due to the impacts of beam broadening (eg Smyth and Illingworth, 1998; Vignal et al., 1999; Kirstetter et al., 2010; Kalogiros et al., 2013). The usual approach to this problem is to include additional *a priori* information in the form of climatological or parameterised vertical profiles to constrain the background VPR.

Linear parameterisations of the VPR

Figure 2.1 shows, in schematic form, the typical shapes used to correct for bright band and VPR in stratiform precipitation. Parameters differ for different schemes, but may include any or all of the following:

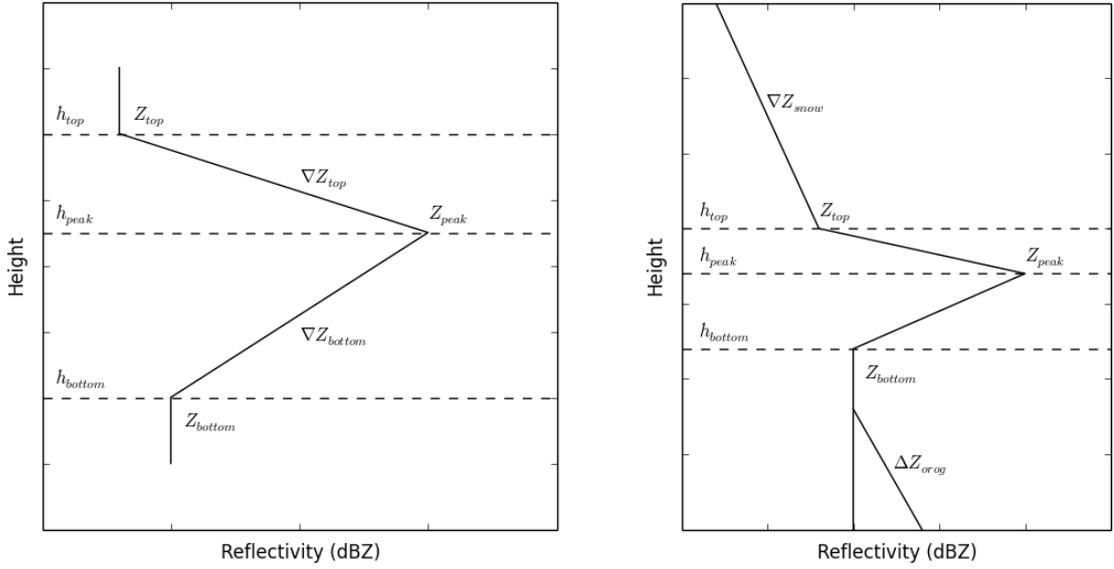


Figure 2.1: Left: schematic linear bright band profile with labelled parameters, for use in correcting radar reflectivity PPIs. Right: extended bright band VPR shape with additional labelled parameters.

- Bright band top height / height of onset of melting (h_{top}) with respect to ground level
- Bright band depth or bright band base height / height of completion of melting (h_{bottom}) with respect to ground level
- Peak reflectivity value (Z_{peak}) or difference between peak and reference level reflectivities (ΔZ , usually expressed as the dB enhancement Z_{peak} over Z_{bottom})
- Position of the peak within the bright band
- Slope of the reflectivity gradient between the melting layer peak and boundaries (∇Z)
- Difference between reflectivities at the top and bottom of the melting layer ($Z_{top} - Z_{bottom}$)

(In this context and throughout, the gradient symbol ∇ is used to denote the derivative of a variable with respect to height only.) The number of variable parameters varies with VPR scheme, with differing use of climatological values and data from ancillary sources (such as model and satellite information). The majority of VPR schemes determine the remaining variable parameters as average values, defining a single background VPR applicable to the whole radar domain.

Kitchen et al. (1994) is a major exception to the “global” rule, having been designed to apply independently on a pixel-by-pixel basis. However, their form of the linear parameterised VPR has formed the basis of a number of global schemes. The idealised VPR for stratiform rain is based on climatological measurements from a high resolution S-band radar (Chilbolton, UK). The profile has a single variable parameter: the linear reflectivity at the reference level. The 0°C isotherm (h_{top}) is determined by extrapolating surface temperature observations assuming a 6° km^{-1} lapse rate. The bright band depth is fixed at 700 m to allow for uncertainties in h_{top} ; the bright band profile is symmetrical around a central peak ($Z_{top} = Z_{bottom}$); and Z_{peak} is related to Z_{bottom} through bright band area:

$$A_{bb} = 0.5 \times (h_{top} - h_{bottom}) \times Z_{peak} \quad (2.10)$$

$$\log(A_{bb}) = 1.42\log(Z_{bottom}) + 2.1 \quad (2.11)$$

Equation 2.11 is derived in Kitchen et al. (1994) using a linear fit to empirical data. An additional correction for orographic enhancement is included in the lowest 1.5 km of the VPR, where required by the terrain. Reflectivities above the melting layer are corrected using a piecewise linear profile derived from the Chilbolton dataset. The authors derive a selection of ice profiles based on the depth of the profile above the melting layer, which is estimated in real time using satellite- or radiosonde-derived cloud top measurements.

Tabary (2007) use a similar linear parameterisation to Kitchen et al. (1994), but allow variation in four different parameters: bright band top, bright band depth, bright band intensity, and the slope of the profile above the freezing level. Climatological information is used to develop 240 *a priori* linear VPR shapes, of which the most suitable is chosen by fitting to observed ratios of hourly rainfall accumulation from different scan elevations. The determined mean field VPR is then used to calculate correction factors for rainfall rates on the basis of the last hours accumulations.

Zhang and Qi (2010) propose a correction for bright band effects in the USA radar network by combining mean apparent VPRs with an idealised bright band profile. A “bright band affected area” (BBA) is first defined as stratiform pixels within a certain range of the 0°C isotherm (obtained from model fields or temperature soundings). Apparent VPRs from multiple scans at pixels within the BBA are then used to constrain a five-parameter bright band model, from which a dB bright band correction factor is derived as a function of range. The method is shown to be effective in reducing positive rain rate bias through comparison with rain gauge accumulations. This work is later refined by Qi et al. (2013b), who demonstrate further improvements to QPEs using ρ_{hv} to determine the location and depth of the bright band.

Matrosov et al. (2007) use an idealised profile with five variable parameters to calculate a mean field VPR across the radar scan. Their method uses ρ_{hv} PPIs to locate the melting layer, and derives the apparent bright band enhancement (in dB) as a function of range (in km) using a parameterised relation derived from range height indicator (RHI) scans:

$$\Delta Z = -0.05r + 6.8 \quad (2.12)$$

While the bright band peak is assumed to be centred in height, Z_{top} and Z_{bottom} are extracted from observations and are generally not equal, unlike the profile of Kitchen et al. (1994). A climatological value $\nabla Z_{snow} = -5.1 \text{ dB km}^{-1}$ is applied to reflectivities above the melting layer. The method is tested at short range ($< 30 \text{ km}$) only, where there is little difference between background and apparent VPRs, and the distinction between the two is neglected in this paper.

Kalogiros et al. (2013) combine dual polarisation information on bright band location with radar reflectivities to determine mean apparent VPRs. The authors use the profile of observed reflectivity along each ray to generate an apparent VPR, similarly to Zhang and Qi (2010), which is normalised by surface reflectivity. However unlike Zhang and Qi (2010), the profiles are also normalised in height using ρ_{hv} -determined bright band locations. Normalised VPRs from all rays are averaged in dBZ units to generate the mean field profile. This method is found to deal well with the bright band, reducing significantly both the bias and random error of surface rainfall estimates. However, as with other apparent VPR methods (eg Vignal et al., 1999), performance is found to degrade significantly for measurements beyond about 60 km in range.

A more complex linear VPR model is presented by Hazenberg et al. (2013). The authors first apply the convective and stratiform diagnosis algorithms of Steiner et al. (1995) and Sánchez-Diezma et al. (2000) to define contiguous rain segments as stratiform, convective or mixed precipitation. The height levels and gradients within the proposed ten-level stratiform VPR model (their figure 4), and the simpler four-level “mixed precipitation” profile, are constrained by theory and fitted to radar observations by a Monte Carlo optimisation approach. Although this method is both complex and computationally intensive, due in part to the number of parameters, the authors identify problems with fitting these profiles for rain segments at long range (beyond 100 km). The temporal averaging and aggregation proposed to mitigate these problems would seem to negate the benefits of the more complex parameterisation, which would be expected to deliver any impact through more precise and representative local detail in the determined vertical profile.

Determining mean field bright band parameters using ρ_{hv}

Tabary et al. (2006) investigate the use of dual polarisation parameters (Z_{DR} , ρ_{hv} and Φ_{dp}) to estimate the location and extent of the radar bright band. They find that the copolar correlation coefficient ρ_{hv} alone is the best indicator of bright band, as it is reliably greater than 0.98 in rain and snow, but falls to around 0.93 in the melting layer. On this basis an idealised linear vertical profile of ρ_{hv} is defined, and PPI measurements at a range of elevation angles are fitted iteratively to derive an average h_{top} and h_{bottom} across the radar domain. The Tabary et al. (2006) algorithm accounts for beam broadening, yielding a mean field background bright band profile, which assumes no variations in melting layer height across the radar domain.

Subsequently, several further VPR papers have used ρ_{hv} to identify properties of the melting layer. Kalogiros et al. (2013) define boundaries of the melting layer such that $\rho_{hv} \geq 0.97$ below h_{bottom} and $\rho_{hv} \geq 0.96$ above h_{top} , and require a melting layer detection at 40% of scan azimuths for a VPR correction to be applied. The method of Matrosov et al. (2007) uses thresholds of 0.95 to locate h_{bottom} and 0.9 for h_{top} , alongside constraints on the range-derivative along the radial (ρ_{hv} decreasing and increasing with height respectively).

Given the closeness of meteorological ρ_{hv} values to their theoretical detection limit, the exact thresholds used to identify rain, snow and melting snow in ρ_{hv} depend in part on the quality of polarisation separation in the radar hardware. This accounts for the minor variations in the literature, which are tuned to different radar networks and meteorological conditions.

2.3.4 Localised VPR methods

Although mean field VPR methods are by far the most common, various approaches have been tested to apply the VPR on a more local scale. Vignal et al. (1999) assess the effectiveness of the ratio method (Andrieu and Creutin, 1995a) over smaller domains, using ratios of hourly rainfall accumulations from all but the lowest elevation scan to estimate VPRs over 15° sectors of azimuth. Corrected reflectivities are then evaluated against measurements from the lowest elevation scan. This method is found to be most effective on domains spanning a radial extent of 30-40 km, over which the assumption of VPR homogeneity can be assumed to hold. However, Vignal et al. (1999) find this method cannot determine local VPRs beyond a range of 70-120 km, due to both decreasing vertical resolution (through beam broadening) and the lack of low-level profile data. Correcting long range (90-120 km) data with VPRs determined at short range was found to degrade the corresponding QPEs.

The retrieval of bright band parameters using ρ_{hv} is most frequently applied on mean field basis (Matrosov et al., 2007; Kalogiros et al., 2013; Qi et al., 2013b). This assumes no variation of the melting layer over regions of several hundred square kilometres, and can perform badly in frontal events, where the freezing level rises or drops sharply within the radar domain. Giangrande et al. (2008) refine this mean field approach, proposing a method to locate bright band boundaries at each azimuth of a single PPI. Melting layer ρ_{hv} values between 0.9 and 0.97 are used as part of a decision tree classification algorithm whose outputs include “ground clutter” and “bright band”. The algorithm makes use of data from multiple elevations, weighting higher elevation melting layer detections more highly due to their shorter range and sharper vertical resolution. There are stringent detection thresholds for bright band location, as well as the option to apply detections in a very few azimuths across larger sectors of the radar domain. Validation against temperature soundings finds a bias of -0.18 km and RMS error of 210 m in the melting layer top height from ρ_{hv} , which is sufficient for VPR corrections (Mittermaier and Illingworth, 2003).

Boodoo et al. (2010) adopt a similar approach to Giangrande et al. (2008) to determine azimuthally-varying melting layer boundaries. The authors set criteria to define melting pixels based on reflectivity and Z_{DR} as well as ρ_{hv} , and define h_{top} and h_{bottom} as the 80th and 20th percentiles of the heights of all “melting” pixels along each ray. This method is shown to be effective in capturing freezing level variations in frontal precipitation, compared with in-situ aircraft observations.

The most local approach to VPR available in the literature is that of Kitchen et al. (1994), whose linear parameterisation of the stratiform profile is described in section 2.3.3. The detailed parameterisation and use of climatological and ancillary data allows the background VPR to be estimated independently at each radar pixel, using an iterative deconvolution method initialised with reflectivity from a single scan. This method is refined by Kitchen (1997) to include freezing level estimates of h_{top} from numerical weather prediction models and an orographic enhancement component, and the correction is currently applied operationally in the UK (section 2.5).

Lewis et al. (2007) propose improvements to the Kitchen et al. (1994) profile by refining the estimate of reflectivity gradient above the melting layer. Whilst Kitchen et al. (1994) use only the lowest altitude reflectivity measurement to determine the profile at each point, Lewis et al. (2007) include information from higher elevation scans to calculate an average ice reflectivity gradient across continuous “rain segments”. Scovell et al. (2008) extend the evaluation of this method to rain rates, comparing hourly radar accumulations with rain gauge measurements over two seasonal periods. Overall, the multi-scan method reduces radar biases at long range, but is computationally expensive, and the benefits were found to be inconsistent over different radars and rainfall events. Therefore these

extensions were not implemented operationally.

Rico-Ramirez et al. (2005) use dual polarisation information to constrain the parameters of their linear VPR. A fuzzy logic algorithm (FLS) is first applied to classify hydrometeors in the radar volume as drizzle, rain, heavy rain, snow or melting snow. Membership functions based on reflectivity, Z_{DR} and LDR show skill in separating melting snow from other hydrometeor types, which is used to determine the bright band limits h_{top} and h_{bottom} for the scan. A mean field background VPR with the variable parameter Z_{bottom} is defined using the empirical constraints:

$$\nabla Z_{snow} = -5 \text{ dB km}^{-1} \quad (2.13)$$

$$Z_{top} = 1.05 + 0.85Z_{bottom} \quad (2.14)$$

$$Z_{peak} = 8.34 + 1.05Z_{bottom} \quad (2.15)$$

This is then fitted iteratively to observed reflectivities following the approach of Kitchen et al. (1994). Although the bright band height and depth are determined globally, in the absence of ancillary freezing level data, the use of an iterative deconvolution method allows local variations in bright band intensity to be captured in the applied VPR. This method is found to reduce surface reflectivity errors from around 7-8 dBZ to an average of 3.5 dBZ in the corrected data.

2.3.5 Combining apparent and parameterised VPRs

In recent years, some “hybrid” approaches to VPR have been proposed with the aim of combining the best aspects of the purely observational and *a priori* parameterised determination methods. These approaches use different weighted combinations of mean apparent VPRs and linear parameterisations depending on range from the radar, and on the quality and quantity of short range reflectivity observations available.

Snow et al. (2011) address the question of how to adapt VPR corrections to situations in which more or less real time radar data may be available. They consider the case in which a band of rain approaches the radar from maximum range, so that there is no short range data available to constrain lower levels of the VPR, and no recent data for temporal averaging. Snow et al. (2011) propose in this case to initialise a five-parameter mean field background VPR with climatological values, and to refine these values iteratively using observational data as the rain band progresses towards the radar. Like Kitchen (1997), the authors find that NWP model freezing level data is needed to determine radar bright band heights at long range.

Kirstetter et al. (2013) combine radar observations with theoretical knowledge of the microphysics underlying certain synoptic situations. The authors use estimates of drop

size distributions, dielectric properties (Szyrmer and Zawadzki, 1999) and particle fall speeds (Heyraud et al., 2008) to construct a simplified one-dimensional model of the VPR, expressed in terms of four variable parameters. These parameters are set initially to climatological values, from which they are evolved for subsequent time steps by assimilating reflectivity observations via an extended Kalman filter. The authors find this approach reproduces well the evolution of physical VPR parameters from ERA climate reanalyses, but do not investigate its skill in the context of instantaneous radar QPEs.

Koistinen and Pohjola (2014) approach the uncertainty in VPR determination through ensembles. A mean apparent VPR is calculated for each radar volume at ranges of up to 40 km. Like Andrieu and Creutin (1995a), the use of very short range data allows the authors to neglect the conceptual difference between the mean apparent and background VPRs, so that $z_a(h) \approx z_b(h)$. A time-weighted average observed VPR is calculated using the 24-member pseudo-ensemble of mean apparent VPRs from 15 minute intervals over the previous 6 hours. A second bias “ensemble mean” is generated from a linear climatological profile, which is constrained using a gridded NWP model freezing level field. The final VPR correction, applied during composite generation, is a spatially varying weighted average of observed and climatological bias estimates.

It has been suggested that NWP models could also be used to support operational determination of radar VPRs. Le Bastard et al. (2018) present an approach applying a radar forward operator to model fields to generate an ensemble of VPRs, and selecting the profile that most closely matches the observed apparent VPR. This approach has been shown to improve radar accumulations in a number of case studies, particularly at long range. However, there are uncertainties inherent in the NWP approach since the rain drop size distribution is not directly modelled, and the reflectivity must therefore be based on assumptions. This limitation of NWP models is likely to limit more complex approaches to QPE, for example based on direct assimilation of radar reflectivities, for the foreseeable future.

2.3.6 Time-averaged precipitation estimates

The majority of VPR literature focuses on obtaining the most accurate possible estimate of the instantaneous precipitation rate at each point in the radar domain. For this purpose, it is generally accepted that the lowest elevation PPI measurement that is free of clutter, once corrected, provides the most accurate estimate of rain rate at the surface. However, an alternative aim of VPR correction is to establish the average rain rate over the period of the radar volume scan. This approach is relevant to hydrological applications, where the desired product is a rainfall accumulation over time. From this alternative perspective, every PPI in a radar volume provides independent and useful

information as to the surface rain rate at different points in time.

Gabell et al. (2017) describe the approach of MeteoSwiss to obtaining an average QPE over a 5 minute period. Each scan is corrected independently using a non-parameteric mean field background VPR retrieved from data within 70 km of the radar (Germann and Joss, 2002). A weighted average of VPR-corrected measurements from all available radars and scan elevations is then used to obtain the final QPE at each point in the national composite.

There are some benefits to approaching QPEs as 5 minute averages rather than instantaneous rates. The use of data from scans separated by 1 minute or less ensures the representation of short-lived features in the QPE; temporal discontinuities between composites are reduced; and the averaging between scans and radars is designed to minimise random uncertainties. However, the fine scale variations captured by a 1 minute temporal resolution are not necessarily well-matched to a VPR correction determined over hundreds of square kilometres, and the need to weight heavily towards the lower scan elevations reflects the consensus that radar QPE errors increase extremely rapidly with height above the ground. In the absence of an objective comparison, it is not clear how much value is added by the inclusion of QPEs from higher elevation scans.

2.3.7 The question of scale

As is clear from the discussion above, the scale on which adjustments for VPR should be determined and applied is an unanswered question in the literature. By far the most common approach is to determine mean field VPRs using short range or idealised information, which are then applied to the whole radar domain. Mean field VPR corrections may be preceded by a classification step (section 2.4), and in some cases include different profiles for different rain types (eg Zhang and Qi, 2010; Hazenberg et al., 2013; Kirstetter et al., 2013). However, the underlying assumption is of spatial invariance of the background VPR over domains spanning many hundreds of square kilometres. Mean field VPR schemes are known to deal badly with frontal events, where improvement has been demonstrated using more local approaches to determine the melting layer profile (eg Giangrande et al., 2008).

Vignal et al. (2000) make a quantitative comparison between VPR corrections determined and applied on three different scales. The simplest is a climatological profile with a constant vertical reflectivity gradient of -1.5 dB km^{-1} . The second is a mean apparent VPR, determined from hourly profiles observed within 70 km range of the radar. Finally, the method of Vignal et al. (1999) is applied to determine “local” VPRs from ratio curves over 15° azimuth sectors. Temporal averaging on sub-hourly scales is shown to have negligible effect on the effectiveness of corrections at this spatial scale, and the

authors find that 90% of the bias reduction achieved by the “local” correction is also achieved by a mean field VPR. However, the scale of the “local” approach evaluated in this case still refers to regions up to several hundred square kilometres, and therefore lacks the fine detail necessary to capture truly local effects, such as rainfall-dependent variations in bright band intensity or embedded convection.

The pixelwise approach of Kitchen et al. (1994) goes to the opposite extreme. The authors observe the necessity for very local corrections (of a few square kilometres, at most) to capture small scale variability in the rainfall field, and to deal with variable melting layer heights in frontal events. Correlations between bright band intensity ΔZ and rainfall intensity (also observed by Hardaker et al. (1995) and Fabry and Zawadzki (1995)) lend weight to the argument that the VPR is not spatially homogeneous, and should be determined as locally as possible to account for these variations. The pixelwise convergence method is replicated in part by Rico-Ramirez et al. (2005), whose local approach has a fixed bright band height, but captures the variations in bright band intensity linked to surface rainfall rates.

Although accurate representation of small scale horizontal variations in the VPR is desirable, and intuitively should improve the resulting surface rainfall estimates, there arises the question of underdetermination at the local scale. A five-parameter linear VPR (eg Matrosov et al., 2007), for example, requires five independent observations to be uniquely constrained. The requirement for five independent observations based solely on meteorological radar data is unrealistic at the kilometre scale. This, along with computational cost, is a driving motivation behind the widespread use of mean field or domain-averaged VPR schemes.

In effective local schemes, underdetermination is addressed using a combination of linear parameterisations, climatological and ancillary data (eg Kitchen, 1997; Tabary, 2007; Koistinen and Pohjola, 2014). This use of alternative data sources can also reduce the uncertainty in cases where very few useful observations are available (Snow et al., 2011). The strength of such local schemes is evident in the magnitude of bias reduction in the resulting QPEs (Kitchen, 1997; Rico-Ramirez et al., 2005). However, there have so far been no studies addressing the question of whether local (kilometre-scale) VPR determination schemes achieve objectively better results than a wider spatio-temporal averaging approaches using radar observations alone.

2.4 Identifying non-bright band VPRs

The different microphysics underlying the formation of convective and stratiform precipitation generates vertical profiles with fundamentally different characteristics. Many of

the VPR correction schemes reviewed in section 2.3 are designed to work in stratiform conditions, and implicitly assume a parameterised shape that includes the bright band. A crucial prerequisite to these corrections is therefore to identify the type of precipitation observed in radar PPIs, in order to exclude non-stratiform VPRs from bright band correction schemes.

This section reviews the various schemes proposed to distinguish between stratiform bright band and convection in radar PPIs. Section 2.4.1 begins with the well established method of Steiner et al. (1995), based on spatial properties of the reflectivity field, which has formed the basis of a large amount of subsequent literature (eg Biggerstaff and Listemaa, 2000; Anagnostou, 2004; Rigo and Llasat, 2004; Delrieu et al., 2009). Other approaches using volume reflectivity measurements are discussed in section 2.4.2. Section 2.4.3 considers the use of dual polarisation in identifying localised convection, whilst section 2.4.4 takes the opposite approach of seeking positive identifications of the bright band.

2.4.1 Convective diagnosis in the SHY framework

A three-step process for identifying convection in radar PPIs is presented by Steiner et al. (1995). The algorithm (hereafter referred to as “SHY”) is applied to radar data on a Cartesian grid at 3 km altitude, which in the studied climate (Darwin, Australia) can be relied upon to be below the melting layer and free of bright band effects. The initial step is based on the premise that there is a limit to the precipitation intensity that can be generated by stratiform processes, and flags all pixels with reflectivities higher than 40 dBZ as convective. In the second step, a “peakedness” criterion is defined to describe the large horizontal reflectivity gradients observed in the vicinity of convective cores. If the reflectivity in a single pixel exceeds the average meteorological (non-clutter) reflectivity measured within an 11 km radius (Z_{bg}) by more than a certain threshold, it is classed as convective. The peakedness threshold depends on intensity as:

$$\Delta Z = \begin{cases} 10, & Z_{bg} < 10 \text{ dBZ} \\ 10 - Z_{bg}^2/180, & 10 \leq Z_{bg} < 42.43 \text{ dBZ} \\ 0, & Z_{bg} \geq 42.43 \text{ dBZ} \end{cases} \quad (2.16)$$

These peakedness and intensity criteria define convective cores.

The third step in the SHY algorithm is the extension of the convective classification from initial “cores” to surrounding regions. Although outside the strongest updrafts, these regions are also affected by vertical motions, including the trailing downdraft behind organised convective structures, and can support the formation by riming of smaller

hail and graupel. Steiner et al. (1995) propose a seeded growth method to classify these peripheral regions based on the maximum reflectivity of the convective core. The classification is extended to all pixels within an intensity-dependent radius of up to 5 km.

The SHY framework has formed the basis of a significant amount of subsequent literature, often seeking to improve on the original algorithm. Biggerstaff and Listemaa (2000) note that SHY has a tendency to misclassify heavy stratiform rain as convective, due to the low reflectivity intensity threshold of 40 dBZ, and to misclassify pixels on the edges of convective showers as stratiform. By using criteria based on horizontal and vertical reflectivity gradients, and the bright band fraction (defined as the proportion of the radar echo in which the column maximum reflectivity occurs within 1.5 km of the 0°C isotherm (Rosenfeld et al., 1995)), a second pass “reclassification” is applied to correct for these particular errors. Whilst the authors conclude, from sensitivity studies, that their modifications improve the stratiform-convective separation, this result is contested by Anagnostou (2004), which finds no statistical improvement by Biggerstaff and Listemaa (2000) over the original SHY algorithm.

The 40 dBZ intensity threshold set by Steiner et al. (1995) is generally found by later authors to be much too low. Implementation of the algorithm based on peakedness alone by Biggerstaff and Listemaa (2000) was found to reduce the incidence of false convective classifications. This finding is supported by Delrieu et al. (2009), who conclude that whilst the intensity and peakedness criteria work well together, the 40 dBZ threshold is responsible for a significant number of spurious convective detections which are then exacerbated by seeded growth. The authors therefore apply a modified version of the SHY algorithm, increasing the intensity threshold to 43 dBZ and removing the seeded growth step.

The variability in criteria and threshold values used to diagnose convection reflect the indirect nature of the link between reflectivity properties and the microphysics underlying convective profiles. While 2D reflectivity properties, gradients and textures can be a useful indicator of precipitation type, they are proxies, based on empirical studies of convective cells and bright bands. All three versions of the two-dimensional SHY algorithm (Steiner et al. (1995), Biggerstaff and Listemaa (2000) and Delrieu et al. (2009)) are equally valid, since there is no direct physical link between the threshold reflectivity value and the presence of convection. Such proxy algorithms are difficult to generalise, and must be carefully tuned to individual radar networks and climatologies.

2.4.2 Alternative approaches using volume reflectivities

Although extensively used, the SHY framework is not the only method available for convective classification. Anagnostou (2004) develops a neural network to classify echoes

for the USA WSR-88D (Weather-Surveillance Doppler Radar) system and scan strategy. Five predictors of convection are identified:

- F1: 15 dBZ echo top height
- F2: Reflectivity at 2 km altitude
- F3: Height difference between the 15 dBZ echo top and the reflectivity peak
- F4: Horizontal standard deviation (texture) of reflectivity
- F5: Vertical reflectivity gradient

A sixth predictor, the product of F1 and F2, is found to enhance the discrimination between stratiform and convective echoes. After training using data from the Tropical Rainfall Measurement Mission (TRMM) satellite, the algorithm was found to match the performance of both the original SHY and Biggerstaff and Listemaa (2000) reclassification methods in stratiform precipitation, and showed increased skill in correctly classifying convection.

Qi et al. (2013a) take a decision-tree approach to precipitation classification. Echoes at extremely short and long range are classified using a temperature-based reflectivity threshold, and additional information is incorporated at intermediate ranges. This includes the vertically integrated liquid water content (VIL), which is calculated as the sum of liquid water content (LWC) inferred from reflectivities at each scan elevation:

$$\text{LWC} \propto Z^{\frac{4}{7}} \times d_B \quad (2.17)$$

where d_B is the vertical extent of the radar beam. The use of VIL assumes a comprehensive scan strategy capturing the full extent of the reflectivity column. Similarly to SHY, convective cores identified by the initial decision tree are extended to neighbouring pixels with reflectivities exceeding 35 dBZ. The application of different reflectivity-rain rate relations to convective and stratiform echoes, as classified by this new scheme, produces radar accumulations in closer agreement with colocated rain gauge measurements than the previous classification algorithm (Zhang et al., 2008; Zhang and Qi, 2010).

The SHY approach itself can also be integrated into framework including other reflectivity parameters. Rigo and Llasat (2004) use a combination of two- and three-dimensional convective diagnosis algorithms to classify structures and events. Their 2D algorithm follows the same three-step framework as SHY, with an intensity threshold increased to 43 dBZ and a modified peakedness criterion. These structures are extended into three dimensions using the Storm Cell Identification and Tracking (SCIT) algorithm of Johnson et al. (1998).

2.4.3 Surface drop size distributions

Recent microphysical studies have investigated in depth the properties of different types of precipitation using dual polarisation parameters. In contrast to the traditional reliance on large scale proxies, the nature of dual polarisation parameters as independent *in situ* measurements allows some microphysical properties to be inferred directly at the local scale. One popular application is to use Z_{DR} measurements to constrain local rain drop size distributions (DSDs) near the surface, in order to select an appropriate rain rate estimator for QPE (eg Brandes et al., 2003).

The differences between convective and stratiform DSDs arise from the precipitation formation processes occurring at high levels in the vertical profile. The dominant ice level processes in convective events are condensation and riming. Since there is little aggregation, melting results in a relatively large number of small rain drops: so N_w is large and D_0 small. In stratiform cases hydrometeor growth is mainly through aggregation, increasing the size and reducing the number of drops. On this basis, Bringi et al. (2009) define and tune a parameter i to separate convective and stratiform precipitation types:

$$i = \log(N_w) - (6.3 - 1.6D_0) \quad (2.18)$$

Positive values of i indicate D_0 is small relative to N_w , and therefore correspond to convection. Conversely, negative i suggests D_0 is large relative to N_w , which is characteristic of stratiform regions. Bringi et al. (2009) also define a “transition” class, which includes decaying convective systems and stratiform precipitation in which a bright band has yet to form. This class is identified where $|i| < 0.1$. Bringi et al. (2009) develop polynomial fits based on disdrometer data to obtain D_0 as a function of Z_{DR} , and use a power law of the form $Z \propto D_0^C N_w$ to estimate N_w at each pixel.

Penide et al. (2013) evaluate the relative skill of the Bringi et al. (2009) (hereafter BAL) and reflectivity-based SHY algorithms. Whilst the two methods show good agreement on stratiform classifications, only 59% of SHY convective pixels are classified as convective by BAL; the remainder are stratiform or transition. The transition class identified by BAL shows distinctly separate characteristics from stratiform or convective classes, with intermediate D_0 , reflectivities and rain rates, and low N_w . The BAL convective and stratiform distributions are much more clearly separated than the SHY cases, and show better agreement with the expected microphysics. On this basis the authors consider BAL the better algorithm, and suggest refinements to the SHY peakedness criterion to improve the agreement between the two methods.

Dual polarisation microphysical methods show promise for investigating rain drop size distributions, and particularly for modifying previously fixed ZR relationships to improve

rainfall estimation. Applying the method of Bringi et al. (2009), Thompson et al. (2014) find that a climatological ZR relationship (equation 1.14) which does not account for precipitation type can overestimate stratiform and underestimate convective QPEs by 50% and 15% respectively. However, DSD methods such as BAL are of limited usefulness for classification as a prerequisite to VPR. Since these methods currently only apply to rain pixels, they cannot classify measurements in and above the freezing level: which are precisely the regions in which VPR classification and correction is required.

2.4.4 Bright band identification

Sections 2.4.1-2.4.3 have described a variety of approaches to identify convective cores based on their observable characteristics. The underlying assumption has tended to be that all non-convective precipitation is stratiform, so that identification of stratiform precipitation is by “default” implication only. However it is also possible to make positive identifications of stratiform regions, specifically where there are indications of a strong bright band.

Sánchez-Diezma et al. (2000) develop a two-step procedure to identify bright band signatures in radar PPIs. The first step locates obvious peaks with the condition that the maximum measured reflectivity must exceed reflectivity in the rain below by at least 5 dBZ. The heights of these initial peaks are averaged across all detections to determine a single “bright band height”. A second pass then sets a less stringent reflectivity peak threshold of 2 dBZ, but requires that this reflectivity be measured within a certain distance of the established bright band height. The authors acknowledge that their method assumes a comprehensive scan strategy, in which the bright band peak is sampled by at least one radar beam at all ranges, and that therefore in reality would likely miss a significant proportion of radar-observed bright bands. By simulating the radar response to an idealised bright band VPR, Sánchez-Diezma et al. (2000) also determine that a typical bright band would not be resolved by a 1° radar beam at ranges exceeding 70 km.

The justification of Steiner et al. (1995) and others in using horizontal reflectivity structure as a framework for convective diagnosis is based on this difficulty in detecting stratiform bright bands at long range. Smoothing of the reflectivity peak with range is identified as a major limitation of a bright band approach, and Steiner et al. (1995) also emphasise that strong reflectivity bright bands are often not measurable until the stratiform system is well developed. With the advent of dual polarisation measurements, however, measurements of reflectivity degraded by beam broadening are no longer the only available option.

Recent papers have shown that the copolar correlation coefficient (ρ_{hv}) can be used to

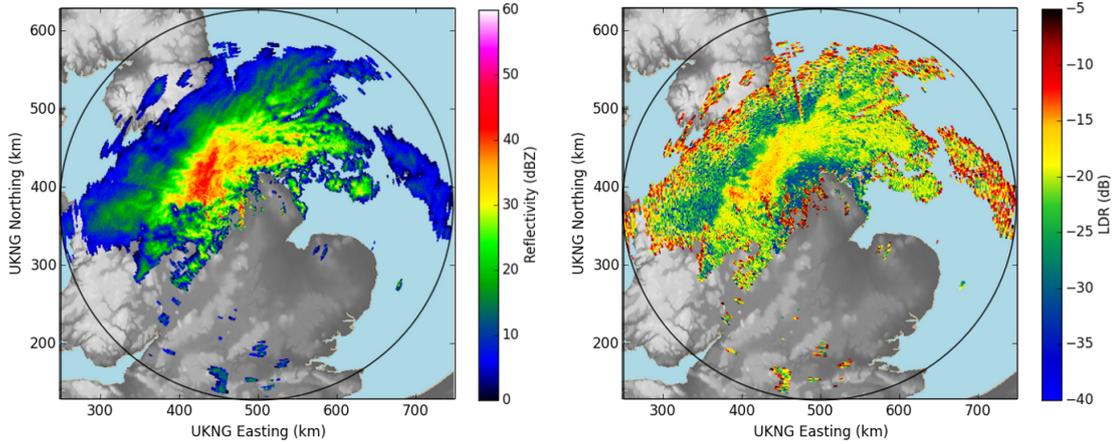


Figure 2.2: Example reflectivity (left) and LDR (right) 0.5° elevation PPIs from Ingham, 21st November 2016 17:09 UTC, with a maximum range of 250 km. The clear bright band in LDR allows high reflectivity values around the radar to be correctly attributed either to bright band (eg North-West of the radar) or to heavy rainfall (eg North and slightly East of the radar, at close range). The high values of LDR at long range are due to low signal-to-noise ratio. This figure has previously been published as figure 2 in Sandford et al. (2017).

locate the melting layer in stratiform rainfall (eg Tabary et al., 2006; Matrosov et al., 2007; Giangrande et al., 2008; Boodoo et al., 2010). The increased variability of hydrometeor sizes, shapes and orientations within the radar pulse volume due to melting is associated with a significant reduction in ρ_{hv} . However, the melting hail and graupel characteristic of convective melting layers also cause a similar reduction in ρ_{hv} . There are no published results to suggest that ρ_{hv} would be significantly different in stratiform bright bands from convective melting layers, where no bright band is present.

Smyth and Illingworth (1998) also propose using dual polarisation parameters to distinguish between stratiform bright band and convective precipitation. Using RHI scans from the high resolution S-band radar at Chilbolton, they show that snow above the melting layer can be distinguished from graupel using a combination of reflectivity and LDR thresholds (section 1.4.3). On this basis they define stratiform precipitation as having $LDR \geq -18$ dB over widespread regions, corresponding to the bright band as sampled by the radar PPI. This high measured LDR occurs as a result of the strong depolarisation properties of large melting snowflakes (section 1.4.3), and is not observed in convective melting layers (Illingworth and Thompson, 2011). A major strength of LDR for bright band detection is that, since an LDR measurement is strongly dominated by the maximum depolarised reflectivity return, long range LDR peaks are much less affected by beam broadening than are reflectivities (eg figure 2.2).

2.4.5 Combining classification algorithms

Although discussed separately in this section, precipitation classification techniques can be used in combination to provide increased accuracy and confidence. In their study of quantitative precipitation estimation in complex terrain, Delrieu et al. (2009) apply a modified version of the SHY algorithm to diagnose convection, followed by the criteria of Sánchez-Diezma et al. (2000) to identify stratiform regions. Remaining areas are classified as “mixed/transition” regions. Hazenberg et al. (2013) adopt the same approach in their development of linear parameterised VPRs. The use of multiple algorithms for VPR type classification tends to be confined to situations pre-dating the widespread availability of dual polarisation.

2.5 Operational VPR correction schemes

Sections 2.3 and 2.4 described the wide range of VPR classification and correction schemes available from over 20 years of theoretical literature. This section revisits in further detail three complete operational schemes as documented in the literature. The importance of operational reliability and resilience are emphasised, along with the need to pair classification and correction schemes that operate on a similar local or global scale. The aim here is to highlight the very different approaches to correction followed by different centres, and to place in perspective and highlight the unique nature of the UK operational VPR procedure, in which context the remainder of this thesis will be based.

2.5.1 Météo France

Tabary (2007) propose a modified version of the ratio method of Andrieu and Creutin (1995a) for application to the French operational radar network. They formulate the VPR as a ratio of Cartesian-gridded rainfall rates, as:

$$R(h, \mathbf{x}) = R_s(\mathbf{x})z_a(h) \quad (2.19)$$

where R is the rain rate calculated from a reflectivity at height h , R_s the true surface rain rate, and $z_a(h)$ is therefore the apparent VPR (cf equation 2.6), modified by a Marshall-Palmer equivalent ZR exponent (equation 1.14). As part of the determination process, “ratio curves” as a function of range are calculated using ratios of hourly rainfall accumulations from different elevation scans.

The method of Tabary (2007) relies upon the generation of a family of *a priori* linear VPRs. Their generalised stratiform profile consists of a symmetrical bright band peak

(as for Kitchen et al. (1994), figure 2.3), constant reflectivity below the melting layer, and a constant dBZ reflectivity gradient in the ice. The profile has four variable parameters, each of which has only a limited range of permitted values for computational efficiency:

- The bright band top h_{top} can be located at the model freezing level height h_{fl} , or at either of $h_{fl} \pm 200$ m
- The bright band peak strength (as a rain rate ratio) can take any integer value from 1 to 5 (corresponding to a maximum reflectivity enhancement of around 11 dB (Tabary et al., 2007))
- The bright band depth ($h_{top} - h_{bottom}$) is 200 m, 400 m, 600 m or 800 m
- The gradient above h_{top} can take any of five values between -1.5 and -6 dB km⁻¹

(see annotated VPR sketches, figure 2.1). This gives a total of 240 *a priori* linear profiles. The best of these linear profiles for each hourly time step is chosen via a least squares optimisation using the measured ratio curves, and is applied to Cartesian rain rate accumulations. The final surface rain rate is computed as a weighted linear combination of all values available from different elevation scans (Tabary et al., 2007).

The Météo France operational scheme is global in nature, assuming no change in the background VPR over the radar domain and hourly timescales. This approach can be problematic in frontal situations, where the height of the freezing level changes during the event (Tabary et al., 2007), and the authors acknowledge other potential weaknesses in cases of low level growth due to orography (Tabary, 2007). No explicit attempt is made to identify or exclude convection at the local scale, although it is possible for the domain averaged VPR to be “non-bright band” to the extent that the bright band peak strength can take a value of 1 (no enhancement).

Tabary et al. (2007) perform a detailed evaluation of the VPR scheme described in their earlier paper. The implementation details are slightly different in this validation, with a larger number of permitted values for each of the four variable parameters, and in particular h_{top} is not constrained by the model-predicted freezing level. This evaluation is also run over daily, rather than hourly accumulations, in order to facilitate processing of a large historical dataset. This thorough climatological validation showed significant skill in reducing systematic accumulation biases on the daily timescale. However, no evaluation of instantaneous QPEs or hourly accumulations at the local scale have been published at this time.

2.5.2 Finnish Meteorological Institute (FMI)

An operational approach to VPR correction in Finland has recently been developed by Koistinen and Pohjola (2014). Unlike most regions, in which the dominant VPR bias error is overestimation due to bright band, at FMI the purpose of VPR is largely to correct for the underestimation of precipitation rates from measurements significantly above the melting layer. The high latitudes covered by the Finnish radar network result in an observed pre-VPR reflectivity underestimation beyond 100 km of between 4 and 15 dB.

The method of Koistinen and Pohjola (2014) uses both parameterised climatological profiles and radar measurements to generate a weighted ensemble of possible VPR bias corrections at each point in the radar domain. This global approach does not explicitly classify VPRs as stratiform or convective, but moulds the applied profile shape to the observed data. First, a mean field apparent VPR is determined for each radar using multi-elevation reflectivity ratios at 2-40 km range. This VPR is used to estimate the bias correction that should be applied to longer range measurements, in a way that accounts for additional broadening of the radar beam beyond the 40 km determination limit. At any given point in the radar domain, the 0-24 member ensemble of measured VPR biases is made up of the bias ratio calculated at this grid point at 15 minute intervals over the previous 6 hours.

The ensemble of parameterised climatological VPRs also has 24 members. The profile shape is similar to that of Kitchen et al. (1994), with a symmetrical bright band 800 m deep, whose top is located using a gridded NWP model freezing level. The bright band intensity ΔZ is fixed at 7 dB, and the profile has climatological reflectivity gradients below and above the bright band. The different ensemble members are determined by the height of the NWP model freezing level at each point, at 15 minute intervals over the previous 6 hours, with the climatological gradients also varying depending on whether precipitation at the ground is rain, wet or dry snow.

Once both sets of ensembles have been calculated, the bias correction for each point in the radar domain is established as a time- and quality-weighted mean of the two ensembles, with more recent ensemble members being given a greater weight. The correction applied to the composite is then a distance-weighted mean of the bias corrections from each radar within 300 km of the composite grid point, which is applied during compositing to the reflectivity measurement from the closest available radar. This means that although based on spatially averaged observations for each radar domain, the VPR correction is applied in a way that is tuned locally for each radar composite grid point. The authors find this scheme reduces the mean VPR underestimation bias to less than 2 dB for measurements within 200 km of the nearest radar.

Like many authors, Koistinen and Pohjola (2014) begin from the premise that mean apparent VPRs provide better information as to the true background VPR than convergence with an idealised profile. From this perspective, the use of short range radar-determined ratios is taken as read; but the authors explicitly justify the addition of climatological profiles and temporal smoothing in terms of reducing the random variability apparent in VPRs determined purely through observed reflectivity ratios. Koistinen and Pohjola (2014) go further than most in making no attempt to determine a separate VPR for convective precipitation, or to exclude convective regions from correction using a bright band profile. The authors justify this choice on the grounds that unlike bright band, the dominant underestimation bias at long range is equally a problem for both stratiform and convective rainfall. It is acknowledged in the paper that the inappropriate correction for bright band can cause problems in short to medium range convective QPEs.

2.5.3 Met Office (UK)

The Met Office radar processing software (Radarnet) implements a pixel-by-pixel VPR scheme originally developed by Kitchen et al. (1994), and refined by Kitchen (1997). The mean stratiform profile shape (figure 2.3) was derived from a three year climatological sample of high resolution range height indicator scans (RHIs) observed with the 25 m S-band dish at Chilbolton. The profile has a fixed bright band depth of 700 m, and uses the Euro4 forecast model wet bulb freezing level (Brown et al., 2012) to define the top of the bright band. Mittermaier and Illingworth (2003) compared the forecast freezing level height with observations of the melting layer top from a vertically pointing radar, and found an RMS error of less than 150 m, confirming that the model height is sufficiently accurate for use in VPR correction. A single variable parameter in reflectivity is used to scale the idealised profile to the measured reflectivity at each radar pixel, using a known beam power profile to simulate the observed reflectivity measurement and adjusting the variable scaling parameter until the simulated reflectivity matches the observation. The surface reflectivity is then

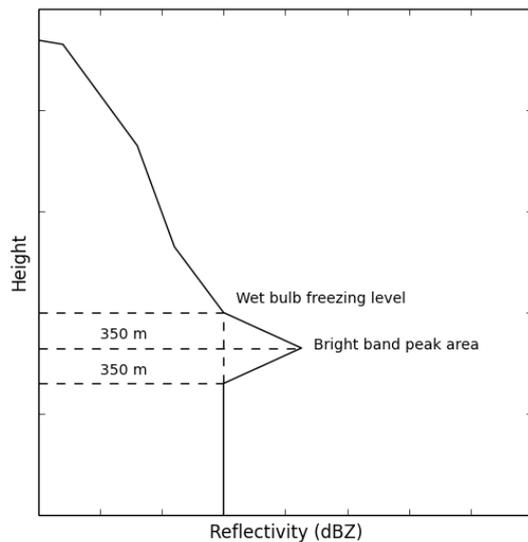


Figure 2.3: The idealised stratiform VPR shape, derived by Kitchen et al. (1994), which is used operationally in the Met Office centralised radar processing system (Radarnet). The wet bulb freezing level is derived from the 5 km gridded operational forecast model output.

extrapolated from the fitted profile.

A significant strength of the Met Office scheme is its ability to account for sub-kilometre-scale variability such as changes in bright band height and intensity and the presence of embedded convection, as it responds to local conditions at the radar radial resolution (600 m for the standard UK QPE) along each azimuth. This allows additional local information, such as a change in freezing level height with frontal passage or low level orographic precipitation growth, to be included in the correction (Kitchen, 1997). For the stratiform profile in its current form, Kitchen et al. (1994) demonstrate a 60% overall reduction in QPE error for a number of light stratiform cases, and emphasise that greater gains would be expected in heavier frontal rain. The overall design of the iterative convergence scheme also provides a flexible framework for assimilating additional information, which could be used to adapt or update idealised parameterised profiles for further improvements to real time QPEs.

The Kitchen et al. (1994) idealised stratiform profile is not suited to cases without bright band, such as occurs for example in embedded convection with graupel. The underestimation in surface rainfall caused by erroneous bright band correction disproportionately affects estimates of the intense, often flood-producing rainfall associated with convective cores. It is therefore important to identify where these profiles occur, to avoid errors in high impact situations.

The current UK VPR scheme uses a high level reflectivity threshold to identify local profiles without bright band in radar data. If a reflectivity exceeding 30 dBZ is measured at a height exceeding 1 km above the wet bulb freezing level (criterion hereafter referred to as Z_1), the pixel is classed as convective, and the VPR at that pixel is set to be constant with height. This draws on the assumption that high reflectivities above the zero degree isotherm can proxy for the strong updrafts associated with convection and non-bright band VPRs (Smyth and Illingworth, 1998).

2.6 Summary

The difficulty of producing accurate QPEs from long range radar measurements is an established and ongoing challenge for the radar community. The nature of the vertical profile of reflectivity (VPR) causes both positive and negative biases in uncorrected data, from order of magnitude enhancements in the bright band to the much lower reflectivities of snow compared to rain. This problem is compounded by broadening of the radar beam, which smears out fine vertical structures in radar PPIs and prevents direct determination of the VPR for correction in real time.

The science of stratiform precipitation, with a particular focus on the bright band, has

been the subject of many studies. The layered structure of cold stratiform rain lends itself well to microphysical modelling (eg Hardaker et al., 1995; Szyrmer and Zawadzki, 1999; Fabry and Szyrmer, 1999; Zawadzki et al., 2005), and the tendency towards spatial homogeneity in these layers has led to much research into correction methods based on large scale averaging. However, modelling of the bright band itself is extremely sensitive to the underlying microphysics, with small inaccuracies in model assumptions leading to large differences between modelled profiles and observations (eg Hardaker et al., 1995).

Given the complexities of modelling, a large body of research has focused on deriving the VPR from observations. Purely empirical ratio-based methods (Andrieu and Creutin, 1995a; Vignal et al., 1999) to determine domain-averaged stratiform VPRs over hourly periods have evolved into more complex linear and microphysical parameterisations (Rico-Ramirez et al., 2005; Matrosov et al., 2007; Tabary, 2007; Kirstetter et al., 2013), which have been shown to improve QPE statistics on daily and climatological timescales. Yet the majority of these methods, however complex the implementation, rely on domain- and time-averaged radar volume parameters to produce what is essentially a global correction for bias with range.

There are issues with correcting globally for a VPR which in reality varies, both temporally and over the radar domain. Tabary (2007) acknowledge the difficulties their method has in dealing with frontal passages, where the height of the bright band changes abruptly within a radar scan. In these cases a domain-averaged VPR can place the bright band incorrectly at ALL points, leading to worse QPEs than those obtained from uncorrected data. Local VPR corrections can provide better responsivity to fronts, variations in bright band intensity, or the presence of embedded convection. Consideration of local features can also provide the opportunity for refining the profile shape, such as the predictive relationship discovered by Kitchen et al. (1994) between surface reflectivity and bright band intensity at each pixel. Such refinements improve determination and correction for VPR, and hence the accuracy of surface rainfall estimates, at the local scale.

The use of a locally adaptive VPR presents both challenges and opportunities. Since beam broadening renders vertical structures unresolvable beyond very short range, radar observations alone do not provide sufficient information to determine the VPR at the pixel scale. This results in the need for ancillary and climatological information to constrain local profiles. However, the flexible nature of some local parameterisations allows new sources of information to be incorporated easily into these existing frameworks. Thus such frameworks and parameterisations are easy to refine, and even simple approximations have the potential to exceed the performance of more complex global schemes.

The advent of dual polarisation as the new standard for operational radar networks has provided an abundant new source of information on the local scale. In the context

of VPR, a promising area for development is the use of the linear depolarisation ratio (LDR) to distinguish between different types of profile. The potential of LDR to identify widespread regions affected by bright band has been acknowledged since the 1990s (Smyth and Illingworth, 1998), and more recently has been shown to have promise at a more local level (Illingworth and Thompson, 2011). However, this knowledge has yet to be applied towards improving VPR or bright band corrections.

This PhD aims to use dual polarisation information to improve local profile classification and correction in the context of the UK operational VPR scheme. The Kitchen et al. (1994) framework already makes use of information as to the stratiform or convective nature of the local profile, which is currently provided by a proxy indicator that is known to underdiagnose convection. The use of LDR to improve the separation of different profile types could significantly improve corrected QPEs in non-stratiform regions. Improved classification could also provide opportunities to refine details of the idealised profiles used for correction, through observing patterns in different precipitation types. Such refinements would be facilitated by the more accurate discretisation of vertical structures, whose features are both similar to each other and distinct from those of other groups.

Chapter 3

A high resolution VPR dataset for testing and verification

The aim of this PhD is to improve local corrections for VPR in the context of the recently upgraded Met Office dual polarisation radar network. To do this, it was first necessary to develop a climatologically representative dataset of observed VPRs. This dataset will initially support the evaluation of new VPR classification methods using the linear depolarisation ratio (chapter 4), and will later provide “ground truth” for simulation studies of improvements to the VPR correction algorithm (chapters 5-7). This short chapter describes the radar hardware, data collection, quality control and calibration methods used to develop a high resolution VPR dataset for this thesis.

3.1 The Wardon Hill research radar

The data for this thesis were provided by the Met Office C-band research radar at Wardon Hill (figure 3.1). The radar is located in the South West of England, Devon (UK National Grid Easting 360907, Northing 102337), at an altitude of 242 m. Due to surrounding trees and the higher altitude of nearby topography, the radar tower is relatively tall, at 12.4 m, so that the antenna (accounting for dish diameter) is centred at 256 m AMSL.

The Wardon Hill radar was built as a prototype for the upgraded Met Office dual polarisation radars. The hardware is based on a Plessey Systems type 45C pedestal, which was refurbished and upgraded to dual polarisation using the method of splitting transmitter power through the single existing waveguide (Darlington et al., 2016). The radar operates a 250 kW magnetron transmitter at a C-band wavelength of 5.3 cm. Scans are taken either at long pulse, with a $2 \mu\text{s}$ (300 m) pulse length and 300 Hz pulse repetition

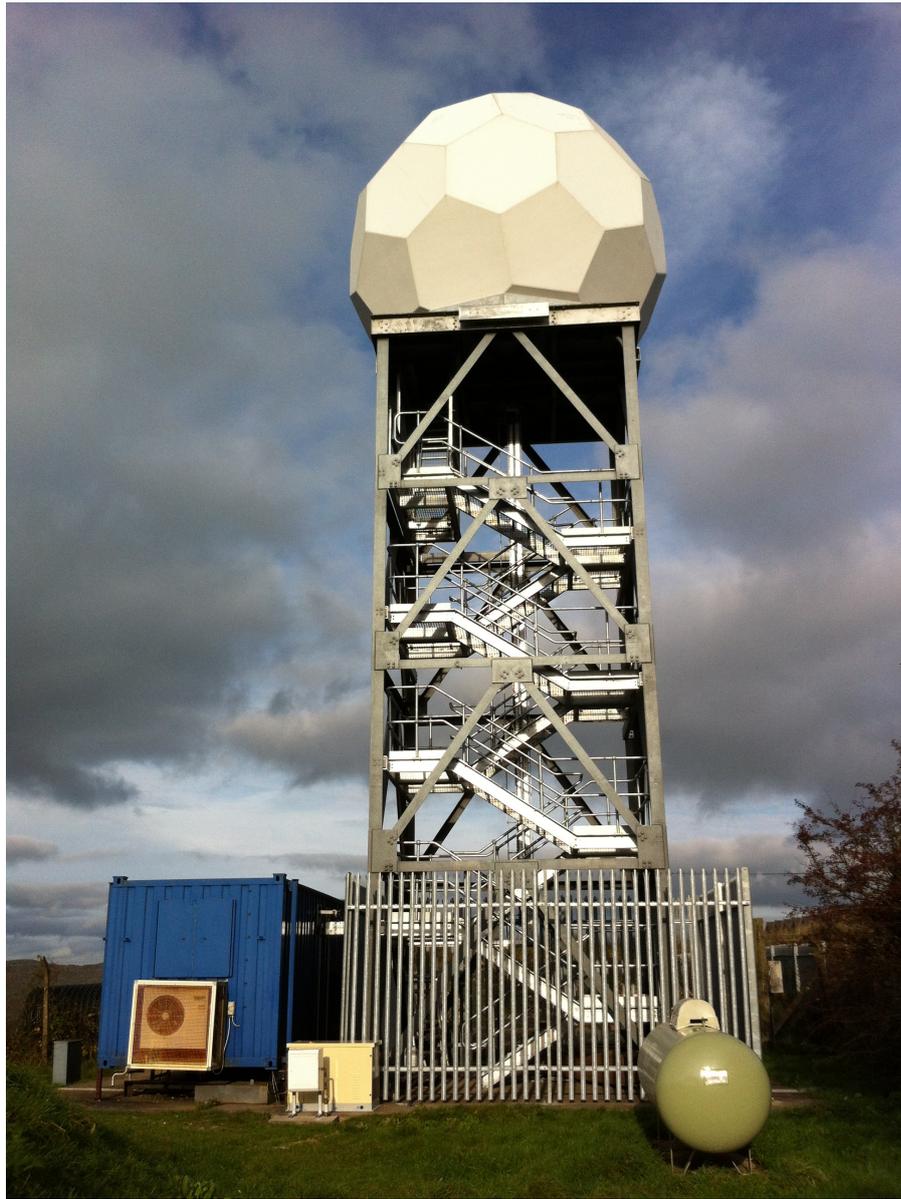


Figure 3.1: The 12.4 m radar tower at Wardon Hill, WMO site number 03854.

frequency, or short pulse ($0.5 \mu\text{s} / 75 \text{ m}$), where the short pulse “Doppler” data alternate between PRFs of 900 and 1200 Hz every 8 pulses (Darlington, 2014).

Like all radars in the Met Office network, Wardon Hill is capable of scanning in two dual polarisation modes. SHV refers to simultaneous transmission and reception in the horizontal and vertical channels, and is used throughout the Met Office operational network to supply real time data for QPE (see section 1.4). Additional LDR mode scans transmit in the horizontal polarisation only, but receive in both H and V channels.

3.2 Data collection strategy

The majority of operational radar scans are taken as polar “plan position indicators” (PPIs) in SHV mode. This is for two main reasons: both to maximise spatial coverage, and to minimise the strain on the drive system caused by frequent changes in antenna rotation. However, the availability of a dedicated research radar at Wardon Hill provides opportunities for investigations using a more flexible scan strategy.

To examine the vertical structure of precipitation, an experimental LDR mode range height indicator (RHI) scan was incorporated into the continuous strategy at Wardon Hill. RHI data were collected by scanning down from 90° to 0° in elevation at a fixed azimuth every 10 minutes, when the radar moves down from the zenith after obtaining a vertically pointing scan. Each RHI therefore provides meteorological data for a vertical slice through the atmosphere, resolving vertical structures that cannot be captured by PPI measurements.

The data available from each RHI consist of the horizontally polarised reflectivity return (Z), linear depolarisation ratio (LDR), and clutter phase alignment (CPA) (Hubbert et al., 2009). CPA is a measure of the phase stability of samples in the radar pulse volume, and is significantly higher for stationary clutter than precipitation echoes (section 3.4.1). An archive of RHI data was built up over two study periods - from late September to November 2014, and from April to August 2015 - which ensured representation in the dataset of the range of different seasonal rainfall types typically observed in the UK.

3.2.1 Scan details

Wardon Hill RHIs were taken in short pulse dual PRF Doppler mode (section 3.1), with a radial resolution of either 300 m or 75 m, to a maximum range of 60 km from the radar. The difference between these data is that the 300 m scans were averaged at site, with values at each range gate calculated using pulses from four adjacent 75 m bins. The duration of each scan was 17 seconds, with an average slewing rate in elevation of 5.3° per second, and azimuths were sampled at random over the study period.

Due to the experimental nature of this scan and limitations of the on-site driver software, it was not possible to record the actual azimuths at which these RHIs were taken. However, for the purposes of this investigation there was no need to compare these data with any fixed geographical reference. Although the ability to compare RHI data with observations from PPIs could potentially have opened up new lines of enquiry and cross-validation, RHI scan data without precise geographical reference can still be used to build up climatologies and investigate the nature of VPR shapes in a statistical manner.

3.2.2 Quality of the LDR measurement

Since this thesis is largely invested in evaluating the deterministic properties of LDR, the quality of the LDR measurement in this dataset is crucial. LDR is defined as the fraction of a horizontally polarised reflectivity transmission that is returned in the vertical polarisation (equation 1.22). Given a depolarised echo component typically 2-4 orders of magnitude smaller than the background reflectivity, LDR is extremely sensitive to the quality of the radar hardware, as the signal is easily contaminated by cross polar noise.

The cross polar isolation at the Wardon Hill radar, determined using the median LDR measurement in light rain (20-25 dBZ), is -36 dB. Light rain at close to the limit of the radar's detectability can be used to obtain empirical estimates of cross polar isolation, because the very small rain drops that contribute to such low reflectivities are close to spherical, and therefore should not depolarise the incident wave. The estimate is not a perfect measurement of isolation, since the drops are not perfectly spherical, but gives a close approximation of the minimum detectable LDR. The Wardon Hill value of -36 dB is typical of UK network radars upgraded to dual polarisation through the recent Weather Radar Network Renewal project (appendix A).

Given a noise threshold of -20 dBZ at 15 km range, a minimum meteorological reflectivity threshold of 10 dBZ would correspond to a minimum signal-to-noise ratio (SNR) of 30 dB. Since this is much higher than the cross polar isolation, the impact of cross polar noise can be considered negligible in terms of the random error on LDR. Extending this sensitivity to a noise threshold -3 dBZ at 100km range, the SNR in the co-polar channel for a reflectivity of 23 dBZ (equivalent to a rain rate of about 1 mm h^{-1}) would be 26 dB; so an LDR of -26 dB can be detected at 100 km. For more significant rain rates of 3 mm h^{-1} (about 31 dBZ), this LDR sensitivity would be achieved out to a range of 250 km.

The random error on LDR is related to the random error on reflectivity discussed in chapter 1. For a total number of pulses $M_t = 37$ (section 1.3.2), the random error on reflectivity is approximately 0.8 dB (Doviak and Zrnić, 1993, figure 6.2). Since LDR is a ratio of reflectivities received in the two channels the random error is twice that of reflectivity, or 1.6 dB. It will be shown in chapters 4 and 5 that these LDR measurements are of sufficient accuracy and precision to support the proposed application.

3.3 Calibration

As a test and development radar, the calibration status of Wardon Hill is not expected to meet the standards required of the operational network. In particular, over both RHI

study periods the radar was running in low power mode without the corresponding on-site reflectivity calibration. This meant a significant reflectivity offset expected in the raw incoming data. After data collection, therefore, calibration factors for both reflectivity and LDR were determined and applied retrospectively to each RHI. This section describes the methods used to determine suitable calibration factors, and places their uncertainties in context for this thesis.

3.3.1 Reflectivity

Reflectivity calibration factors for the Wardon Hill were calculated using a self-consistency relation. Self-consistency methods exploit the redundancy of Z , Z_{DR} and KDP measurements in rain to derive a reflectivity offset based on the agreement between theoretical and measured values (Goddard et al., 1994). The Radarnet auto-calibration scheme compares a theoretical definition of KDP with the measured gradient of differential phase shift (Φ_{dp}) to determine a multiplicative calibration factor for reflectivity (Adams, 2015; Gourley et al., 2009):

$$\text{KDP} = 10^{-5} Z (6.746 - 2.97 Z_{DR} + 0.711 Z_{DR}^2 - 0.079 Z_{DR}^3) \quad (3.1)$$

The comparison is made on a per-ray basis, using KDP calculated from the first 25 range bins (15 km) where the total Φ_{dp} over this range is between 8 and 12° (Adams 2017, personal communication). If five or more ray calibration values are available, these are averaged to derive a calibration factor for the whole scan.

This auto-calibration scheme is not currently applied operationally. The data quality requirements for a stable offset calculation are stringent, including precise calibration of Z_{DR} and high enough Φ_{dp} for KDP to be calculated (since Φ_{dp} is noisy), but not so high that attenuation or differential attenuation would significantly impact the results. There must be enough rain close to the radar to calculate the offset for multiple rays, so that the average is stable; but not over the radar itself (Adams 2017, personal communication), since during rain events, water running off the radome causes additional attenuation that impacts equation 3.1. This “wet radome” attenuation is not easily quantifiable, and cannot be completely avoided even through the best hydrophobic radome design; and the effect increases as the hydrophobic coating degrades over the lifetime of the radar. In practise, this means that good quality calibration factors cannot be calculated frequently or reliably enough to inform an operational adjustment scheme. The reflectivity calibration offset is therefore calculated for monitoring purposes only, as part of the Radar Data Quality Management System (RDQMS).

For the Wardon Hill RHI dataset, reflectivity calibration factors were calculated using SHV mode PPIs for scans matching the two study periods. The initial auto-calibration

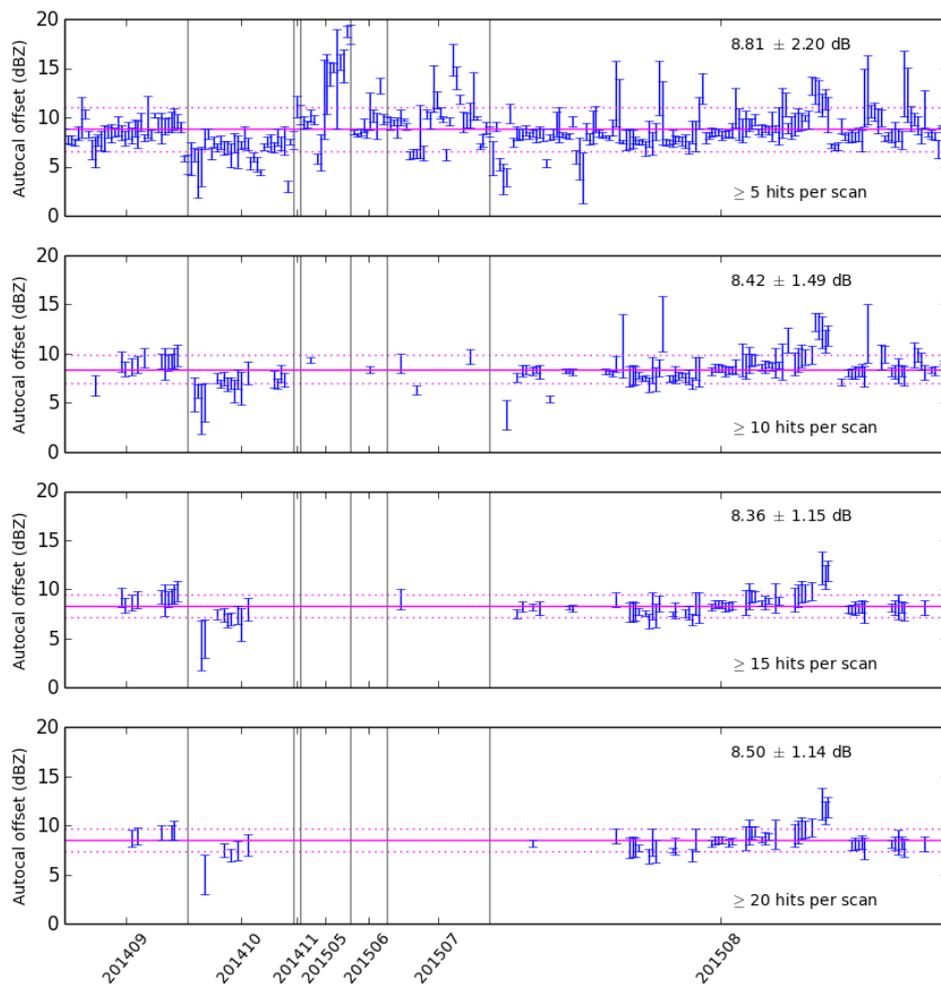


Figure 3.2: Reflectivity calibration offsets calculated using the self-consistency scheme for Wardon Hill data, during the study period over which RHIs were collected. The four plots represent different requirements, in terms of “hits” per scan, to accept a given offset value. Grey vertical lines delineate months over which calibration hits were calculated. The different window lengths are indicative of the number of RHI scans included in the sample over each month.

outputs, requiring a minimum of five values per scan, are shown in the top panel of figure 3.2. Even with the stringent data quality requirements, the variability in the calculated offset within and between SHV scans is extremely high, providing further justification for the decision not to apply this scheme in real time.

For the purposes of this thesis, it was decided that since no on-site calibration adjustments were performed at Wardon Hill during the study periods, a single reflectivity offset could be applied for the whole year. This allowed the available data to be further reduced to improve the quality of the calibration, by requiring a larger number of valid rays and better agreement to accept the calibration factor from any given scan. The lower panels of figure 3.2 show the reflectivity offsets calculated using successively higher numbers of ray values per scan. On the basis of these data, the final offset applied to the RHI dataset was 8.5 dB, with a nominal uncertainty of up to 1.5 dB.

This method of calculating reflectivity offsets clearly has substantial uncertainty, and some judgement has necessarily been applied in choosing a final value for calibration. However, the absolute calibration of reflectivity in fact has very little bearing on the investigations performed in this thesis. In evaluating the behaviour and skill of LDR in classifying VPRs, the reflectivity parameters used are restricted to dBZ differences between various height levels in the profile, which are insensitive to a multiplicative calibration offset. The absolute value of reflectivity is used only in comparing the skill of LDR with that of high-level reflectivity in classifying VPRs (chapter 4). Since a range of reflectivity threshold values are used to support the analysis, it can be stated with confidence that the uncertainty in absolute reflectivity calibration does not affect the conclusions of any part of this thesis.

3.3.2 LDR

Measurement bias in LDR results from sensitivity differences of the radar receiver in the H and V channels. This offset can be measured using interference from sun detections (Frech et al., 2017). Radiation emitted from the sun and detected by the radar is completely depolarised, so sun measurements should have equal power in all polarisations. Any difference in sun reflectivity between the H and V channels is therefore directly attributable to differences in sensitivity between the receivers, and can be used to calculate a calibration offset for LDR.

Long range noise emissions are collected from Met Office radars in SHV mode as part of the operational PPI scan strategy. The sun appears in these emissions data as a completely depolarised target, with an intrinsic LDR value of zero. These emissions therefore provide a direct measurement of any system offset in LDR.

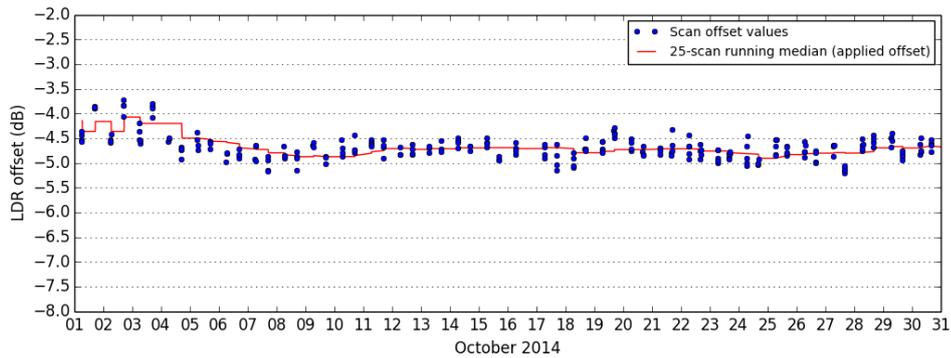


Figure 3.3: LDR offset timeseries and applied calibration factor for Wardon Hill, generated by reprocessing SHV mode data from October 2014. The full set of LDR offsets for the RHI dataset is included in appendix C.

Figure 3.3 shows an example timeseries of LDR offsets calculated using measurements of the depolarised emissions from sun “hits”, where the sun is in the radar’s direct field of vision. These hits appear in higher scan elevations ($2\text{-}4^\circ$) at predictable times and coordinates around sunrise and sunset. For temporal stability, the applied LDR calibration factor is a running median of the previous 25 hits. This calibration scheme is now running live in the operational Radarnet system.

Emissions measurements are not collected in the short pulse mode used for RHI scans. To calibrate the Wardon Hill dataset, an archive of SHV mode PPIs from September 2014 to August 2015 were reprocessed to calculate the LDR offsets, which usually updated around twice per day (at sunrise and sunset). Each RHI was calibrated using the most recent calculated offset from that validity time.

An example timeseries of LDR offsets for Wardon Hill for October 2014 is shown in figure 3.3. (The complete series, for both RHI study periods, is included in appendix C.) There is clearly a difference between LDR offsets calculated at sunrise and sunset, although which offset measurement is higher does not remain consistent over long periods. This fluctuation occurs to a greater or lesser degree for all radars in the UK network, and may be due to temperature effects, although it has not been possible to demonstrate a connection to any obvious physical factor. However, although the raw measured offsets fluctuate, the applied LDR offset is stable to within approximately ± 0.2 dB for the majority of the study period. This uncertainty is negligible in the context of other expected errors on LDR, which are discussed in chapter 5.

3.4 Extracting meteorological profiles

3.4.1 Removal of non-meteorological echoes

After calibration of both reflectivity and LDR, the RHI dataset was quality controlled to remove non-meteorological echoes. The first step in this exercise was filtering by inspection to remove both “dry” RHIs and those containing undesirable wide-ranging effects, such as radio frequency interference. The remaining scans were then examined using a pixel-by-pixel algorithm to identify meteorological echoes.

Since the RHI scans are in LDR mode, the dual polarisation Bayesian classifier described in section 1.5.1 cannot be used to filter out non-meteorological echoes. Many of the single polarisation filters used in Radarnet are based on gridded spatial information that cannot be mapped onto a “vertical slice” RHI scan, even if the precise azimuth were known. The information available to quality control this data was therefore limited to the datasets available within each scan: reflectivity, LDR and CPA.

Met Office single polarisation radars have traditionally used a measure of signal amplitude variability to identify ground clutter. The “clutter indicator” (CI) is an average of the pulse-to-pulse power difference in dB between successive echoes at the same range gate:

$$\Delta_k = 10\log P_k - 10\log P_{k-\tau} \quad (3.2)$$

$$\text{CI} = \frac{1}{M-1} \sum \Delta_k \quad (3.3)$$

where M is the number of pulses per range gate and τ is the inverse of the pulse repetition frequency (Sugier et al., 2002). This captures variability even in spatially slow-moving precipitation, by responding to the changes that result from hydrometeors falling into and out of the radar pulse volume. The distributions of CI in clutter and rain are relatively well separated, with peaks in the region of 0-2 dB and 6-7 dB respectively. However, there is still some overlap between the distributions.

Clutter phase alignment (CPA) offers a similar method of distinguishing between rain and clutter, using the phase stability of an echo rather than its magnitude. The phase of a radar echo is defined as:

$$\phi = \arg(x_i) \quad (3.4)$$

where x_i is the sum of the real (“in-phase” I) and complex (“quadrature” Q) components of the received signal ($x_i = I_i + jQ_i$). The phase is constant in time for fixed ground clutter targets, but varies for weather echoes. Hubbert et al. (2009) suggest a measure

of phase variability normalised by power:

$$\text{CPA} = \frac{\sum x_i}{\sum |x_i|} \quad (3.5)$$

A comparison of the two papers Sugier et al. (2002) and Hubbert et al. (2009) indicates that the use of phase stability (CPA) rather than amplitude (CI) can improve the separation between rain and clutter histograms. A mapping between these two variables using average detection statistics at UK radars (Norman, 2016, personal communication) shows that a CPA value of 0.525 corresponds approximately to the CI value of 3.5 dB previously used to identify clutter in the UK radar processing system (Sugier et al., 2002).

Using this information, the criteria used to identify meteorological pixels within each RHI scan was as follows:

- Reflectivity greater than 10 dBZ
- LDR less than -5 dB
- CPA less than 0.525

The reflectivity and LDR thresholds are qualitative and conservative - they were chosen by observation to exclude noise and clutter, but to include all meteorological features, such as strong bright bands (with potentially very high LDR) and the low reflectivity values associated with ice at high levels in the profile. The CPA threshold corresponds to that used in the Met Office Radarnet quality control system. All other pixels were identified as noise ($Z < 10$ dBZ) or clutter, and excluded from the final dataset.

3.4.2 Regridding to vertical profiles

Following quality control, the polar RHI data were interpolated onto a 100 m by 100 m resolution Cartesian grid. A very fine Cartesian grid was chosen so that this initial regridding could be done using a simple “nearest neighbour” algorithm. From this intermediate grid, the data were averaged to 1 km resolution in the horizontal. Reflectivity was averaged as a linear quantity; for LDR, the two linear components of the ratio were averaged separately (using the colocated reflectivity), and then recombined to give the average LDR. The averaging done at this stage reduced the random noise in reflectivity, as well as negating any possible differences between scans at different radial bin lengths (section 3.2.1). With the data in this form, vertical profiles of reflectivity and LDR could be extracted directly from the final 1 km (horizontal) by 100 m (vertical) Cartesian grid.

To preserve fine vertical structure and to minimise the effect of non-zero elevation on

LDR, the vertical profiles for this study were taken only from ranges between 5 km and 15 km from the radar location. Given the proposed application of this research to QPEs, profiles having a reflectivity of less than 20 dBZ at the lowest levels (equivalent to a very light rain rate of approximately 0.6 mm h^{-1}) were excluded from the sample. Then from the initial RHI sample, taken at continuous 10 minute intervals over more than 8 months total duration, the various stages of quality control and refinement resulted in a final dataset of 6680 high resolution vertical profiles of colocated reflectivity and LDR, including data from 2283 RHIs taken on 104 different days. This final dataset underpins the investigations of chapters 4-7 of this thesis.

Chapter 4

Using the linear depolarisation ratio to identify non-bright band VPRs

4.1 Introduction

An important step in radar quantitative precipitation estimation (QPE) is the conversion of a beam broadened reflectivity measurement aloft into an estimate of the reflectivity at the ground. This process is known as correcting for the vertical profile of reflectivity, or VPR.

Of the many existing and few operational VPR schemes reviewed in chapter 2, the majority are “global”. This means that the atmospheric “background” VPR (section 2.3.1) is assumed to be the same over large spatial and temporal domains. A significant strength of the Met Office operational VPR scheme (section 2.5.3, based on Kitchen et al. (1994)) is the use of ancillary model and climatological information that allows the profile to be determined at the pixel level, so that local differences in precipitation structure can be resolved. The method includes, for example, an adaptive bright band, which varies in strength according to precipitation intensity and is matched in height to the estimated freezing level at a spatial resolution of 5 km. However, this scheme is not designed for convective precipitation, and such profiles must be identified and excluded prior to correction for stratiform bright band.

Although several methods are available in the literature to distinguish stratiform and convective precipitation (reviewed in section 2.4), there are few designed to resolve convection on the scales required for a local pixelwise correction. The algorithm employed

by the Met Office operational radar processing software (Radarnet) uses a simple reflectivity threshold to diagnose convection. If a reflectivity exceeding 30 dBZ is measured at a height exceeding 1 km above the wet bulb freezing level (hereafter referred to as Z_1), the pixel is classed as convective. This is based on the work of Smyth and Illingworth (1998), who find that reflectivities above 30 dBZ do not occur at this level in stratiform precipitation. However, the operational threshold $Z_1 > 30$ dBZ is known to underdiagnose convection, leading to many non-bright band cases being “corrected” for an assumed bright band profile.

Figure 4.1 shows an example of the impact of inappropriate bright band correction in moderate intensity rainfall. The true VPR in this case is an average of the 5-15 km data from the RHI scan at Wardon Hill at 13:46 on 9 October 2014, where there was no bright band. The Radarnet stratiform VPR shape (Kitchen et al., 1994; Kitchen, 1997) is fitted to simulated measurements from this VPR at different ranges from the radar, and the resulting surface rainfall estimates compared to the true value of 10.1 mm h^{-1} . At short range, where the radar beam is below the melting layer, the rainfall estimate is accurate. However, at longer ranges, the assumption of a bright band profile causes underestimation of up to 66% in regions where the melting layer is sampled (70-150 km range).

The implications of inappropriate correction for bright band are significant in an operational context. While reflectivities exceeding 30 dBZ are not observed at high levels in stratiform precipitation, the apparent corollary - that for all convective profiles $Z_1 > 30$ dBZ - is not necessarily true. Misdiagnosed profiles such as those illustrated in figure 4.1 are likely to include some intense precipitation resulting from localised embedded convection and squall lines, which nevertheless do not quite meet the stringent reflectivity criterion for convective diagnosis. Correcting inappropriately for bright band can therefore cause significant underestimation in the high impact rainfall cases which are of particular importance for hydrological modelling, pluvial flood forecasting and event management.

The recent upgrade of the Met Office radar network to dual polarisation provides opportunities to refine the identification of convective profiles. The linear depolarisation ratio (LDR) has shown promise in identifying widespread bright band in radar plan position indicator (PPI) scans (Smyth and Illingworth, 1998), due to its response to the high scattering cross-section and canting angle of the large melting snowflakes that cause the radar reflectivity bright band (section 1.4.3). Additional results from Illingworth and Thompson (2011) illustrate different melting layer LDR signatures in a small sample of stratiform and convective cases, raising the possibility that melting layer LDR could be systematically different in stratiform and convective precipitation. The potential to use LDR to identify bright band *in situ* could reduce the incidence of inappropriate bright

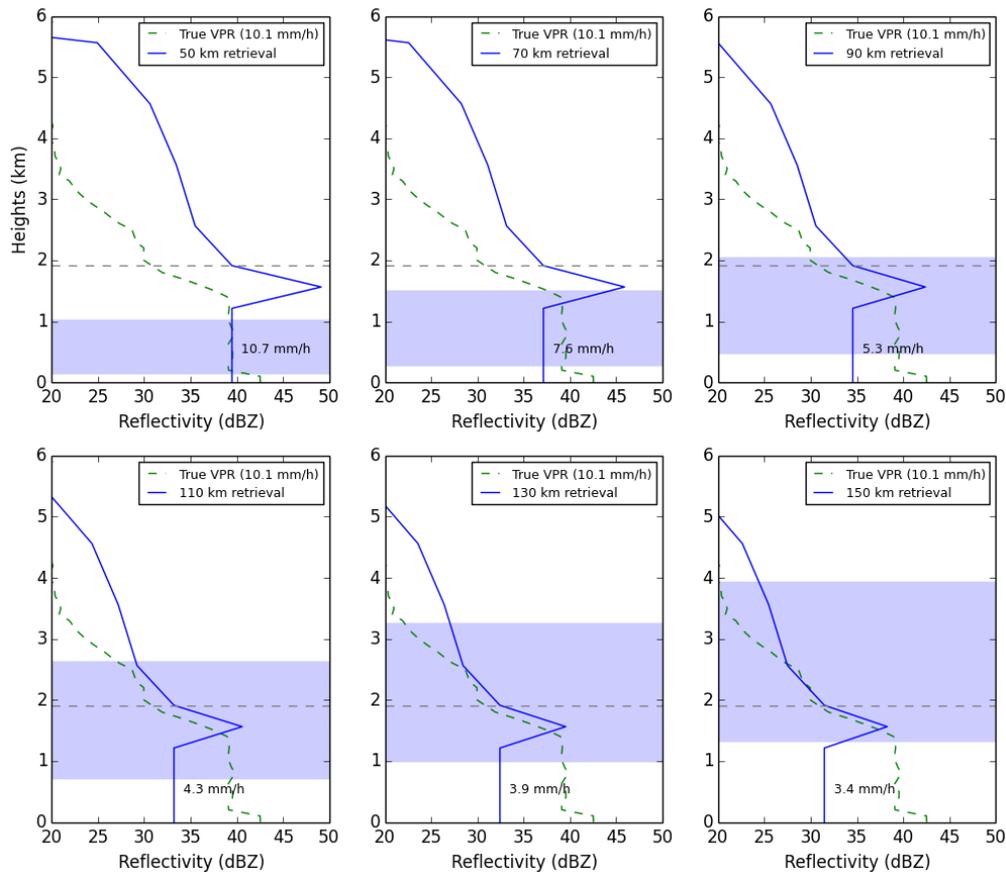


Figure 4.1: Example of the impact of inappropriate bright band correction on surface QPEs in a non-bright band case. The blue line shows the fitted VPR, the green dashed line the true atmospheric VPR, and the purple shading shows the position and extent of the (broadened) radar beam at that range. The dashed horizontal line shows the height of the 0°C isotherm.

band corrections, and conserve more of the high rainfall rates associated with localised convective cores.

Using the high resolution vertical profile dataset described in chapter 3, this chapter investigates the skill of the linear depolarisation ratio in distinguishing between stratiform and non-bright band reflectivity profiles. As an initial investigation to quantify the results of Illingworth and Thompson (2011), this chapter focuses on the skill of the LDR peak measured at short range in identifying non-bright band melting. Any findings will be developed in chapter 5 to account for beam broadening in long range PPI measurements, with the aim of providing improved classification information to inform the Met Office operational correction for VPR.

This chapter is organised as follows. Section 4.2 describes the types of VPR observed in the high resolution dataset, and how this sample was divided into “bright band” and “non-bright band” categories. In section 4.3, relative operating characteristic (ROC)

curves are developed for both peak LDR and Z_1 . These results are discussed in terms of the respective discriminatory skill of the two criteria, as well as suggesting optimal LDR thresholds for VPR classification. Section 4.4 concludes with a summary and future aims. This work (including all figures except 4.1) has recently been published in the Journal of Applied Meteorology and Climatology (Sandford et al., 2017).

4.2 Defining VPR types within the high resolution dataset

To assess the skill of LDR in distinguishing between profile types, a large dataset of high resolution vertical profiles was derived from Wardon Hill range height indicator scans (RHIs). The scan strategy, calibration, quality control and regriding methods applied to construct this dataset are described in chapter 3. The final dataset, comprising 6680 high resolution profiles of colocated reflectivity and LDR, was derived from measurements at ranges of between 5 and 15 km from the radar. In this section, the properties of the high resolution dataset are described in terms of “true” VPR type, which is based on the reflectivity behaviour in the melting region. These classifications provide the baseline for assessing the discriminatory skill of peak LDR.

4.2.1 Locating the melting layer

The vertical profile dataset was sorted into classes based on the shape of the reflectivity peak in the vicinity of the melting layer. This required the development of an automated melting layer detection algorithm based on the data available in the RHI scans. Although several dual polarisation melting layer detection algorithms exist in the radar literature (eg Matrosov et al., 2007; Giangrande et al., 2008), these methods require measurements of the copolar correlation coefficient (ρ_{hv}), which cannot be measured in the “LDR mode” in which the RHI scans were taken (section 1.4). For the purposes of this study only, therefore, a simple LDR-based algorithm was defined to locate the boundaries of melting in RHI profiles.

The melting layer was defined as the region around the LDR peak where the LDR values and gradients met certain conditions. By experimenting with thresholds for a selection of vertical profiles, the following identification procedure was developed:

1. Find the maximum measured LDR in the profile above the expected clutter height (300 m). This peak is assumed to occur within the melting layer.
2. Search downwards from this peak for the melting layer base, where $\nabla\text{LDR} < 20$ dB km⁻¹ and LDR < -25 dB. The reflectivity at this point is Z_{rain} .

3. Search upwards from the LDR peak for the melting layer top, where $\nabla\text{LDR} > -20$ dB km⁻¹ and LDR < -25 dB. The reflectivity at this point is Z_{ice} .

(The gradient symbol ∇ is used to denote the derivative of a variable with respect to height only.) The maximum reflectivity in the melting layer (between the heights of Z_{rain} and Z_{ice}) is designated as Z_{peak} , which is not usually colocated with the peak in LDR. The LDR gradient for this algorithm was calculated over a 200 m height window, between the values immediately above and below the point of interest. The clutter height, LDR and gradient thresholds were determined by inspection.

It should be reiterated that this method was used only to automate the processing of this particular dataset, and there is no intention to extend it to any other context, given that robust dual polarisation melting layer detection algorithms already exist for PPI measurements.

4.2.2 Observed profile types

Initial analysis of a selection of RHIs and profiles suggested that this dataset should not simply be classified as “stratiform” or “convective”, but that three categories would be more appropriate. This contrasts with the majority of previous literature, which bases VPR classification around the two categories of stratiform “bright band” and convective rain (eg Steiner et al., 1995). However, this finding of more than two types of precipitation profile is far from new (eg Fabry and Zawadzki, 1995; Bringi et al., 2009; Delrieu et al., 2009; Matrosov et al., 2016). In this case, the three observed classes in the high resolution dataset align well with three of the five VPR types identified by Fabry and Zawadzki (1995) using a vertically pointing X-band radar:

Low-level rain: shallow, light rainfall developing below the zero degree isotherm in stratiform conditions.

Rain with bright band: cold rain developing above the zero degree isotherm in stratiform conditions. This profile shows a clear increase in reflectivity with onset of melting and decreasing Z below the melting layer, forming the traditional reflectivity bright band.

Rain from compact ice: similar to the “rain with bright band” profile, in that increased reflectivity occurs with the onset of melting, but no decrease in Z is observed below the melting layer. Fabry and Zawadzki (1995) speculate that this profile shape “is likely caused by the melting of fast-falling snow pellets or dense graupels”. This is supported by later DSD analyses of Matrosov et al. (2016).

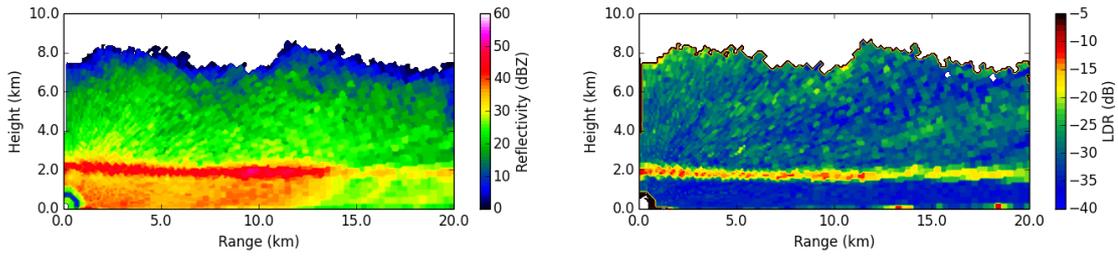


Figure 4.2: Example stratiform RHI: reflectivity and LDR from 14:16 UTC, 13th October 2014, truncated at 10 km height and 20 km range. The bright band is clearly visible as a region of enhanced reflectivity and LDR at 2 km altitude, just below the 0°C isotherm (see also figure 4.5, top panel). Note that the bright band is clearly visible in LDR even at 15-20 km range, contrasting with the weaker reflectivity bright band in this region.

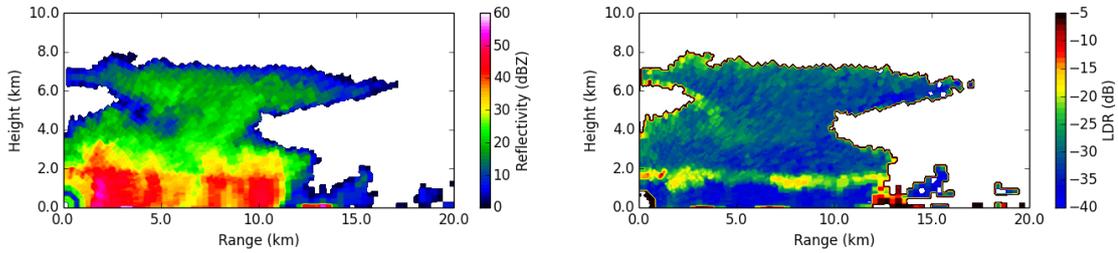


Figure 4.3: Example “rain from compact ice” RHI: reflectivity and LDR from 13:46 UTC, 9th October 2014, truncated at 10 km height and 20 km range. The “compact ice” region around 5-7 km range shows no clear bright band in reflectivity, and correspondingly lower LDR than in the surrounding bright band regions. However there is a sharp increase in reflectivity in the melting layer at 2 km altitude (see also figure 4.5, middle panel), which is not consistent with convective updrafts.

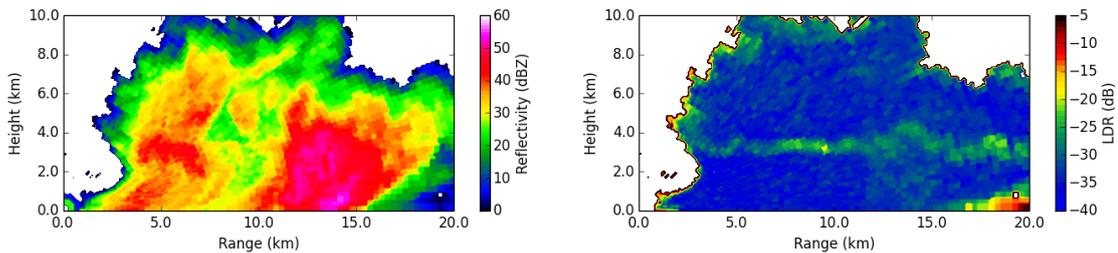


Figure 4.4: Example convective RHI: reflectivity and LDR from 22:20 UTC, 3rd July 2015, truncated at 10 km height and 20 km range. A weak bright band in the region of 8-10 km range, marking the 0°C isotherm at 3.5 km altitude (figure 4.5, bottom panel), contrasts sharply with the convection at 10-15 km. High reflectivity in this region extends consistently around 2 km above the melting layer. LDR values in the convective region are lower than in the weak bright band region, and much lower than in the strong bright band case of figure 4.2.

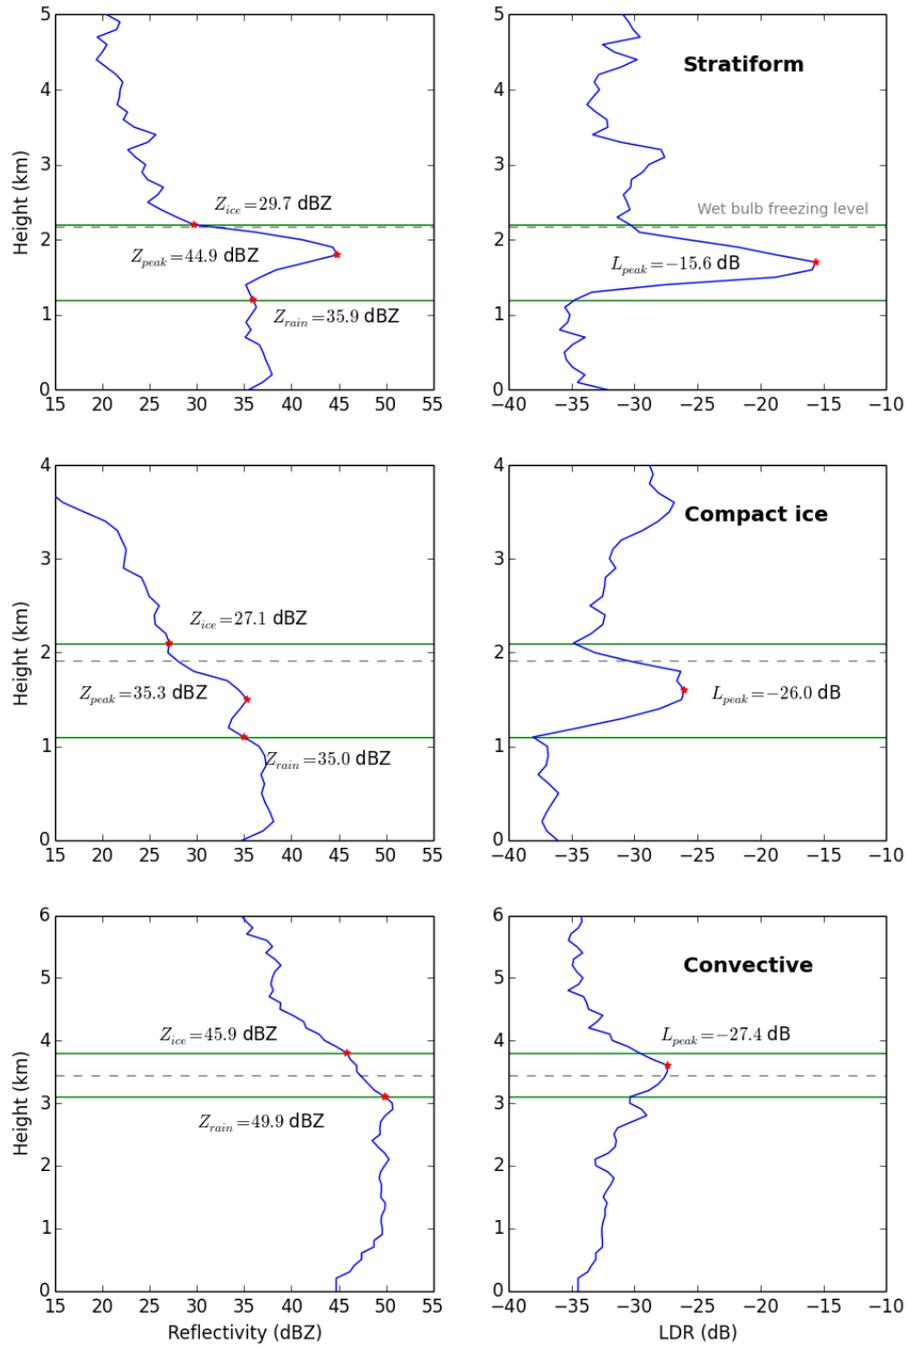


Figure 4.5: Top: vertical profiles of reflectivity and LDR from the stratiform RHI in figure 4.2 (7.5 km range). Middle: vertical profiles of reflectivity and LDR from the compact ice RHI in figure 4.3 (5.5 km range). Bottom: vertical profiles of reflectivity and LDR from the convective RHI in figure 4.4 (12.5 km range). Limits of the LDR-determined melting layer are shown in green, and the wet bulb freezing level by the dashed grey line. Annotated red stars show values at the key levels: reflectivity at the top (Z_{ice}) and bottom (Z_{rain}) of the melting layer, peaks (Z_{peak}) for stratiform and compact ice cases, and the peak melting layer LDR (L_{peak}).

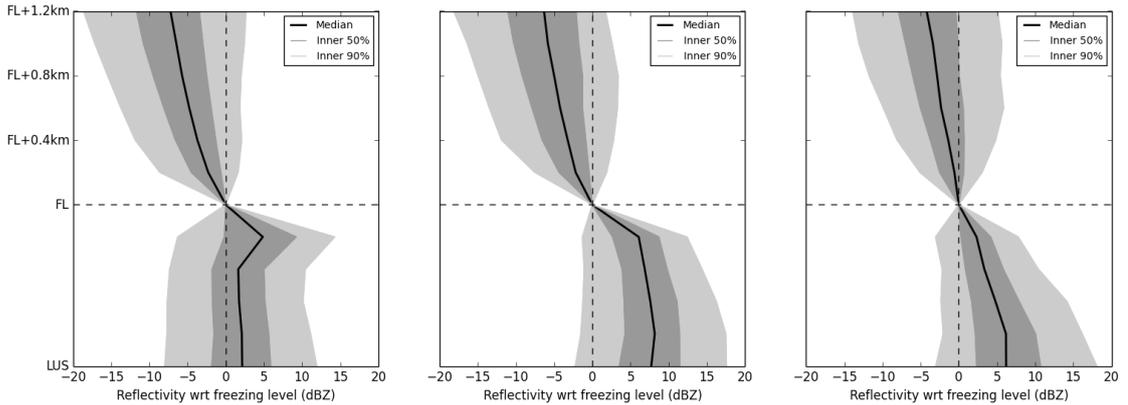


Figure 4.6: Average stratiform, compact ice and convective reflectivity profiles with height relative to the model derived (Brown et al., 2012) wet bulb freezing level. Height levels are at six evenly spaced intervals between the lowest usable reflectivity (LUS) and the freezing level (FL), and then in 200 m steps above the freezing level.

Showers: shallow, light rainfall developing below the zero degree isotherm in convective conditions.

Deep convection: the unstratified profiles observed where updrafts are present in thunderstorms, squall lines and embedded convective cells.

On the basis of Fabry and Zawadzki (1995), the VPR dataset was sorted into “rain with bright band” (hereafter “stratiform”), “compact ice”, and “convective” classes. Examples of each profile type are shown in figures 4.2-4.4. It is entirely possible that both of the other profile classes (“low level rain” and “showers”) were also observed by Warden Hill during the study period. However, both of these profile types were excluded from this investigation by design, since in these cases precipitation formation and growth occurs in the liquid phase and there is no melting layer to classify.

Using the three observed categories of VPR, the “true” classification for each profile in the dataset was determined based on the shape of the melting layer peak. Figure 4.5 shows how the maximum reflectivity in the melting region (Z_{peak}) and at the top (Z_{ice}) and base (Z_{rain}) of the LDR-determined melting layer were compared. Classification rules were applied as follows:

1. If the peak-to-rain reflectivity difference $\Delta Z = Z_{peak} - Z_{rain} \geq 3$ dB, the profile has a bright band and is therefore “stratiform”
2. If there is no bright band but the peak-to-ice reflectivity difference $Z_{peak} - Z_{ice} \geq 6$ dB, the profile is “compact ice”
3. Otherwise, the profile is “convective”

The choice of quantitative ΔZ and peak-to-ice classification thresholds is discussed in section 4.2.4.

4.2.3 Properties of the classified dataset

As expected for a high latitude climate, the most prevalent VPR type in the Wardon Hill dataset was found to be stratiform, accounting for 84% (5600 profiles) of the total 6680 sample. Compact ice profiles accounted for a further 10% (677), and 6% of profiles (403) were classed as convective. Examples of individual RHIs and profiles of each type are shown in figures 4.2-4.5. The average shapes of these profiles, along with quantiles to illustrate spread, are shown in figure 4.6.

Although this chapter is focused on LDR, many pre-existing classification algorithms have used properties of the rain-level reflectivity field to distinguish between stratiform and convective precipitation. In particular, several methods are based on the reflectivity “intensity” and “peakendess” criteria of Steiner et al. (1995) (as discussed in section 2.4.1). Although reflectivity texture (“peakedness”) information is not readily available, comparing the intensity of rain-level reflectivities in this classified VPR dataset provides some insight into whether convective rain intensities are systematically higher than those of stratiform precipitation.

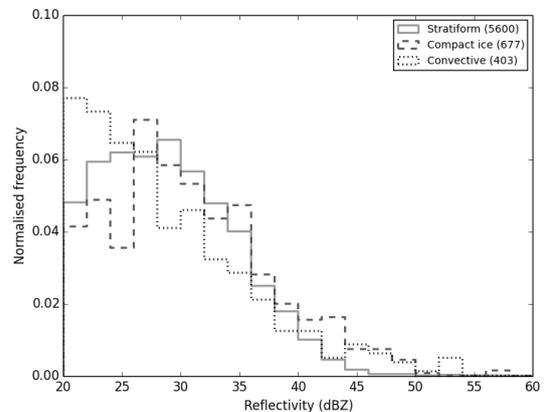


Figure 4.7: Rain reflectivity frequency distributions for stratiform, compact ice and convective VPRs. Numbers in brackets give the total number of each type of profile in the dataset.

Figure 4.7 shows the frequency distributions of rain reflectivity for the stratiform, compact ice and convective profiles used in this study. While small differences exist, the majority of these distributions occupy the same reflectivity region, with no significant difference in modal or mean values between the three profile types. This suggests that approaches based on reflectivity intensity may not be universally reliable in distinguishing between different types of VPR.

4.2.4 Quantitative peak size thresholds for “true” VPR types

Section 4.2.2 listed the minimum peak-to-rain and peak-to-ice reflectivity differences used to classify VPRs in the dataset as stratiform, compact ice or convective. The justification

for those thresholds, and their impact on the investigation of LDR skill, is outlined below.

The reflectivity peak size used to classify stratiform “rain with bright band” is based on inspection of individual profiles, which showed some noise at the dBZ level. In order to account for this when determining the reflectivity peak, a ΔZ value exceeding 3 dB (a factor of 2 in linear reflectivity) was judged suitable for a profile to define a bright band.

The separation of compact ice profiles from convection was more theoretically based. The microphysical explanation proposed by Fabry and Zawadzki (1995) for rain from compact ice is similar to that of “rain with bright band”. In bright band rain, the sudden five-fold increase in dielectric factor $|\kappa|^2$ when snowflakes begin to melt causes a sharp increase in reflectivity below the 0°C isotherm (sections 1.2.1 and 2.2.2). This effect is counteracted by the decrease in diameter and increase in fall velocity to completion of melting, so that the reflectivity decreases, forming the characteristic bright band peak (section 2.2.2). Fabry and Zawadzki (1995) suggest that in rain from compact ice, the drop distribution above the melting layer could be dominated by smaller, high density ice particles, with few or no air inclusions. The increase in reflectivity at the onset of melting, through increasing κ , would then be similar to that in bright band rain. However, since this compact ice is closer in both fall speed and diameter to liquid water, there is no decrease in reflectivity with size or fall speed as melting continues. In this way compact ice profiles would show the pronounced increase in reflectivity observed at the top of the melting layer, but not the characteristic “bright band” peak of stratiform precipitation.

Recent observations of surface drop size distributions by Matrosov et al. (2016) provide additional evidence in support of the “compact ice” model. The authors find that rain from compact ice (which they call “non bright band”), where vertical velocities are not convective but no bright band is observed, has a much higher proportion of small drops than either “stratiform bright band” or “convective” rain. This characteristic rain DSD would be consistent with a population of relatively small, dense ice or snow pellets in the region directly above the melting layer.

Adopting the compact ice model, and using the standard assumption of no aggregation or breakup across the melting layer (eg Szyrmer and Zawadzki, 1999; Hardaker et al., 1995), the change in reflectivity with melting can be calculated directly using known changes in hydrometeor diameter, fall speed and dielectric factor. The small, high density ice particles responsible for compact ice profiles are similar in both diameter and fall speed to liquid water, so the increase in reflectivity is almost entirely due to the increase in dielectric factor. The ratio of the dielectric factors of ice and water (section 1.2.1) suggests a reflectivity increase in compact ice cases of 7.2 dB with melting, or 6.5 dB if the slight difference in density between solid ice and water is accounted for. By contrast, the vertical velocities present in convective updraughts (Steiner et al., 1995) are expected to disrupt the clear stratification of structures around the melting layer, so that there is no sudden

increase in reflectivity at the 0°C isotherm. To avoid false classifications of convection, and to accommodate the likelihood of some lower density snowflakes occurring in compact ice cases, a slightly lower peak-to-ice threshold of 6 dB was chosen to distinguish compact ice profiles from the remaining convective precipitation.

The use of approximate thresholds for “true” classification has the potential to impact the results of this study. If the reflectivity peak size thresholds are not well matched to the underlying physics, any discriminatory skill that LDR is found to have could be spurious, and the confidence in the best LDR peak threshold for profile identification would be low. To increase the robustness of the study results, a range of peak size thresholds will be tested around the selected values. These results are discussed alongside the main outcomes in section 4.3.

4.3 Results and discussion

Having classified the VPR sample according to the reflectivity peak, the maximum value of LDR in the melting layer was extracted from each profile to assess its skill in distinguishing between profile types. Different values of LDR were tested as thresholds for identifying “non-bright band” (compact ice and convective) and convective profiles. These results were compared with the skill of a reflectivity threshold 1 km above the model wet bulb freezing level (Brown et al., 2012), where $Z_1 > 30$ dBZ is the criterion for convective identification currently applied in the UK operational system.

4.3.1 Comparing LDR skill with a high level reflectivity criterion

The intrinsic skill of LDR as a classification criterion was examined by comparing relative operating characteristic (ROC) curves for LDR with curves for high level reflectivity (Z_1). ROC curves above the 1:1 line demonstrate that a quantity has discriminatory skill, and the point furthest above the 1:1 line is the threshold at which that quantity most skilfully distinguishes between profile types.

Contingency tables were generated for a range of melting layer LDR thresholds between -25 dB and -15 dB at 0.5 dB intervals, and for Z_1 from 20 to 36 dBZ at 1 dB intervals. Two sets of tables were generated: one for diagnosis of all non-stratiform profiles (convection and compact ice), and the other for identifying convection only. Correct diagnosis of a “non-stratiform” (or convective) profile is termed a “Hit”; false identification a “False alarm”; incorrect default to “stratiform bright band” a “Miss”; and correct default to a bright band profile is a “No detection”. ROC plots of hit rate (HR) against false alarm rate (FAR) (see appendix B for definitions) were constructed for both classification tests,

with each point on the curve corresponding to a different threshold. These plots are shown in figure 4.8.

The left hand panels of figure 4.8 show the skill of LDR and Z_1 in distinguishing between stratiform and non-stratiform profiles. The best threshold for peak LDR, where a value of less than -20 dB indicates “no bright band”, has a hit rate of 0.42 and a false alarm rate of 0.06. By contrast, high level reflectivity Z_1 has virtually no skill in distinguishing between stratiform and non-stratiform profiles. A threshold of 34 dBZ gives a hit rate of 0.03, which is the highest above the corresponding rate of false alarms (0.01). The operational threshold of 30 dBZ has a higher hit rate (0.06), but this is equal to the rate of false alarms (0.06). From these results alone, the LDR criterion gives a seven-fold increase in non-bright band profile detections over the current operational criterion $Z_1 > 30$ dBZ, with no cost in terms of false alarms.

The right hand panels of figure 4.8 show the respective skill of LDR and Z_1 in separating convection from other types of VPR (stratiform and compact ice). Convective VPRs are skilfully identified by a peak LDR of -20 dB, with a hit rate of 0.58 and false alarm rate of 0.09. By comparison Z_1 shows minimal skill. The best threshold for identifying convection, which is equivalent to the current UK operational convective diagnosis criterion ($Z_1 > 30$ dBZ), has a hit rate of 0.10 and a false alarm rate of 0.05.

It is important to put this detection rate of 0.10 in context, in terms of the quantifiable impacts on operational QPE. Although the RHI profile sample is not directly comparable with data from PPIs, if taken literally this result would suggest that around 90% of convection in the UK is being misclassified, and the corresponding QPEs largely underestimated through inappropriate correction for a non-existent bright band. In fact, the high level reflectivities measured in PPIs will generally be lower than those extracted from short range vertical profiles. The impact of beam broadening at long range generally acts to reduce the measured reflectivity, particularly if the radar beam is only partially filled (that is, the top of the precipitation profile is lower than the top of the broadened beam). In addition to this, the sampling geometry of the Met Office QPE scan strategy (figure 1.2) requires that for reasonable spatial coverage, some flexibility must be allowed in the height used to obtain the convective diagnosis criterion. Practically, this means that the measurement used for convective diagnosis will be from 1 km OR HIGHER above the wet bulb freezing level. So the reflectivity at EXACTLY 1 km above freezing level (from the high resolution VPRs) presents effectively an upper bound on the actual reflectivity measurement that would be used for operational classification. By this reasoning, we would expect the proportion of missed convective diagnoses to be higher than 90% using the current operational scheme.

High reflectivity kilometres above the freezing level is considered a proxy for strong updrafts and large rimed hydrometeor species, such as hail (Smyth and Illingworth, 1998).

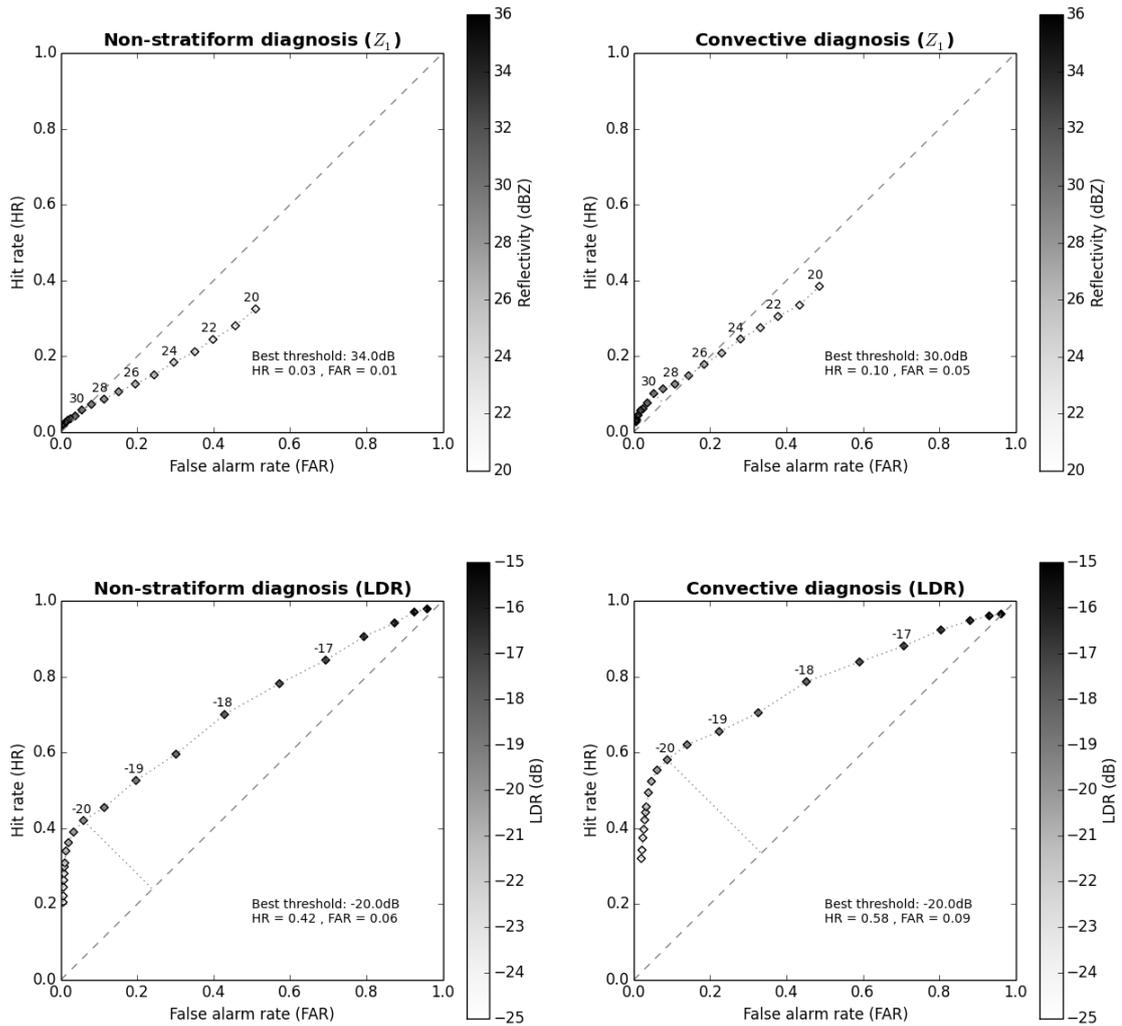


Figure 4.8: ROC curves for Z_1 thresholds from 20 to 36 dBZ (top), and peak LDR from -25 to -15 dB (bottom). Left: skill in identifying non-stratiform profiles (compact ice and convective). Right: skill in identifying convective profiles only. Plot points are colour-coded by the threshold used to classify profiles as non-stratiform (left) or convective (right), and some points have also been labelled by threshold.

The lack of skill of Z_1 in identifying compact ice profiles is consistent with the expected microphysics, since compact ice profiles arise from a different DSD than other types of profile (Fabry and Zawadzki, 1995; Matrosov et al., 2016), and not from convective processes. However, the lack of skill in identifying convective profiles is unexpected, and suggests that reflectivity at this height is, in fact, a poor proxy for updrafts and riming.

4.3.2 Bright band intensity

The distinguishing feature of stratiform profiles is a high peak-to-rain ΔZ value (section 4.2.2), which defines both the presence and intensity of the bright band. This feature

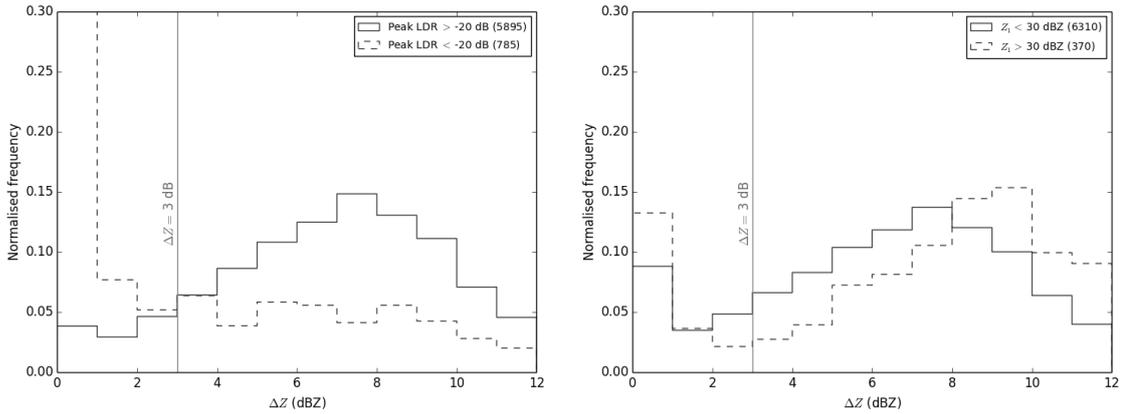


Figure 4.9: Normalised frequency histograms of peak-to-rain (ΔZ) value for VPRs classified as “bright band” or “non-bright band” by LDR (left) and Z_1 (right). Dashed lines correspond to a “non-bright band” classification, solid lines to stratiform VPRs. The figure in brackets in each legend is the total number of profiles included in each histogram.

provides another method by which to illustrate skill in distinguishing between “bright band” and “non-bright band” profile types. If a criterion has skill, the frequency distribution of ΔZ for profiles diagnosed as “non-bright band” should be peaked at very low values. For profiles with bright band, the intensity distribution should peak at higher positive values.

Figure 4.9 shows normalised histograms of ΔZ in the high resolution profile sample for stratiform and non-stratiform profile types, as diagnosed by LDR and Z_1 respectively. The ΔZ threshold for “true” profile classification of 3 dB is shown for reference. The histograms following LDR-based classification are well separated, with the distribution of bright band ΔZ peaked at 7-8 dB, and non-bright band profiles heavily skewed towards $\Delta Z < 1$ dB. However, the Z_1 histograms show a significant amount of overlap. The bright band histogram is similar to that obtained from LDR; however the profiles diagnosed as non-bright band by Z_1 show a bimodal distribution in ΔZ , with equally high normalised frequencies in both $\Delta Z < 1$ dB and the 8-10 dB window. This further demonstrates the lack of intrinsic skill in Z_1 , and the improvements achievable by using LDR.

4.3.3 Sensitivity to VPR type definitions

Section 4.2.4 presents the peak size thresholds used to define reflectivity profiles as stratiform, compact ice or convective. The thresholds of 3 dB and 6 dB for ΔZ and peak-to-ice reflectivity differences respectively were subjectively determined, and therefore have some inherent uncertainty. The “best” LDR threshold of -20 dB, and the qualitative skill of LDR as a parameter (as measured by an ROC curve consistently above the 1:1 line), should not be sensitive to small changes in VPR peak size thresholds within this un-

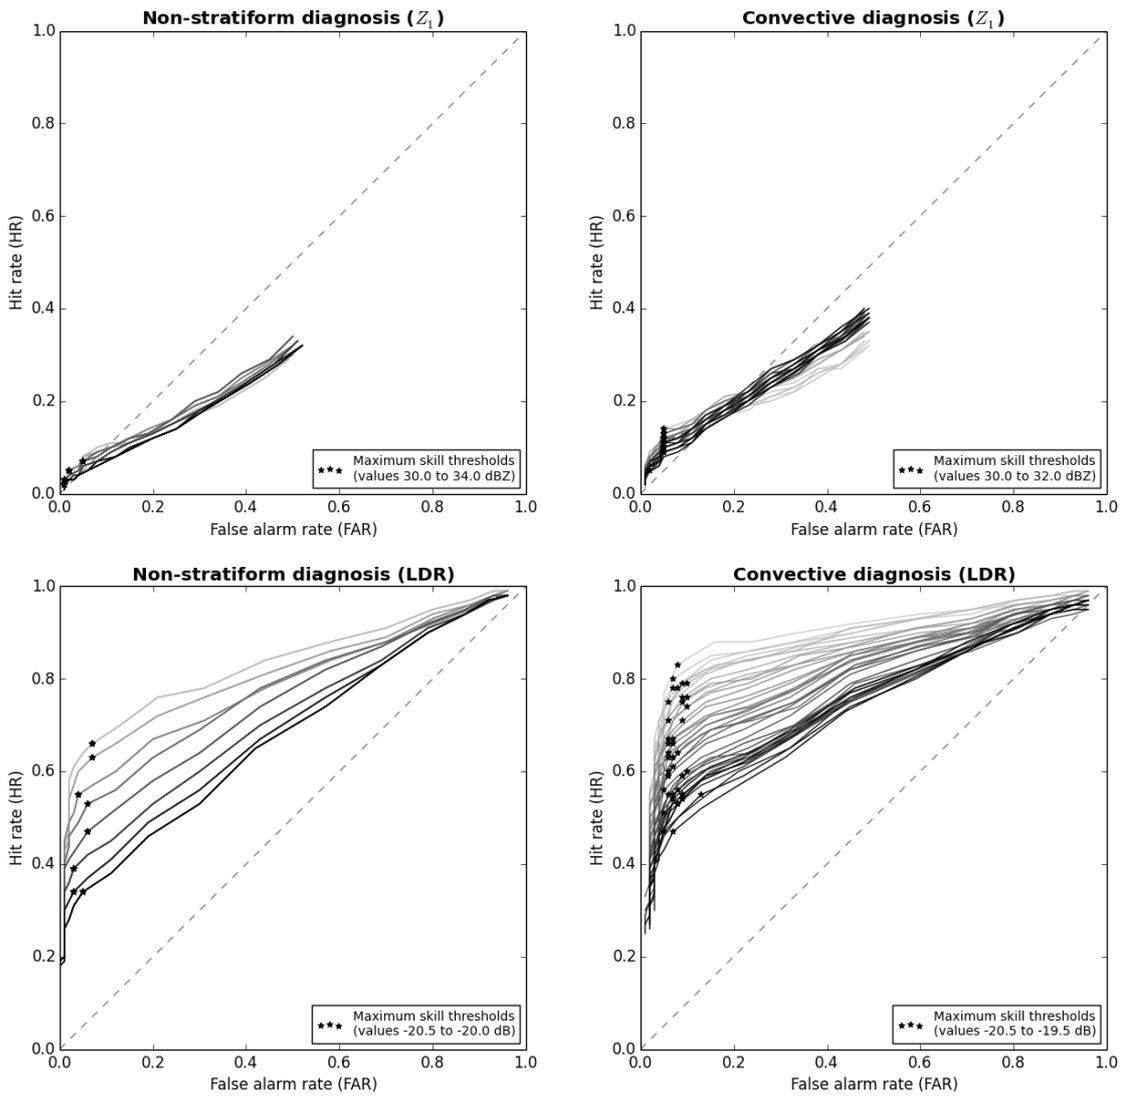


Figure 4.10: As for figure 4.8: ROC curves for Z_1 thresholds from 20 to 36 dBZ (top), and peak LDR from -25 to -15 dB (bottom), for the range of “true” classification ΔZ and peak-to-ice values tested in the sensitivity study. Lighter colours correspond to smaller ΔZ and peak-to-ice thresholds (see also figure 4.11).

certainty range. Consistency in the “best” LDR value for profile discrimination would increase confidence in the finding that LDR has skill, and in the inference that this skill is through response to a physical process: that is, the melting of large snowflakes.

To test the sensitivity of LDR skill to the precise definition of VPR types, additional ROC curves were generated for a range of reflectivity peak size thresholds. The minimum ΔZ to define a profile as “stratiform” was varied between 0.5 dB and 4.0 dB (8 test values), and peak-to-ice thresholds from 4.0 dB to 8.0 dB (5 test values).

Figures 4.10 and 4.11 show the range of ROC curves generated from the 40 different

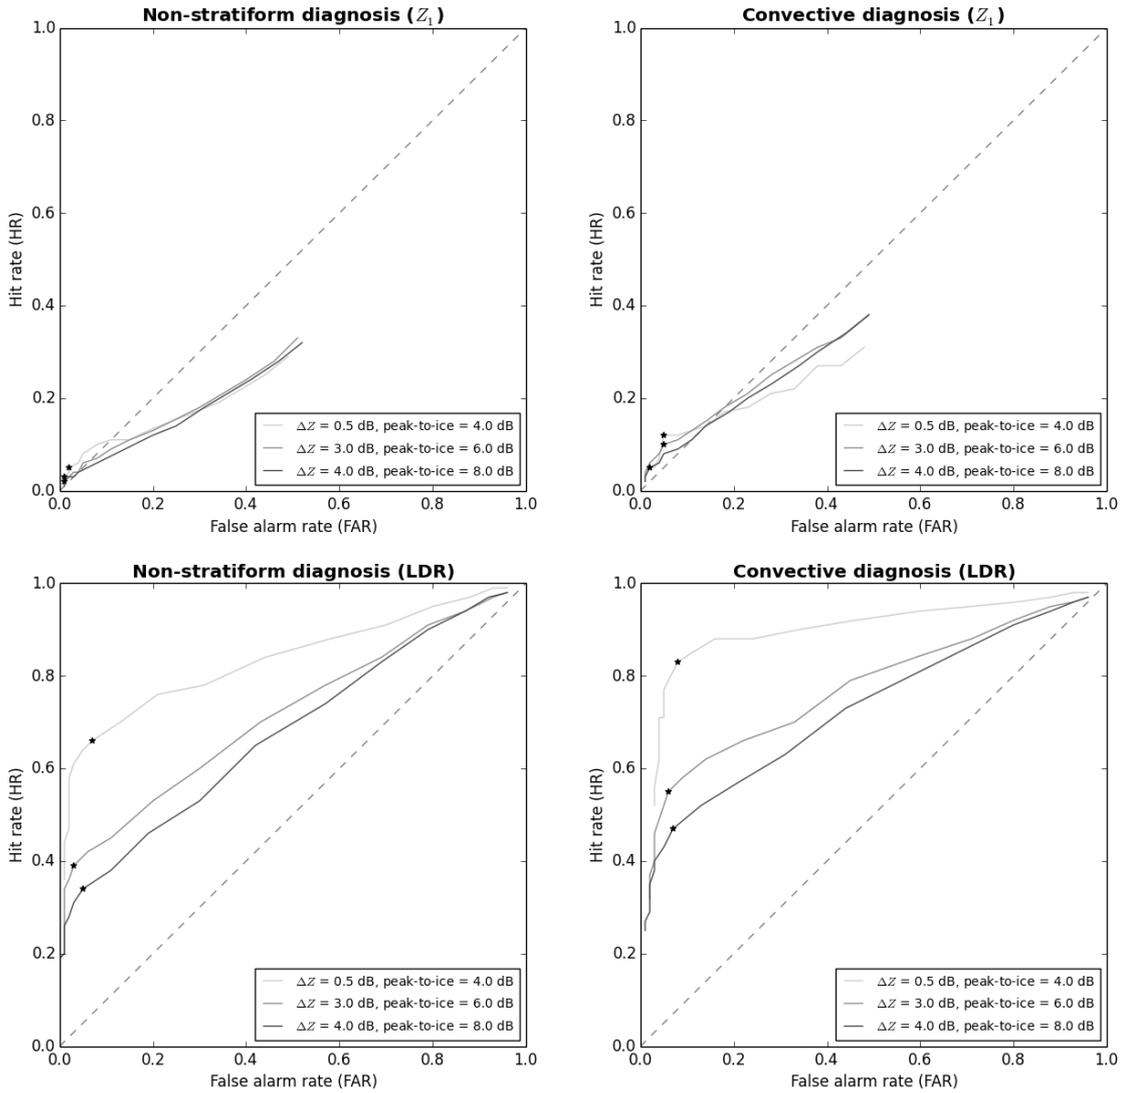


Figure 4.11: As for figures 4.8 and 4.10: ROC curves for Z_1 thresholds from 20 to 36 dBZ (top), and peak LDR from -25 to -15 dB (bottom), for extremes of the sensitivity study shown fully in figure 4.10. Curves are labelled by threshold values to illustrate trends. (Note that greyscale here does not match figure 4.10.)

sets of thresholds used in the sensitivity study. (For the identification of “no bright band” profiles, only 8 of the 40 curves are visible in figure 4.10. This is because the “no bright band” identification is not sensitive to peak-to-ice threshold, so the 5 peak-to-ice curves for each ΔZ threshold are identical.) It can be seen that although the position of the curves differs, all LDR threshold curves are consistently above the 1:1 line. The maximum skill threshold for identifying non-stratiform and convective profiles in LDR is robust to uncertainty in “true” profile definitions, having a value -20 ± 0.5 dB for all combinations of tested thresholds. By contrast, there is no combination of thresholds for which Z_1 shows skill in distinguishing between profile types.

From these data it can be concluded that LDR skill is not sensitive to the precise dividing lines between VPR types as defined by the size of the reflectivity peaks, but is a robust indicator of the presence of large melting snowflakes leading to stratiform bright band.

4.4 Conclusions

This chapter assessed the potential skill of LDR measurements in distinguishing between different types of VPR on a local scale. LDR responds directly to the presence of the large melting snowflakes responsible for the reflectivity bright band. This chapter tested the hypothesis that values of LDR in the melting layer should vary with precipitation type, showing lower values in non-stratiform regions where melting species originate from higher density ice.

In this work a large sample of high resolution vertical reflectivity and LDR profiles were collected and classified into three types, corresponding to three of the types identified by Fabry and Zawadzki (1995). This dataset of 6680 profiles was used to assess the skill of peak melting layer LDR in distinguishing between VPRs with and without bright band. Peak LDR was shown to have diagnostic skill over a range of threshold values, significantly greater than the skill of the high level reflectivity threshold currently used in the UK. A peak LDR value of -20 dB was found to maximise the probability of detection of non-bright band profiles for a given false alarm rate. By contrast, results from the current operational reflectivity-based criterion suggest that 90% or more of convective precipitation may be misclassified as “rain with bright band”. These outcomes illustrate the potential for LDR to effect large improvements in the operational identification of non-bright band reflectivity profiles, thereby reducing the underestimation of rain rates due to inappropriate bright band correction, and with particular benefits expected through preserving the high precipitation intensities associated with convective cores.

Chapter 5

Applying an LDR-based classification criterion to operational data

5.1 Introduction

Chapter 4 demonstrated the ability of the linear depolarisation ratio (LDR) to distinguish between different types of melting in vertical reflectivity profiles (VPRs). This result has implications for surface rain rate estimation from weather radar. The VPR shapes used for correction in both operational and research contexts are often selected based on proxy criteria (section 2.4), which may not be representative of the true underlying microphysical situation. The responsiveness of LDR to the specific hydrometeors responsible for the radar reflectivity bright band makes it inherently a more accurate and reliable diagnostic criterion than many of those currently available for classification. If LDR measurements can be used to classify VPRs in real time, this could greatly improve the representativity of determined VPRs and the accuracy of surface quantitative precipitation estimates (QPEs).

The skill of LDR peaks in VPR classification is not in itself sufficient to deliver benefits to operational QPEs. Real time radar measurements are a convolution of the actual atmospheric profile of the measured parameter (reflectivity or LDR) with the off-axis power profile of the radar beam (section 2.3.1). Since the radar beam broadens with increasing range, any radar measurement will show an increasingly smoothed response to the true atmospheric profile. The peak LDR threshold for classification established in chapter 4 does not, therefore, translate directly into a threshold suitable for LDR measurements at all ranges. The aim of this chapter is therefore to develop the findings of chapter 4 into

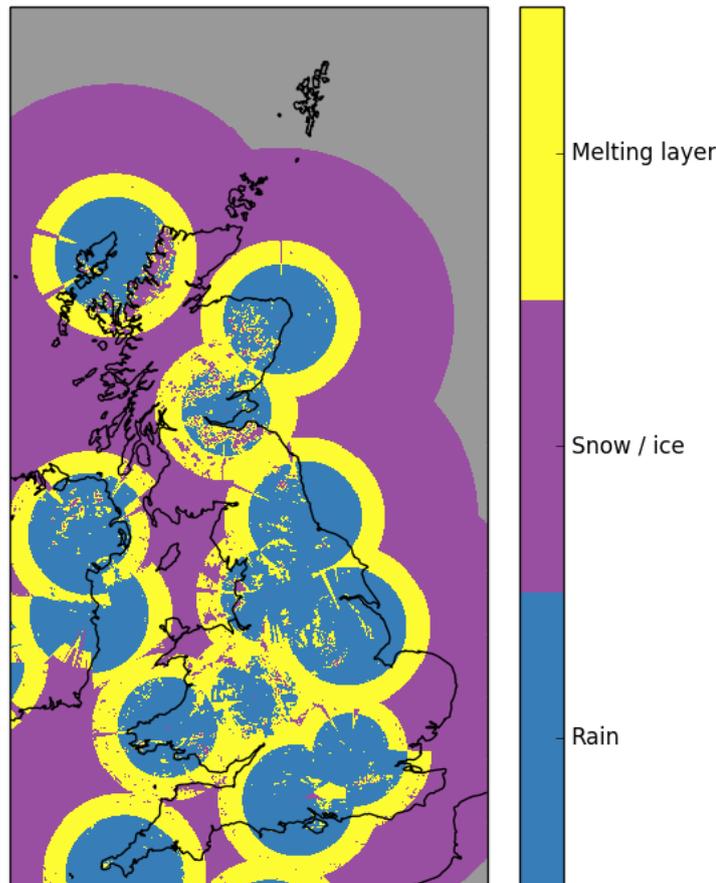


Figure 5.1: Example of precipitation types at the height of the radar beam **axis** over the UK mainland. This example assumes a freezing level (onset of melting) at 2 km altitude and a 700 m melting layer depth (cf Kitchen, 1997). The true influence of the melting layer “bright band” impacts a larger area of the composite due to beam broadening (chapter 1, figure 1.8).

an LDR-based VPR classification algorithm suitable for real time application to radar PPIs.

The primary aim of the suggested algorithm, given the three distinct profile shapes identified in the high resolution dataset, was decided according to operational concerns. Under typical conditions, with a wet bulb freezing level at 2 km (Kitchen et al., 1994), the Met Office radar network covers almost all of the UK mainland at altitudes within or below the bright band (figure 5.1). Regions where the radar beam is habitually centred above freezing level are largely confined to East Anglia and some areas of the Scottish Highlands. It is therefore more important in this context to identify whether or not to correct for bright band than to constrain the shape of the VPR above the melting layer. For these reasons, the implementation developed in this chapter focuses solely on the distinction between VPRs with and without bright band, and not on identifying explicitly the small minority of convective profiles.

The use of only two classes of VPR is a simplification. In reality, the spectrum of VPRs is likely to be more continuous, with no sharp cutoff between “bright band” and “non-bright band” precipitation. This is the principle behind schemes such as that of Tabary (2007), which develop a spectrum of VPR shapes to fit to the data available, in order to avoid any artefacts or discontinuities caused by sharp boundaries in classification. On the other hand, the results of chapter 4 (figure 4.9) clearly show distinct and different distributions of ΔZ as separated by melting layer LDR. The distribution of bright band enhancements around a peak of 6-8 dB contrasts markedly with the non-bright band distribution, which is strongly biased towards a ΔZ of less than 1 dB, and there is minimal overlap between the two distributions. This supports the idea of a clear distinction between these precipitation types, rather than a continuous shared distribution of VPR shapes. For VPRs identified as bright band, the ability of the Kitchen et al. (1994) VPR correction to converge on different values of ΔZ with no minimum limit is expected to minimise any spatial discontinuities in QPEs due to boundaries between different precipitation types.

The first part of this chapter develops the theoretical finding of LDR peak skill into a prototype algorithm for VPR classification in the Met Office operational radar data processing chain (described in section 1.5). Section 5.2 describes the methods used to determine “best” LDR classification thresholds to identify non-bright band (compact ice and convective) VPRs, accounting for the effects of beam broadening over operational ranges of up to 250 km from the radar. The impact of LDR-based classification using these thresholds is evaluated in section 5.3, using a radar simulator and implementation of the current Radarnet VPR scheme (detailed in section 2.5.3). These sections establish the potential benefits of beam-broadened LDR measurements in rainfall estimation, and justify the subsequent work to overcome barriers to an operational implementation.

The second part of this chapter addresses the challenges inherent in using LDR measurements for the first time in an operational environment. Calibration of LDR has been addressed in chapter 3, and can be applied directly to operational PPIs. Section 5.4 describes the data availability and quality control required of real time LDR measurements for this application. This section also considers mitigation for noise issues and the limits on accuracy imposed by hardware quality: specifically the quality of polarisation separation. Finally, section 5.5 evaluates the impacts of a prototype LDR-based classification algorithm on QPEs, by comparing radar precipitation accumulations with those of colocated rain gauges for a high impact non-stratiform case study.

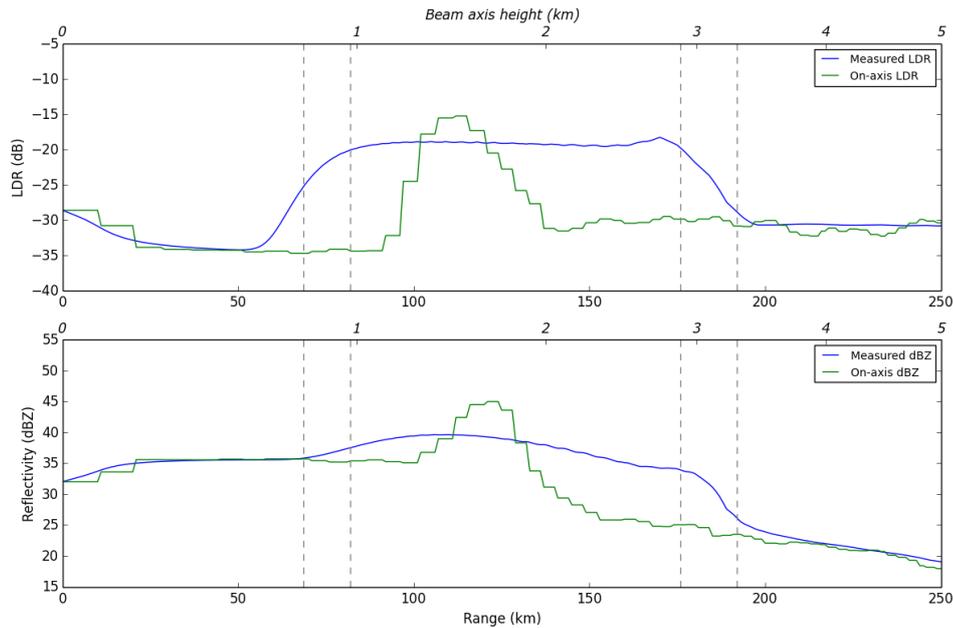


Figure 5.2: An example of simulated measurements and on-axis values of LDR (top) and reflectivity (bottom), for an averaged stratiform VPR from the RHI in figure 4.2, as sampled by a 0.5° elevation radar beam. Vertical dashed lines delineate different melting layer sampling regimes (defined in section 5.2.2). Note that the bright band reflectivity measurement reduces in intensity with range, but the LDR measurement remains constant wherever the peak is sampled.

5.2 Classification of VPRs using simulated LDR measurements

The first step in developing a VPR classification algorithm is to establish whether the diagnostic skill of LDR peaks can be achieved using longer range measurements. Like reflectivity, measurements of LDR are affected by broadening of the radar beam with range. This effect is illustrated in figure 5.2, where it can be seen that the maximum values of reflectivity and LDR measured by a 0.5° elevation PPI in the melting layer are not the same as the “true” peaks (at 100 m vertical resolution) from the high resolution profiles. It follows that the skill and thresholds for VPR classification derived in chapter 4 do not apply directly to LDR measurements at longer range.

The skill of beam broadened LDR measurements for classifying VPRs is assessed in this section using a simulation study. A radar simulator is used to generate measurements of reflectivity and LDR at operational ranges (up to 250 km) from the high resolution vertical profile dataset. Classifications from these simulated measurements are compared with the “true” classifications (described in section 4.2.2), using the same methods applied to evaluate peak LDR skill in chapter 4. The results of classification using different LDR thresholds are plotted as ROC curves, to establish both the level of skill and the

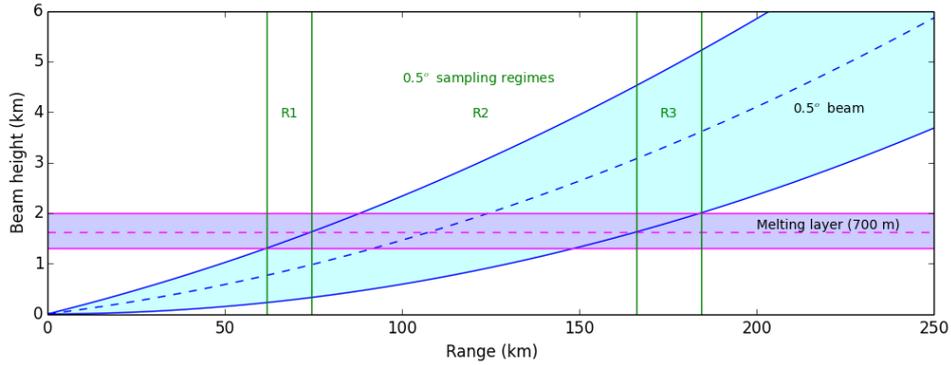


Figure 5.3: Diagram of the position of a 0.5° elevation radar beam with range. Melting layer sampling regimes are shown and labelled in green.

most skilful diagnostic threshold. The results are analysed both as a whole and split by measurement range, to determine whether thresholds should be varied with range to account for different levels of beam broadening.

5.2.1 Radar simulator

For this investigation the radar simulator designed by Kitchen and Jackson (1993) was used to generate data from a beam with a sinc^4 power profile and 1.0° beam width (equation 2.5), to simulate measurements equivalent to those from a UK operational radar. This simulator was extended to generate LDR measurements by first breaking down the LDR profile into vertical profiles of Z_{hh} and Z_{vh} using the known vertical reflectivity profile (equation 1.22). Broadening of the two reflectivity components was then simulated separately, and the broadened values recombined to produce the “measured” LDR.

Measurements of reflectivity and LDR at different ranges were simulated from each of the observed profiles in the high resolution dataset (chapter 3). Each of the 6680 VPRs in the sample is unique, so the freezing level height is different in each case, providing a well spread sample of melting layer measurements at different ranges in the simulated scans. Simulated LDR measurements in the melting layer region were used to classify each VPR using a variety of possible thresholds, to generate statistics for ROC curves and “best” threshold determination.

5.2.2 Sampling regimes: defining the melting layer region

A broadened radar beam at low elevation samples different layers of the atmosphere in varying proportions over a substantial part of the scan domain (figure 5.3). A limitation of LDR as a diagnostic criterion is that it can only be used to classify VPRs where the

beam intersects the melting layer. In evaluating skill in the context of simulated PPIs, it was necessary to consider precisely how the melting layer must be sampled in order to use the LDR measurement in VPR.

The melting layer region is defined as the region over which any part of the radar beam intersects the melting layer. For this purpose the melting layer is assumed to be 700 m deep (Kitchen, 1997), and the melting layer peak (in both reflectivity and LDR) is assumed to be central, at 350 m below the UKV model wet bulb freezing level (Brown et al., 2012). A melting layer measurement can then be assigned to one of three sampling regimes based on the height and breadth of the beam at that location. These regimes were defined as follows:

Regime 1: where the radar beam top is between the melting layer base and central peak (only the lower half of the melting layer is sampled).

Regime 2: where the radar beam spans the melting layer peak.

Regime 3: where the radar beam base is between the central peak and the 0°C isotherm (only the upper half of the melting layer is sampled).

These regimes are illustrated in figure 5.3. By implication there also exists a regime 4, above the 0°C isotherm, where it is desirable to classify VPRs but no melting layer LDR is available; and a regime 0, below the melting layer, where classification has no impact on QPE.

Figure 5.2 shows an example radar response to a bright band profile in the three melting layer sampling regimes. Measured and on-axis values of reflectivity and LDR are plotted with range for an 0.5° elevation beam as it samples an averaged stratiform VPR (from ranges 5-15 km) from the RHI in figure 4.2. The vertical dashed lines in the figure delineate melting layer sampling regimes. Figure 5.2 shows that measured reflectivity values are enhanced throughout the melting layer region (regimes 1-3), but are subject to a gradual decrease with increasing range. By contrast, LDR measurements remain constant with range throughout regime 2, where the beam samples the melting layer peak. This is because unlike reflectivity, which averages linearly, the LDR measurement is strongly dominated by the maximum depolarised reflectivity in the sampling volume. Regime 1 LDR values are considerably higher than LDR in the rain; but regime 3 LDR values are lower and less distinguishable from the ice above. The lower values in regime 3 are due to significant beam broadening, so that high melting layer values of LDR are only sampled by a very small proportion of the radar beam, in the very low power region close to the lower edge of the beam (equation 2.5).

It was decided to use LDR for VPR classification in regimes 1 and 2, but not in regime 3. This is because of the impacts for QPE of false non-bright band diagnoses. In regime

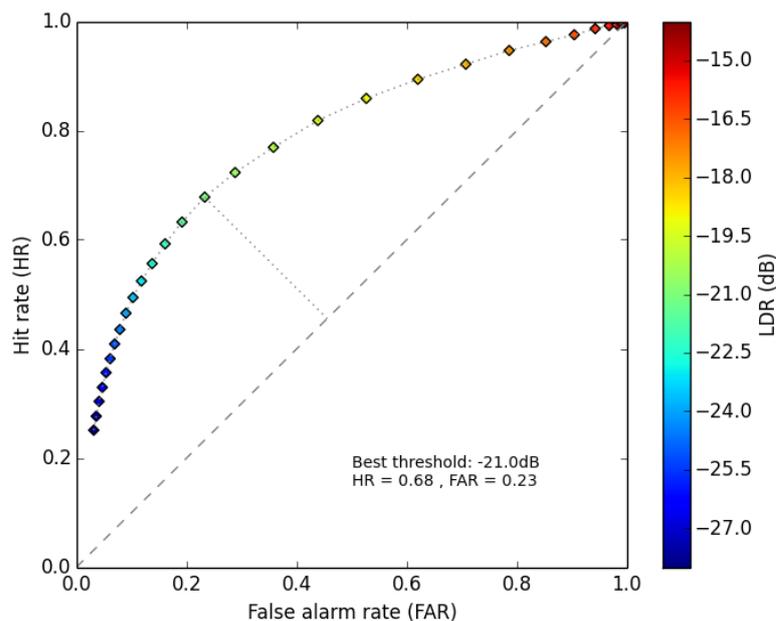


Figure 5.4: Skill of simulated melting layer LDR measurements in identifying non-bright band (compact ice and convective) VPRs, amalgamated across all ranges.

3, where the radar beam is largely above the melting layer, rain rates from profiles incorrectly diagnosed as non-bright band would be significantly underestimated at ground level, as they are assumed to have a constant reflectivity with height (section 2.5.3). By contrast false alarms in regime 1, where the radar beam is mostly sampling rain (figure 5.3), should have negligible impact on surface QPEs. The new algorithm will use LDR to classify VPRs in regimes 1 and 2, and fall back on the original reflectivity-based criterion ($Z_1 > 30$ dBZ, section 4.3) to minimise false diagnoses in regimes 3 and 4.

5.2.3 Overall skill of LDR measurements

Having defined the melting layer region, simulated LDR measurements at 0.5° elevation from regimes 1 and 2 were used to classify VPRs from the high resolution sample as “bright band” or “non-bright band” using a range of possible thresholds. The results for data at all simulated ranges (0-250 km), corresponding to a beam width of up to 5km, are shown in figure 5.4.

Broadened LDR measurements from this simulation show clear skill in distinguishing between VPRs with and without bright bands. When compared to the results from chapter 4 (figure 4.8), the skill in terms of hit rate (HR) minus false alarm rate (FAR) (appendix B) is of a similar magnitude to that of high resolution LDR peaks. For broadened measurements, the slightly lower overall “best” (maximum difference between Hit Rate and False Alarm Rate) threshold value of -21 dB reflects the impact of beam

broadening on LDR over the extent of the radar range.

5.2.4 Range dependent LDR thresholds

To assess in more detail the impact of range, data from the original simulation were divided into 20 km range bins, corresponding to an increase in total beam width of 400 m per bin. The sorted data were used to calculate HR, FAR and the equitable Heidke Skill Score (HSS, appendix B) for non-bright band identification at each range using several different LDR thresholds.

The use of HSS in developing an operational algorithm differs from the ROC curve assessment presented in chapter 4. A simple comparison of hits with false alarms neglects to account for the significant difference in frequency of bright band and non-bright band events, where stratiform bright band occurs over 5 times as often as non-bright band conditions in the Wardon Hill dataset (section 4.2.2). This means that an algorithm that always classified VPRs as “bright band” would have non-zero skill, as it would achieve the correct outcome 84% of the time. The Heidke Skill Score accounts for the lower prior probability of non-bright band precipitation by giving more weight to the correct identification of low probability events. By this measure the “always bright band” classification algorithm would be correctly assessed as having zero skill. The HSS is therefore more likely to be optimised at an objectively “correct” value of LDR (that is, a value with some microphysical significance), and should deliver an algorithm with consistent skill which is not dependent on the relative frequency of events in the underlying climatology.

The best LDR classification threshold as a function of range was chosen to maximise overall HSS, subject to constraints on range dependence. The effect of beam broadening on the measured LDR peak has been shown to be minimal (figure 5.2). This is due to the negligible impact of very low depolarisation from rain and ice phase hydrometeors (section 1.4.3) on an average LDR which includes the much higher melting layer peak. On this basis, the LDR threshold for identifying non-bright band VPRs should be constant with range. However, exceptions to this range independence are likely to occur at short range, where the radar beam width is of order hundreds of metres, so that the beam does not sample the whole of the melting layer. In these cases the measured LDR (and therefore the threshold for non-bright band diagnosis) is expected to be higher, particularly in regime 2.

A constant LDR threshold for non-bright band VPR identification was sought for measurements at ranges beyond 41 km, at which a 1° radar beam width exceeds the 700 m melting layer depth assumed in the current UK VPR (Kitchen, 1997). At shorter ranges the threshold should vary with radar beam width, and have a value of -21 dB at zero

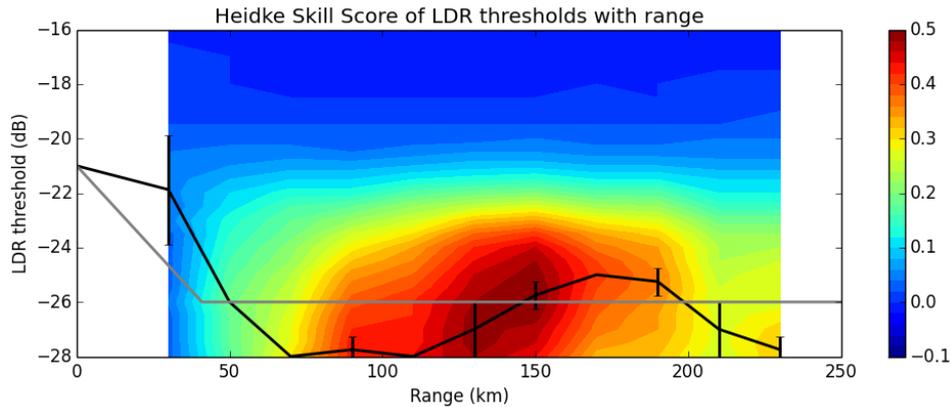


Figure 5.5: Heidke Skill Score of LDR thresholds with range, generated via simulation study using 0.5° and 2.0° elevation beams (20-240 km). The black line joins the LDR thresholds with maximum HSS, the zero-range point is the maximum HSS threshold using high resolution LDR peaks (chapter 4), and the shading shows the actual Heidke skill score value. The grey line shows the “fitted” best skill LDR thresholds to be used for operational convective identification.

range from the radar. This zero-range threshold is slightly different from the -20 dB determined in chapter 4, which maximised the difference between HR and FAR rather than the more equitable HSS, which accounts for the relative frequencies of occurrence of bright band and non-bright band VPRs.

Due to the properties of the high resolution profile dataset, very few short range melting layer measurements could be generated using only a simulated 0.5° elevation beam. Unfortunately, data from within 40 km of the radar is particularly needed to inform how the “zero-range threshold” should transform with range into the lower constant threshold for a broadened beam. A 2.0° elevation beam was therefore added to the simulation to provide extra measurements in this region. Since there is no overlap of a 1° radar beam width between these two elevations, simulated melting layer measurements from the same vertical profile occur at different ranges in each scan. The additional 2.0° data at each range bin are therefore independent of those obtained from the 0.5° simulation.

Figure 5.5 shows the combined results from the 0.5° and 2.0° datasets in terms of HSS. The black line shows LDR classification thresholds with the highest HSS at each range, with error bars set at a minimum of ± 0.5 dB, since this was the difference between tested LDR thresholds. At some ranges a number of LDR thresholds showed the same level of skill. In these cases the median of equally skilful thresholds has been plotted, and the error bar range increased to 0.5 dB above and below the maximum and minimum thresholds.

The non-constant behaviour of the measured HSS with range is unexpected. Given the contrast with the behaviour of LDR in isolated cases, this range behaviour is almost certainly an effect of differences in the sampling of different melting layer regimes. Regime

1 melting layer detections, where the melting layer peak is not sampled, will have lower LDR threshold values than regime 2 detections, because the bright band LDR measurements will be higher when the peak is sampled. Since regime 2 detections occur with higher frequency at longer ranges (see eg figures 5.2 and 5.3), a slight increase in the “non-bright band” detection threshold would be expected at ranges where regime 2 detections are increasing with respect to regime 1. It is therefore likely that a higher proportion of regime 1 detections at intermediate ranges (50-120 km) explains the lower LDR thresholds required to diagnose non-bright band precipitation in this region.

An LDR threshold profile with range was fitted by eye to the HSS dataset and is plotted in figure 5.5. It was decided to vary the LDR threshold linearly in the 0-41 km range bracket, with the linear increase in radar beam width from 0 to 700 m (where the measurement samples the entire depth of the melting layer (Kitchen, 1997)). This resulted in thresholds:

$$\text{LDR}_{\text{thresh}} = \begin{cases} -21 - 5r/41 & r \leq 41 \text{ km} \\ -26 & r > 41 \text{ km} \end{cases} \quad (5.1)$$

where range r is expressed in km and LDR in dB. A melting layer LDR measurement below this threshold is proposed as the new criterion for pixel-by-pixel convective identification in radar PPIs.

5.2.5 Results by sampling regime

Section 5.2.2 defined constraints on the radar beam position for a melting layer LDR measurement to be considered useful for classifying the VPR, and justified the decision to use LDR only from regimes 1 and 2 in classifying VPRs. The following subsections (5.2.3-5.2.4) then used only regime 1 and 2 melting layer measurements to identify LDR skill and thresholds. For completeness, this section compares skill scores for the chosen LDR thresholds in **all** of the three different melting regimes, to illustrate the undesirable behaviour in regime 3.

The classification skill of LDR measurements in each of the three melting layer regimes, in comparison with the skill of Z_1 , is shown in figure 5.6. The LDR-based criterion delivers improvement over the current operational method in all melting layer sampling regimes. As might be expected, the best performance is achieved in regime 2, where the melting layer peak is sampled by the radar beam. Both regimes 1 and 3 have significantly higher false alarm rates, with the equitable Heidke Skill Score for regime 3 being not appreciably greater than that of the current operational criterion. These data, alongside the impacts of false alarms for QPE discussed in section 5.2.2, illustrate the potentially detrimental effects of using measurements from regime 3 in QPE, and support the decision to exclude regime 3 LDR measurements from use in classifying VPRs.

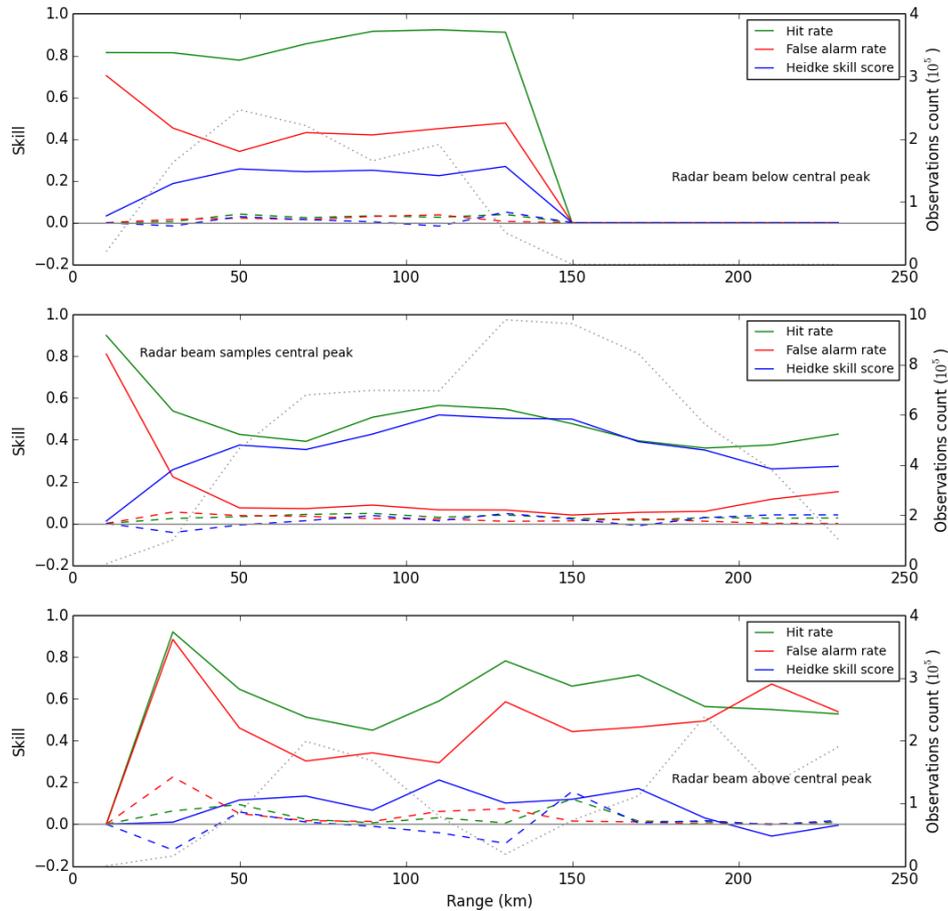


Figure 5.6: Skill by range in the three sampling regimes described in section 5.2.2. Top: regime 1; middle: regime 2; bottom: regime 3 (which was not included in the earlier simulations of this chapter). Solid lines indicate LDR skill, while dashed lines show the skill of the Radarnet operational convective diagnosis criterion ($Z_1 > 30$ dBZ, chapter 4). Observation counts are plotted in grey.

5.3 Benefits for QPE

The benefits of LDR-based classification on rain rates at the surface will be realised indirectly, through application of an appropriate VPR shape to estimate surface reflectivities. Where a non-bright band profile is identified in the UK operational system, no adjustment for VPR is applied. A diagnostic criterion based on melting layer LDR, with a higher detection rate for non-bright band profiles, is therefore expected to improve overall surface QPEs by reducing inappropriate correction for bright band. The quantitative impact of the LDR algorithm on QPEs should be to reduce the underestimation seen at moderate to long ranges in non-bright band rain.

In this section the simulation study setup is extended to evaluate the quantitative benefits of using LDR to identify non-bright band VPRs. The iterative VPR method of Kitchen

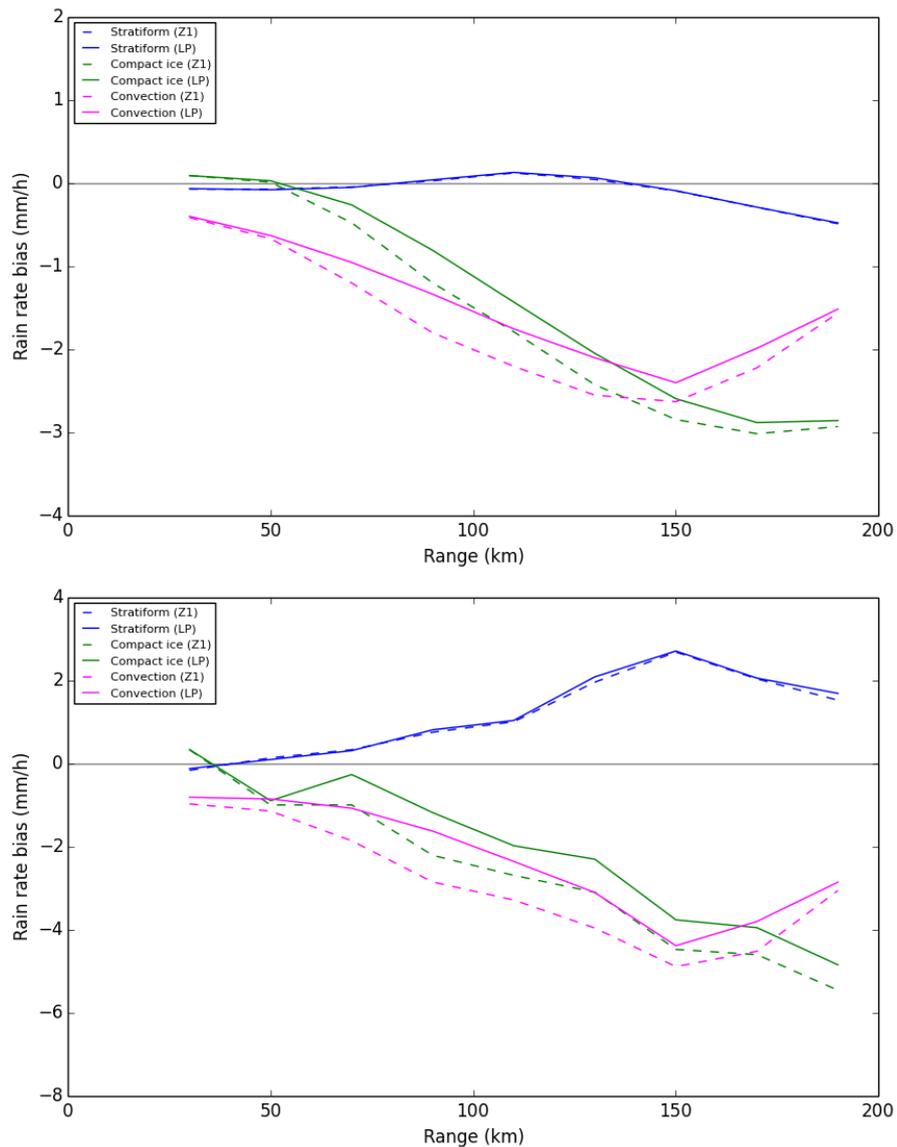


Figure 5.7: Surface rain rate biases with range using reflectivity (Z1: dashed lines) and LDR-based (LP: solid lines) VPR classification schemes. Top: all data; bottom: datapoints where VPR classification differed between the two methods. Note the difference in y-axis scales.

et al. (1994) is used to correct simulated 0.5° reflectivity measurements from the high resolution profiles, where the VPR is classified using either the high level reflectivity ($Z_1 > 30$ dBZ) or the melting layer LDR. If a profile is identified as non-bright band, no VPR correction is applied. Surface reflectivities from the original profiles are then compared with the VPR corrected values. The “surface” reflectivity is defined as the lowest usable measurement in the profile, after the application of quality control as described in section 3.4.1, which must be in the rain below the melting layer. For ease of interpretation this comparison is made in rain rate units, using the Marshall-Palmer relation $Z = 200R^{1.6}$ (section 1.2.2).

The top panel of figure 5.7 shows the bias of simulated surface QPEs with range from the radar using the reflectivity and LDR-based VPR classification schemes. Bright band profiles are dealt with well by both schemes, with overall biases of less than 0.5 mm h^{-1} throughout. For non-bright band profiles, however, rain rates are increasingly underestimated with range. The LDR-based classification scheme clearly reduces these biases, with the greatest improvement in the 50-150 km range, where the melting layer is most frequently sampled by a 0.5° beam.

The bottom panel of figure 5.7 shows biases for the subset of points where LDR and Z_1 classify the same VPR differently. This highlights the type of situation in which using LDR has the greatest impact on QPEs. The differently classified cases show larger overall biases than the dataset as a whole, demonstrating that the impacts of LDR are effectively targeted where the reflectivity-based classification algorithm is not performing well. The reduction in bias using LDR is of order 1 mm h^{-1} over large areas for both compact ice and convective profile types. These data demonstrate the significant potential benefits of using LDR measurements to support radar quantitative precipitation estimation.

5.4 Quality and requirements for operational LDR scans

The first part of this chapter demonstrated the abilities of LDR measurements in principle to identify non-bright band VPRs, using simulations based on a large dataset of high resolution vertical profiles. The physical responsiveness of LDR to the large melting snowflakes that cause the radar reflectivity bright band (chapter 4), combined with the relative insensitivity of LDR peaks to beam broadening (section 5.2), makes melting LDR robustly able to identify non-bright band VPRs even at long range, with corresponding measureable improvements in the resulting QPEs (section 5.3).

Having verified that beam broadened LDR has skill in classifying VPRs, several details remain to be addressed in developing an algorithm for PPIs. It is worth reiterating here that LDR is not typically used by operational radar processing systems. This means that top level quality control issues, such as gross error checks and calibration, have not been addressed thoroughly in the existing literature. A method for calibration of LDR was presented in chapter 3, based on the dual polarisation receiver calibration method of Frech et al. (2017). This method has been successfully applied to real time LDR PPIs.

Aside from quality control and calibration, there are further issues to address. Firstly, there is the issue of representativity. Given that LDR and SHV mode scans are not simultaneous, it is necessary to consider the level of mismatch in time (and equivalently - since precipitation is generally not stationary - in space) that can be tolerated for this application. Secondly, there is the difference between the idealised performance of

simulated data and the realities of a remote sensing measurement. For LDR in particular, because the depolarised component Z_{vh} is so small, the measurement is limited by the receiver noise level. Noise subtraction for Met Office radars is applied at site, and a noise level increasing quadratically with range inevitably leads to some missing data values in LDR, particularly at long ranges. Thirdly, it is necessary to consider the impact of uncertainties in the measured LDR value due to the quality of polarisation separation, and any signal “leakage” between the two polarisation channels. This section addresses the requirements on LDR PPIs and the processing required to be able to use these data for VPR classification operationally.

5.4.1 Hardware quality requirements

A fundamental requirement for the use of any radar parameter is that the measured value provides a reasonably accurate representation of the intrinsic conditions at the target location. For LDR, this is complicated by the minimum detection limit. The minimum LDR that can be measured by a given radar is dictated by the hardware quality: specifically the separation of signal power transmitted in the horizontal and vertical channels, or cross-polar isolation.

For a given cross-polar isolation $10\log(|\epsilon|^{-2})$, the LDR measured by the radar relates to the intrinsic LDR of the scattering target through:

$$\text{LDR}_i = 10\log\left(\frac{\langle S_{vh} \rangle^2}{\langle S_{hh} \rangle^2}\right) \quad (5.2)$$

$$\text{LDR}_m = 10\log\left(\frac{\langle S_{vh} \rangle^2}{\langle S_{hh} \rangle^2} + 4|\epsilon|^2\right) \quad (5.3)$$

where S_{hh} and S_{vh} are the amplitude components of the target scattering matrix (Bringi and Chandrasekar, 2001). These relationships can be used to define a theoretical error on the LDR measurement for any given intrinsic LDR as a result of imperfect isolation.

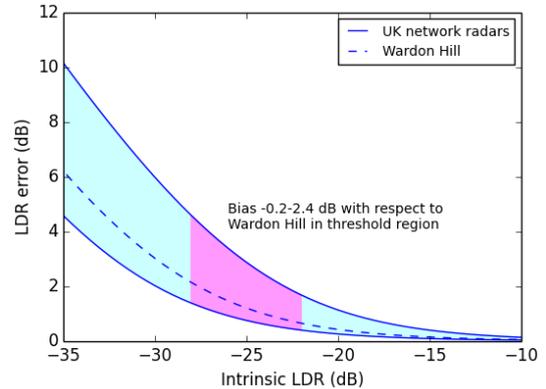


Figure 5.8: LDR measurement error due to limits on cross-polar isolation for a given intrinsic LDR. The blue shaded region represents values for the upgraded UK radars.

Figure 5.8 shows the bias error on LDR due to cross-polar contamination for different intrinsic LDR values (neglecting the impact of noise). The blue shading and lines represent upgraded Met Office dual polarisation radars, with cross-polar isolation values ranging from 31.3 to 38.3 dB (appendix A); and the dashed blue line shows the bias on Wardon Hill measurements, with an isolation of 36.0 dB. Since the LDR thresholds

derived in this thesis were determined empirically using Wardon Hill data, it is the LDR bias with respect to the dashed line rather than total error (with respect to the intrinsic LDR) which is significant in this context. The pink shaded region highlights the values of interest for VPR classification, where Wardon Hill measurements lie between -21 and -26 dB. In this range, the cross-polar isolation of upgraded UK radars results in theoretical biases of -0.2 to 2.4 dB with respect to the same measurement from Wardon Hill. Given the uncertainty on determination of optimal LDR classification thresholds (figure 5.5), and in the context of other uncertainties on the LDR measurement, these errors are considered acceptable.

5.4.2 Scan strategy

Given the multiple parameters available to support quality control and non-reflectivity based rainfall estimation, SHV mode scans (section 1.4) are preferred over LDR mode by the majority of centres for operational QPE (eg Figueras I Ventura and Tabary, 2013; Helmert et al., 2014; Qi and Zhang, 2017). The competition for time in a continuous operational scan strategy means that even the networks capable of measuring LDR typically do not do so operationally.

In the UK, the scan strategy in 2017 (when this work was done) included an LDR mode PPI at 0.5° elevation every 10 minutes. This is not easily integrated with the Radarnet QPE processing cycle, which is based around five minute frequency SHV mode radar volumes. Implementing a suitable scan strategy will be a crucial prerequisite to using LDR in QPE.

Representativity

Given that LDR requires a different scan from the SHV measurements used for QPE (section 1.5), it is not possible to obtain an LDR value that is perfectly representative of the point for which it is required (the reflectivity measurement). Minimising the difference, both spatially and temporally, between LDR and the corresponding SHV scan will be essential to maximising the benefits available from LDR-based VPR classification.

To obtain an LDR measurement that corresponds as closely as possible to the reflectivities used for QPE, the timing and elevation of operational LDR mode scans is important. Of particular concern in the Met Office scan strategy was the 10 minute frequency with which LDR data were available, and the resulting time offset between LDR and the corresponding SHV mode scans.

Timing

The main operational output from the Met Office Radarnet system is the 1 km gridded UK rain rate composite. To generate this composite, complete QPE coverage from SHV mode scans is required from each radar in the network every 5 minutes. LDR scans are currently available only every 10 minutes, at 0.5° elevation, immediately before the corresponding SHV mode scan. This results in a time offset between LDR and SHV mode scans of up to 6 minutes.

Figure 5.9 illustrates the limiting impact a 5-6 minute time offset could have on the representativity of an LDR-based algorithm. Panels (a) and (b) show 0.5° elevation reflectivity and LDR scans of a fast moving shower case near Holehead radar (appendix A). This is one type of case in which LDR is expected to be of significant benefit to QPEs, by correctly identifying very localised regions in which bright band corrections are not required.

The lower panels of figure 5.9 show how the available LDR values map onto “wet” pixels in two SHV mode scans: one the following minute (panel (c)), and the other 5 minutes later (panel (d)). The melting layer is marked by purple rings, and areas of rainfall where LDR is not available are visible as white gaps inside the black contoured regions. Figure 5.9c, with a 1 minute time offset, shows good agreement in rainfall positioning between the LDR and SHV mode scans. Raw LDR data is available for 79% of “wet” pixels in the SHV mode scan, and for 71% of the pixels in the melting layer region. A 71% availability of the criterion where it is needed is somewhat low, but is likely to increase with pre-processing (section 5.4.4, figure 5.11). However for the later image, where the scans are offset by 6 minutes, only 55% of “wet” pixels have a valid matching LDR value, and only 49% of the pixels in the melting layer region.

This example is illustrative of typical behaviour in fast moving convection, and justifies concern regarding the 10 minute frequency of operational LDR scans in the UK. In meteorological situations where LDR could deliver significant benefits through correct identification of convective showers, a time offset of 5-6 minutes between scans significantly limits the representativity of LDR, to the point where valid data are available for less than half of the spatial coverage required. For these reasons, an increase in LDR scan frequency to every 5 minutes is recommended for this application, ideally with a time difference of not more than 1 minute from the corresponding SHV mode scan.

Elevations

Due to the significant number of competing data requirements in the Met Office operational scan strategy, it is not feasible to generate a complete LDR volume to match every

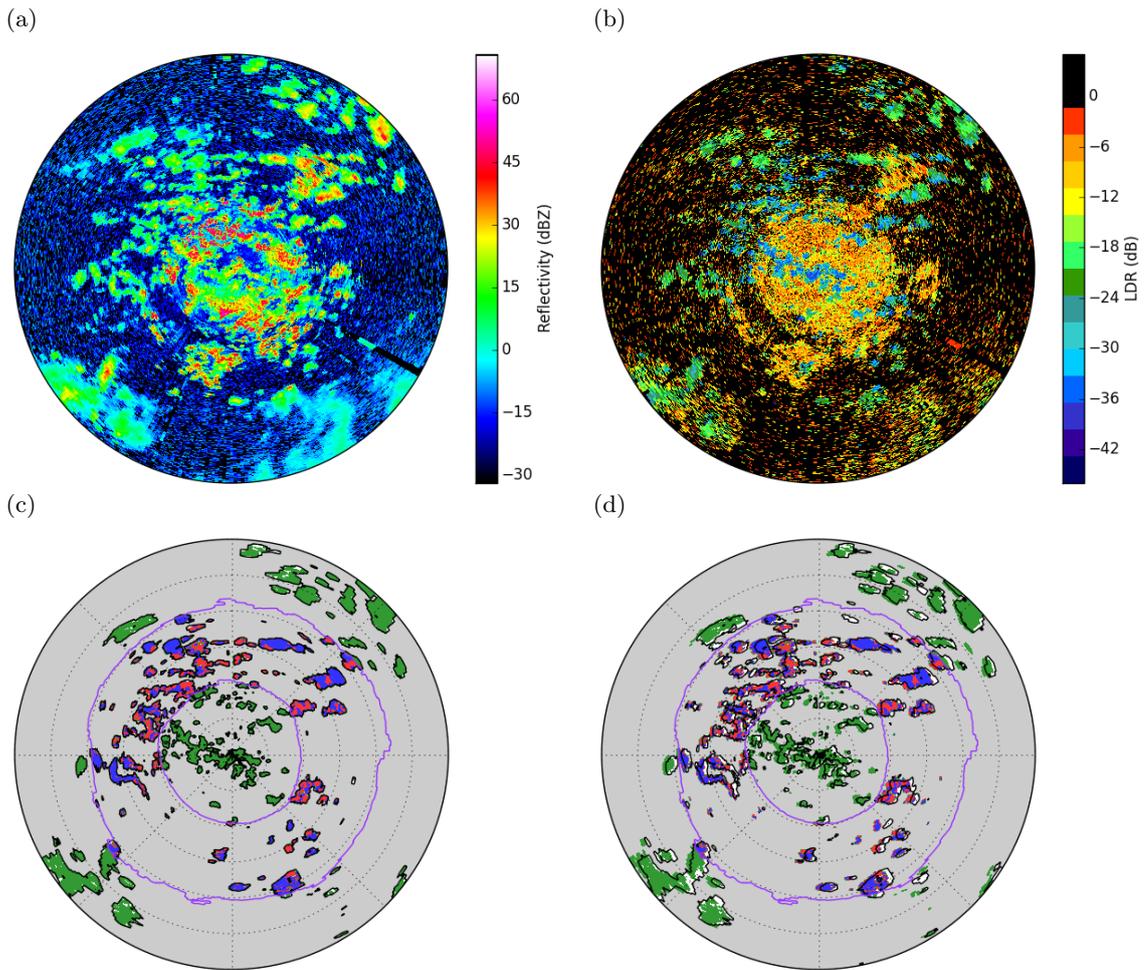


Figure 5.9: Top: Raw (pre-quality control) LDR mode parameters, (a) reflectivity and (b) LDR, from the scan at 13:49 on 25th July 2016. Bottom: mapping between LDR and SHV reflectivities at (c) 13:50 and (d) 13:54. All panels are plotted out to 180 km range (of the maximum 255 km). Black contours in panels (c) and (d) outline rainfall echoes (according to the SHV mode scan), showing their movement over the five minute period, and non-rain regions are shaded grey. Purple rings show the position of the melting layer. Within the melting layer blue regions have “bright band” LDR values and red regions are “non-bright band”. Valid LDR values outside the melting layer are shown in green. Areas of rainfall with no valid LDR are shown in white, inside the black contours.

operational QPE scan. To determine a minimum recommendation for scan elevations, LDR availability was considered in terms of coverage of the 1 km operational rain rate composite. “Coverage” is defined here as the proportion of pixels in the 1 km composite with access to a “matching” LDR measurement.

The 1 km operational QPE composite uses rain rates computed from the lowest usable scan at each radar. In a typical case, 90% of composite grid cells use data from 0.5° elevation radar scans. A further 7% of data is contributed by 1.0° scans. A 5 minute scan strategy including one 0.5° LDR scan from all radars would therefore provide 90%

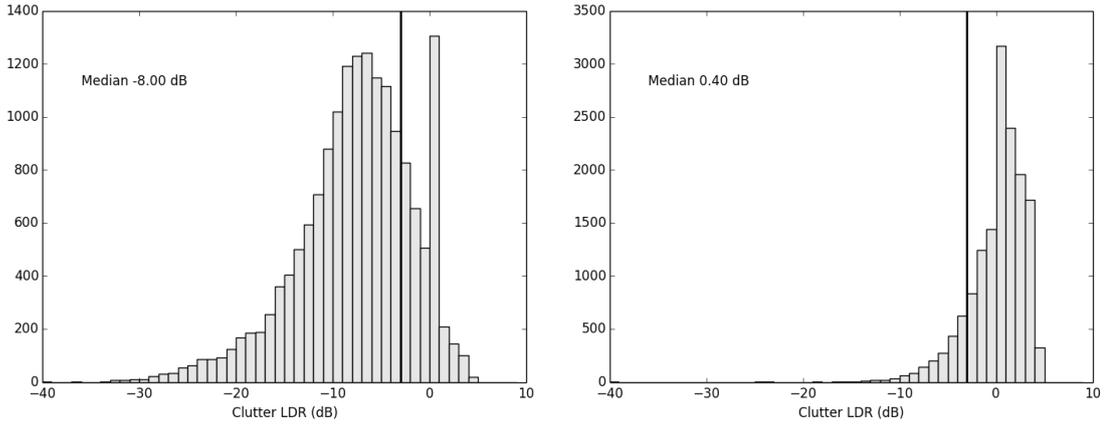


Figure 5.10: Example clutter LDR histograms for a functioning (left) and broken (right) waveguide switch. The threshold of -3 dB is shown by the black vertical line.

LDR coverage of the composite area. This can be considered sufficient to realise the majority of the benefits of LDR for operational QPE.

An additional 1.0° scan at all radars would increase operational LDR coverage from 90% to 97%. This is a significant enhancement, exploiting an increasing proportion of the benefits from LDR-based algorithms, but may not be achievable given the operational constraints. There also exist intermediate options which would increase coverage significantly without requiring two LDR scans at all radars. For certain radars - notably Chenies and Castor Bay (see appendix A) - a large proportion of the lowest elevation scan is unusable due to beam blockages. Including an 0.5° LDR scan at all radars, with additional 1.0° scans at Chenies and Castor Bay (but not elsewhere), would provide 93% LDR coverage of the 1 km composite area. Although not crucial to the success of the algorithm, such an increase in the availability of LDR is likely to increase the benefits to QPE across the UK network.

5.4.3 Quality control

An important part of radar data processing is quality control, in which echoes from non-meteorological sources are flagged for exclusion from further processing. The dual polarisation echo classifier applied to SHV mode data cannot be used for LDR scans. However, sophisticated quality control procedures for single polarisation reflectivity data are already implemented within Radarnet (Harrison et al., 2012; Sugier et al., 2002), and can be applied to the Z_h data collected in LDR mode. Since quality control flags (“meteorological” or “non-meteorological”) are associated with all data fields in a scan, the flags generated from reflectivity will also apply to the colocated LDR field.

An additional consideration for LDR data quality is the state of the waveguide switch,

which alters the radar transmission between single polarisation (LDR) and simultaneous transmission (SHV) modes. A broken waveguide switch produces a sensible LDR mode reflectivity scan, but because of contamination from the vertically polarised transmission no meaningful data is collected in the LDR field.

Waveguide switch problems can be identified by examining the measured values of LDR in clutter. Clutter can be identified in LDR mode data using the clutter phase alignment $CPA > 0.525$ (section 3.4.1). While the true LDR distribution in clutter is slightly negative, peaking between about -5 and -10 dB, the Z_{DR} distribution (absent a significant calibration offset) is centred on zero. Figure 5.10 shows an example of the differences in LDR clutter histograms for functioning and broken waveguide switches, which illustrates the very clear distinction between these cases. Considering the respective shapes of the clutter distributions, a median LDR value exceeding -3 dB in clutter was chosen to warn of problems with the waveguide switch. In fact, any threshold between -1 and -4 dB is likely to distinguish reliably between these cases. In real time processing, any scan with a median clutter LDR value greater than -3 dB raises a warning and is rejected from further processing.

5.4.4 Pre-processing to mitigate noise issues

LDR is measured by transmitting in the horizontal polarisation and comparing the magnitude of received echoes in the horizontal and vertical channels (equation 1.22). The depolarised echo component, Z_{vh} , is typically 2-4 orders of magnitude smaller than the background reflectivity, and is often close to the radar noise level. Where Z_{vh} is below the noise level, the noise subtraction applied at the radar site replaces the measured LDR with a “missing data” value.

For operational VPR classification, LDR values which are missing due to noise subtraction can be interpolated from surrounding valid LDR measurements. The radius of interpolation was chosen to be 2 km, which by inspection was sufficient to infer useful values at the majority of missing missing pixels.

Partly due to the closeness of Z_{vh} to the reflectivity noise threshold, gate-to-gate variation in LDR can also be very high. This is undesirable where fluctuations might lead to spurious classification of single pixels as non-bright band. A median filter was applied along each ray to smooth out this noise, while aiming to retain genuine variations in LDR due to atmospheric phenomena. An appropriate filter length was derived by maximising the covariance between LDR and a log-transformed version of ρ_{hv} ($L = -\log_{10}(1 - \rho_{hv})$), as defined in Keat et al. (2016)) from the adjacent SHV mode scan, since LDR and ρ_{hv} respond similarly to atmospheric phenomena but their noise characteristics should be independent. From testing on 3 upgraded radars with varying antenna quality, the best

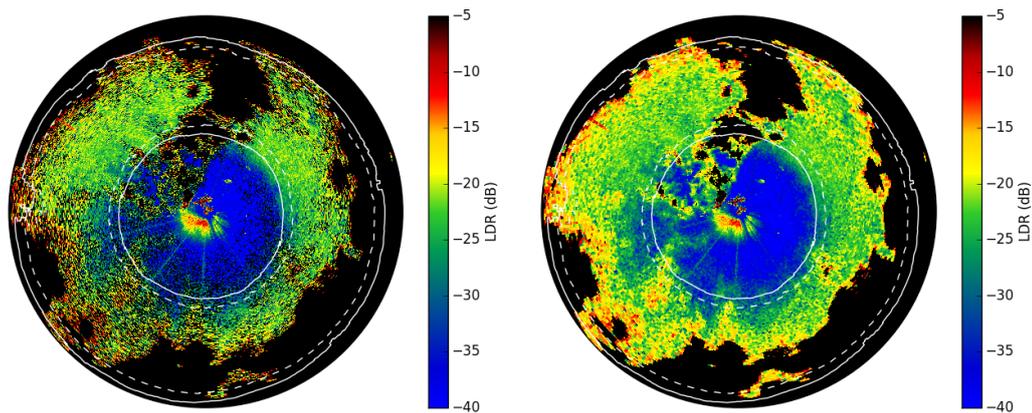


Figure 5.11: LDR from Predannack radar, 0.5° scan at 10:29 on 1st August 2016. Left: raw data; right: after interpolation and filtering. The white solid and dotted lines delineate expected melting regions by marking the ranges (from centre) at which the beam a) top is 700 m below freezing level, b) axis is 700 m below freezing level, c) axis reaches freezing level and d) bottom reaches freezing level (sampling regimes 1-3, see section 5.2.2).

results were obtained using a 3-gate (1.8 km) median filter.

An example of LDR before and after pre-processing is shown in figure 5.11. It can be seen that local features - both desirable and undesirable, such as the spokes in the South-West sector - as well as larger scale range behaviour are retained in the images after smoothing. Missing values of LDR at shorter ranges, due to the very low depolarisation in the rain, are also correctly interpolated. Note that this pre-processing only interpolates to pixels with missing data, and so does not replace the high LDR values measured at the edges of rainfall due to low signal-to-noise ratio.

Performing this interpolation and smoothing on LDR has an effect on the precision and spatial resolution of the processed fields. The raw measurement resolution of 600 m by 1° is smoothed to a resolution of approximately 2 km. The impact of this smoothing should be considered in light of the time offset between LDR measurements and the corresponding SHV mode scan, and in the context of other data used as part of correction for VPR.

Low elevation (0.5 and 1.0°) scans at the Met Office are performed at 8.4° per second, with a total duration of 43 seconds. If an 0.5° elevation LDR mode scan is performed at the optimal time, immediately before the 0.5° SHV mode scan, then the values are offset in time by 43 seconds. For precipitation moving at a typical wind speed of 5 m s^{-1} (Sinden, 2007), the spatial offset in 43 seconds is 215 m - around ten percent of the 2 km smoothing window. So the proposed smoothing results in a loss of precision that is larger than the expected spatial offset between the two scans. However, there is precedent for using lower resolution metadata to support correction for VPR, for example the use of 5 km gridded model wet bulb freezing levels to locate the top of the melting layer

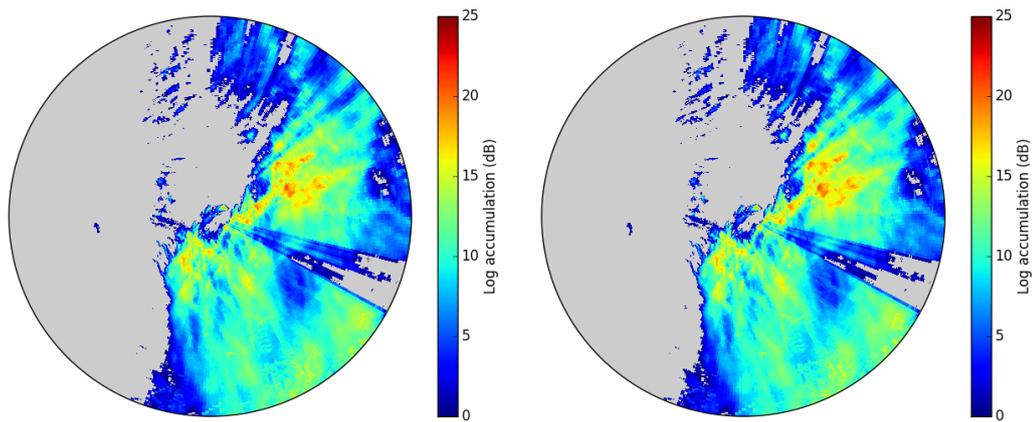


Figure 5.12: 24 hour rain accumulations from 16:00-16:00 on 15th-16th September 2016 for the control (left) and LDR algorithm trial (right). There is little visible difference between the two accumulations - the differences are analysed more closely in figure 5.14. Accumulations are given in dB units: $10\log(R/\text{mm})$. The maximum point radar accumulation for the event was 146.9 mm.

(Kitchen, 1997; Mittermaier and Illingworth, 2003). Given the alternatives of missing data and spurious misclassifications, the degradation of LDR measurements to a 2 km equivalent spatial resolution is judged acceptable for this application.

5.5 Evaluation of Radarnet implementation

Having defined the data requirements and necessary preprocessing of LDR, a trial implementation of LDR-based classification was coded into the Radarnet development environment. The impacts on QPE of this prototype algorithm were evaluated both for a case study and statistically over a longer trial period, by comparing radar-derived rain rates with colocated hourly accumulations from a network of 0.2 mm tipping-bucket rain gauges. The analysis was performed using PPI data from the Wardon Hill research radar, as this was the only installation from which LDR data availability met the requirements of section 5.4.2.

5.5.1 Case study evaluation

A 24 hour rainfall event beginning at 16:00 UTC on 15th September 2016 was used to analyse the impacts of LDR-based VPR classification on surface rain rates. This convective event included rainfall intensities that would be considered extreme in the UK, from a relatively stationary weather system that caused very high accumulations in some locations (figure 5.12). The position of maximum rainfall intensities with respect to

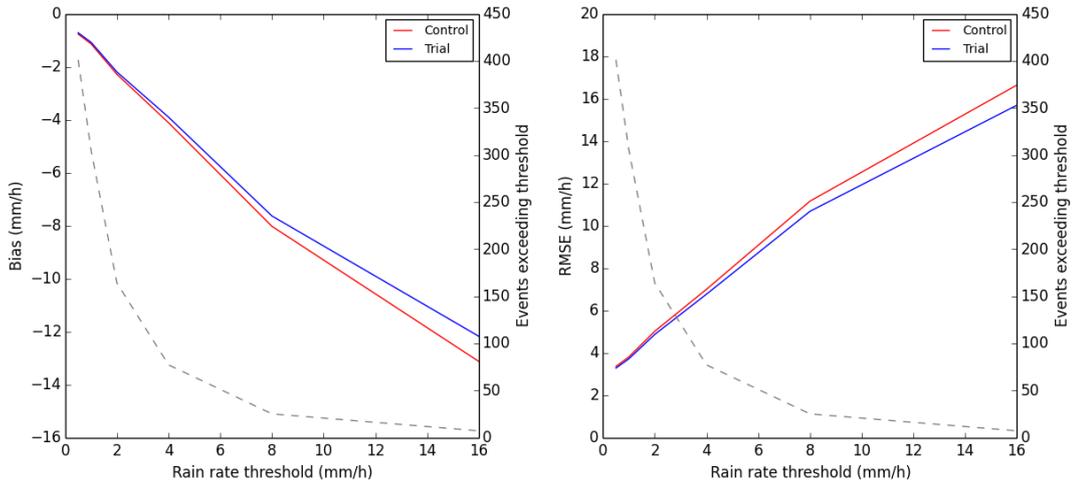


Figure 5.13: Bias (left) and RMSE (right) of hourly radar accumulations over the 24 hour trial compared to those of colocated gauges, where the gauge accumulation exceeds a certain threshold. Event counts are shown by the dashed grey line. There is a small but noticeable reduction in both bias and RMSE for the LDR algorithm trial at all intensities.

Wardon Hill, coupled with a 0° isotherm height of around 3 km, provided ample sampling of the melting layer for the duration of the event.

Reflectivity volumes from Wardon Hill were processed first using the current Radarnet convective diagnosis algorithm ($Z_1 > 30$ dBZ) only, and then with the new algorithm: using melting layer LDR (in regimes 1 and 2) to identify the presence or absence of bright band. In each case, where the operational or LDR-based criterion identified “no bright band”, a constant VPR ($Z(h) = Z_s$) was assumed.

A pilot investigation using this case study discovered some instances in which the LDR-based criterion caused a reduction in surface rain rates. On inspection of other radar PPIs covering this event, which included high reflectivity values above and low ρ_{hv} values below the melting layer, it was deduced that these results were likely attributable to hail being sampled at the melting level. In this case high reflectivity measurements persisted at high levels above the freezing level, but the LDR in sampling regimes 1 and 2 (section 5.2.2) did not distinguish correctly between melting hail and stratiform bright band, leading to a bright band correction being applied. For this reason it was decided to allow positive identification of non-bright band precipitation from either LDR or Z_1 in regimes 1 and 2, in order to diagnose convection correctly in the specific case where hail is present.

The results of the prototype LDR-based algorithm (including the hail contingency) compared to the operational reflectivity-based convective diagnosis scheme are presented in figures 5.13 and 5.14. Figure 5.13 shows the bias and RMSE (appendix B) of hourly radar accumulations from the event. The LDR algorithm reduces negative biases, partic-

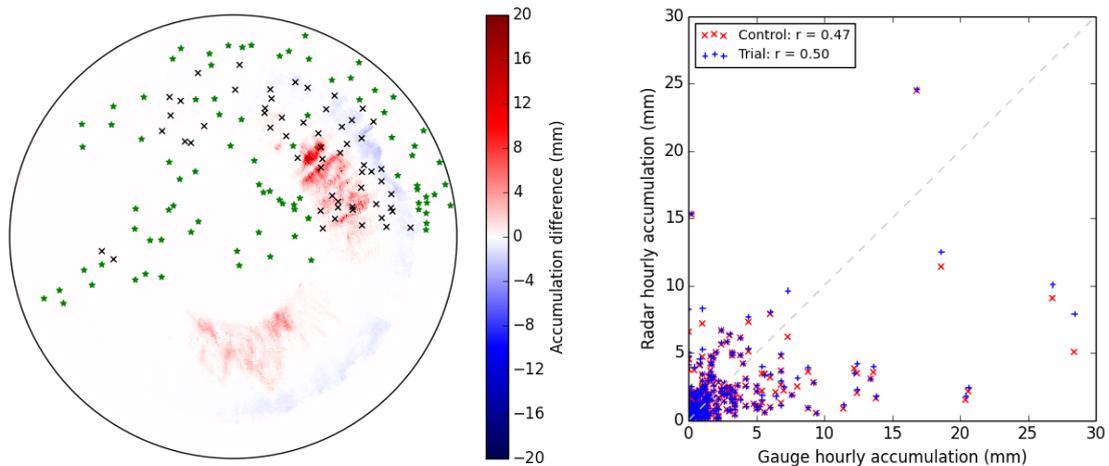


Figure 5.14: Left: total radar accumulation differences (trial minus control) with all rain gauge positions marked, to a maximum range of 255 km. Black crosses show the 58 gauges with near-complete timeseries at which the total radar event accumulations differed. Right: scatterplot of the hourly accumulations at these 58 gauges where LDR-based VPR classification produced a change in the radar accumulation.

ularly at high rain rates. This is expected, since the primary impact of LDR should be to reduce the inappropriate correction for bright band which contributes to these large negative biases. The impact of this on RMSE is noticeable, with the error on the highest intensity events reduced by almost 1 mm h^{-1} .

It should be remarked at this point that the biases in this case, although large, are not atypical of statistics for the UK network. This is partly due to the use of a fixed Marshall-Palmer (Marshall and Palmer, 1948) reflectivity-rain rate (ZR) relationship to calculate rain rate. The true ZR relation is dependent on rain drop size distribution (DSD, section 1.2.2), and the Marshall-Palmer relation is known to underestimate rain rates in compact ice and convective conditions, which contain a relatively large population of small rain drops (Bringi et al., 2009; Matrosov et al., 2016). For the event of 15th-16th September, independent disdrometer measurements suggest the DSD was more uniform than Marshall-Palmer (Thompson 2016, personal communication). This goes some way to explaining the particularly large negative biases during this event.

Figure 5.14 evaluates the changes in total event rainfall accumulation when LDR is used to inform the VPR correction. The left hand panel shows the change in accumulation (in mm h^{-1}) resulting from LDR-based classification. Differences are confined to the region (approximately 100-200 km in range) in which the radar sampled the melting layer. As expected, LDR-based VPR classification generally results in increased rain rates through avoiding inappropriate correction for bright band. The left hand panel of figure 5.14 also shows the positions of all rain gauges used in the evaluation of QPEs from Wardon Hill. Black crosses indicate the 58 gauges at which radar accumulations differed for this case

study, and which had near-complete timeseries of hourly accumulations throughout the event (at least 22 out of the 24 hours available). Green stars show the locations of all other reporting rain gauges.

Alongside the preservation of high rain intensities from non-bright band events, there is a small reduction in this case in rainfall accumulations at longer ranges. These represent an area over which LDR diagnoses “non-bright band” conditions, but the current Radarnet non-bright band VPR shape does not account for any decrease in reflectivity above the melting layer. This result, coupled with the average shapes of VPRs observed in the high resolution dataset (chapter 4, figure 4.6), indicates a need to reconsider the shape used for VPR correction in non-bright band conditions. The potential to correct more effectively for non-bright band VPRs is explored in chapter 6.

The right hand panel of figure 5.14 plots the hourly radar accumulations that differ between trial and control against the colocated rain gauge accumulation. This is to determine whether the changes to QPE accumulations - both positive and negative - represent an improvement in radar accuracy with respect to “ground truth”. The majority of changes due to LDR, particularly at the higher hourly values, move the radar accumulation closer to that of the colocated gauge. This can be seen both in individual points on the scatterplot and the increase in radar-gauge correlation for the LDR trial.

5.5.2 Gauge-radar statistical trials

In addition to the case study presented above, a controlled trial was run to evaluate QPEs over the month of September 2016. This longer period (totalling approximately 3 weeks, due to radar outages) was used to verify that the overall statistical performance of Radarnet QPEs was not in any way degraded by the new LDR algorithm. This is a standard final check before submitting code to the operational Radarnet system.

Figure 5.15 shows the impact of the new algorithm on performance over the longer term. Although changes are small, there is evidently no detriment to rainfall statistics over this monthly period, with some indication that errors are reduced at higher accumulations. The algorithm can therefore be considered safe for operational implementation.

5.6 Conclusions

The use of LDR for VPR determination has the potential to deliver significant benefits to operational QPEs. This chapter developed the theoretical finding from chapter 4 into an LDR-based classification algorithm suitable for use in operational radar data processing

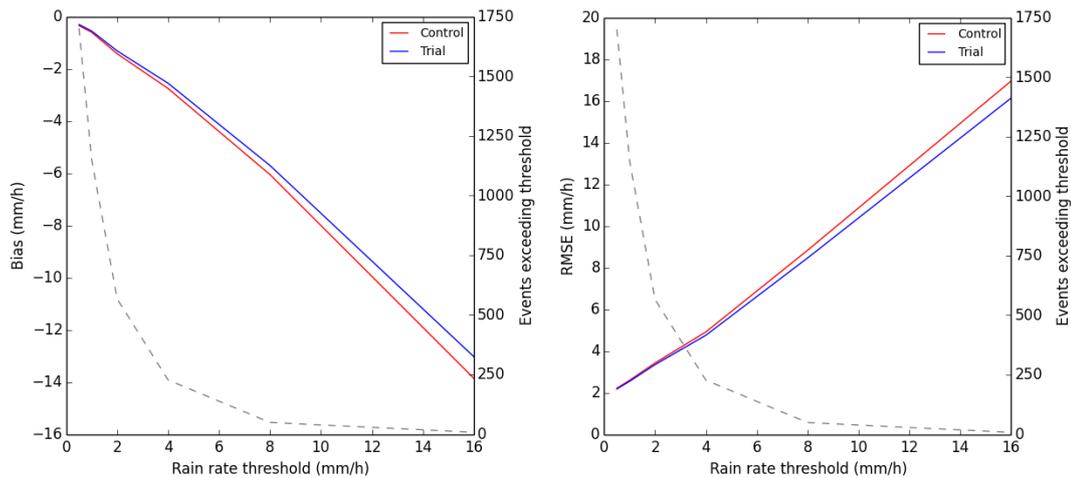


Figure 5.15: Bias (left) and RMSE (right) of hourly radar accumulations over the month of September 2016 compared to those of colocated gauges, where the gauge accumulation exceeds a certain threshold. The improvement seen in the 24 hour case study persists over this longer trial period, showing that the benefits of the new algorithm outweigh any potential adverse effects.

systems. As a first step, section 5.2 established the skill of broadened LDR measurements in identifying non-bright band VPRs through a simulation study. A classification threshold of -26 dB was derived, with higher thresholds indicated within 41 km of the radar, where the radar beam is likely narrower than the depth of the melting layer. The impacts of using these classification thresholds were simulated in section 5.3 using a high resolution vertical profile dataset from the Wardon Hill research radar (chapter 3), and showed significant potential for benefits to real time QPEs.

Section 5.4 considered the data quality, availability and quality control procedures required for operational implementation of an LDR-based algorithm. Gross error checks and noise mitigation were discussed, along with minimum requirements on scan strategy and timing. A prototype classification algorithm and associated pre-processing were implemented within the Radarnet development framework, and were shown to benefit Wardon Hill QPEs in both a high impact case study and over a longer evaluation period (section 5.5). Network wide testing and a potential operational implementation of this algorithm are contingent on a change to the Met Office radar scan strategy to include an 0.5° LDR scan at 5 minute frequency.

The correct identification of a substantially increased proportion of non-bright band cases using LDR has demonstrable benefit to QPEs through avoiding inappropriate correction for bright band. However, the case study in section 5.5 suggests that further work may be needed to reap the full benefits of LDR-based classification. The decrease in rain accumulations at longer range, where the radar samples only the very top of the melting layer, is an inevitable result of the constant reflectivity profile ($Z(h) = Z_s$) assumed for

non-stratiform precipitation in the current UK VPR scheme. Although there are not enough rain gauges in this case to comment on the nature of this change (as good or bad), given the average profile shapes observed in the high resolution dataset (figure 4.6), it is clear that a constant reflectivity with height is not representative of the majority of non-bright band situations in the UK. An obvious area for future work, therefore, is to investigate a more suitable non-bright band shape for the correction of reflectivity measurements in an operational environment.

Chapter 6

Introducing a VPR correction for non-bright band precipitation

6.1 Introduction

The standard framework for vertical reflectivity profile (VPR) classification and correction assumes that precipitation is either stratiform (with bright band) or convective. Given that stratiform bright band is the most prominent feature affecting radar precipitation estimates at high latitudes, substantial efforts have been made to characterise and correct for this feature (eg Kitchen et al., 1994; Andrieu and Creutin, 1995a; Tabary, 2007; Zhang and Qi, 2010; Koistinen and Pohjola, 2014), some of which have been incorporated into operational radar rainfall estimation (QPE) schemes.

By contrast, the lack of research into non-stratiform reflectivity profiles is striking. Smyth and Illingworth (1998) show that convective precipitation profiles are extremely variable and difficult to characterise, which is consistent with the definition of convection as dominated by strong vertical motion (Steiner et al., 1995). It follows from such papers that correction for VPR in convective conditions is unlikely to be feasible. As a result, and in the context of a simple stratiform-convective classification framework, the nature of non-bright band precipitation profiles has been largely neglected by the existing literature.

Recent studies have called into question the validity of the assumption that the range of observed reflectivity structures can be fully described by only two generic types of VPR. As early as the 1990s, Fabry and Zawadzki (1995) observed and characterised five different types of VPR using vertically pointing radar. More recently, Matrosov et al. (2016) presented observations of a “non-bright band” (NBB) VPR distinctly different from both stratiform and convective profiles. The drop size distribution of rain at the surface from NBB profiles was found to be dominated by smaller drops than either

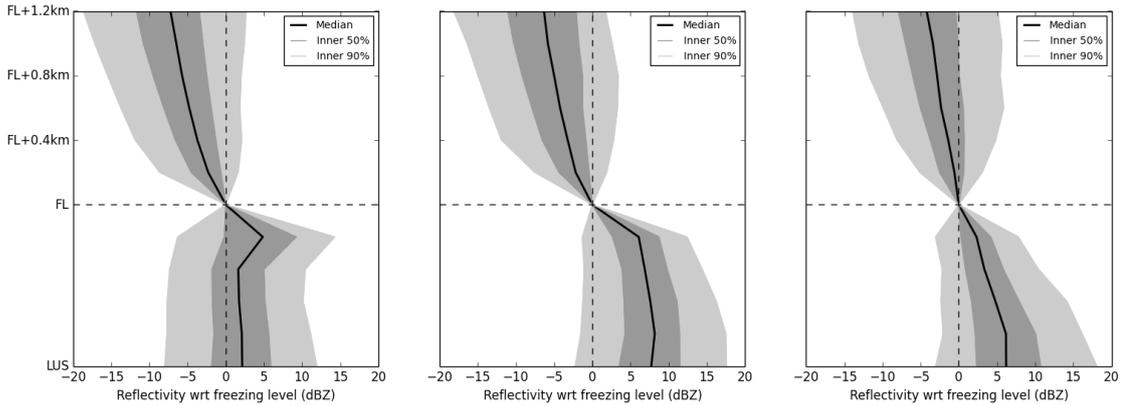


Figure 6.1: Average stratiform, compact ice and convective reflectivity profiles with height relative to the model derived (Brown et al., 2012) wet bulb freezing level. (Repeat of figure 4.6.)

of the other two rainfall types. Not only is this finding consistent with the “compact ice” hypothesis of Fabry and Zawadzki (1995), but the NBB profile shape is also in good agreement with independent observations of the compact ice profile (Fabry and Zawadzki, 1995; Sandford et al., 2017). This research suggests that VPRs without bright band do not only occur in cases of deep convection, but in more structured meteorological conditions, which could potentially be characterised and constrained.

In chapter 4 of this thesis, VPRs measured at high resolution from the Wardon Hill research radar were shown to have a bright band in 84% of cases. In the operational Met Office radar processing system (Radarnet), the remaining 16% of profiles would be corrected for beam broadening assuming a constant reflectivity profile with height ($Z(h) = Z_s$). However, of the observations, neither compact ice (10%) nor convective (6%) VPRs are well described by this constant reflectivity profile (figure 6.1). The Wardon Hill dataset provides opportunities to observe and investigate this lack of correspondence between the corrections applied and the true nature of non-bright band VPRs.

As in the case of inappropriate bright band correction, the assumption of a constant reflectivity profile in correcting for cases with more systematic height-based behaviour can have a significant adverse impact on QPEs. In chapter 4 the impact of inappropriate bright band correction was illustrated by fitting the Kitchen et al. (1994) stratiform profile to a compact ice VPR at different ranges from the radar (figure 4.1). Figure 6.2 shows QPEs from the same compact ice VPR when corrected using a constant reflectivity profile, which replicates the current operational treatment of non-bright band profiles in Radarnet. Using this “convective” profile shape there is no correction for bright band, leading to better rainfall estimates than figure 4.1 in the 70-130 km range. However, failure to represent the changes in reflectivity within and above the melting layer leads to increasing underestimation at long range. In this case, rain rates at 150 km are no better estimated using the convective shape than from a stratiform bright band profile. The

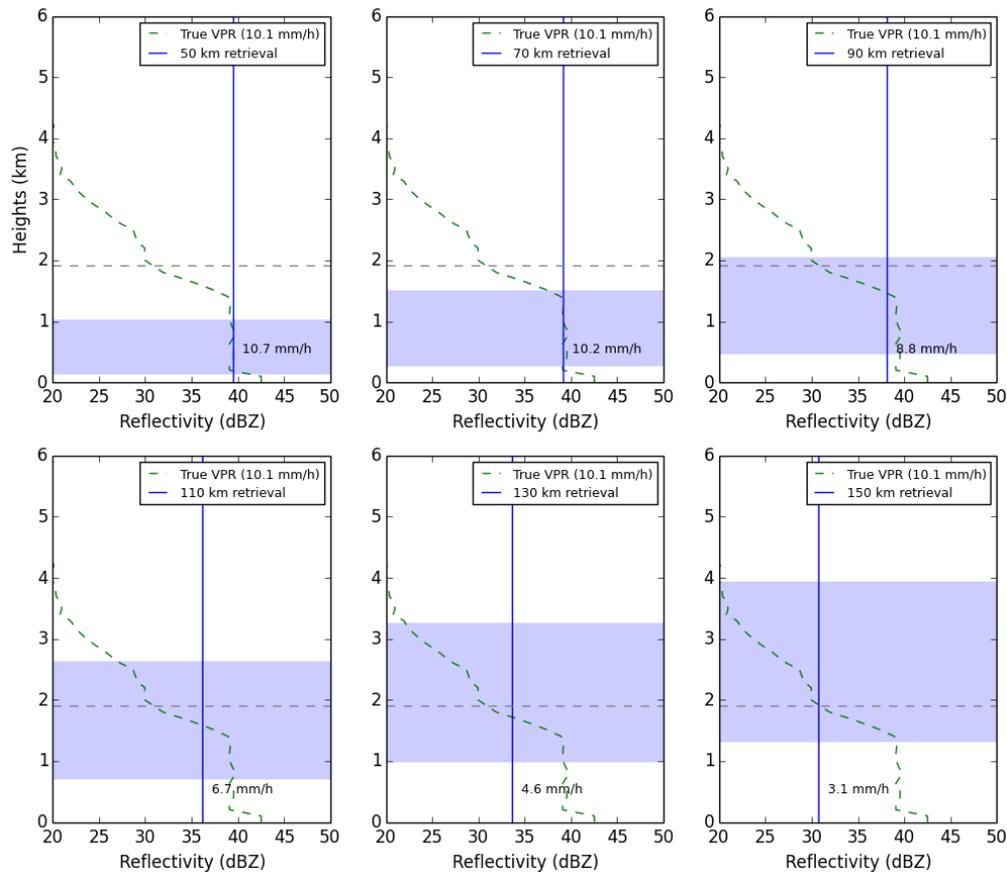


Figure 6.2: Example of the impact of a badly-fitted VPR shape on surface QPEs in a non-bright band case. QPEs from the same atmospheric VPR as in figure 4.1 are derived using the iterative method of Kitchen et al. (1994) to account for beam broadening, using a constant reflectivity with height idealised profile. The blue line shows the fitted VPR, the green dashed line the true atmospheric VPR, and the purple shading shows the position and extent of the (broadened) radar beam at that range. The dashed horizontal line shows the height of the 0°C isotherm.

same impacts from neglecting the shape of non-bright band profiles were also apparent in the case study analysis of chapter 5, where inaccurate representation of the VPR above the melting layer was responsible for a reduction in long range QPEs in non-bright band conditions.

It is clear from figure 6.2, along with the average profile shapes from the Wardon Hill sample (figure 6.1), that the current treatment of non-bright band profiles could be responsible for significant inaccuracies in real time QPEs. To gain the full benefits of improved VPR classification using LDR (chapters 4 and 5), there is a clear need to improve the characterisation and correction for non-bright band precipitation profiles.

This chapter develops and evaluates a selection of possible shapes for determination and correction of non-bright band (compact ice and convective) VPRs. For consistency with the LDR-based classification algorithm (chapter 5) on which this chapter builds, the aim

is to develop a single idealised profile shape to represent both types of non-bright band profile. Section 6.2 reviews the limited existing literature in relation to non-bright band VPRs, and uses this to suggest a number of possible profile shapes for use in non-bright band conditions. In section 6.3 the candidate profiles are evaluated via a simulation study, using the 1080 high resolution non-bright band VPRs available from the Wardon Hill dataset (chapter 3). The most skilful of the new VPR shapes is then implemented within the Radarnet development framework, where its impact on QPEs is evaluated with respect to rain gauge accumulations in a known convective case study. Conclusions are presented in section 6.5.

6.2 Constructing profile shapes

The current VPR literature (discussed in chapter 2) is substantially focused on correction for bright band in stratiform conditions. This is understandable, as the bright band is the most frequently occurring feature of vertical profiles with a significant impact on surface QPEs. However, there is a small subset of the literature that does address non-bright band profiles, usually as part of a more comprehensive determination or correction scheme. Some observational studies have also characterised or described a complete set of different types of VPR.

This section uses information from the current literature and the dataset from Wardon Hill to suggest profile shapes that could more accurately represent the reflectivity structure of non-bright band precipitation. In section 6.2.1 qualitative observations from the literature are used to describe and justify four candidate VPRs (figure 6.3). Quantitative features of these profiles are then estimated using the high resolution Wardon Hill dataset. The candidate profiles are evaluated with respect to each other and to the current Radarnet profiles in section 6.3, where the profile with the highest skill is selected for further testing.

6.2.1 Existing literature

Some of the earliest studies of VPR indicate a high level of variability and low predictability for convective reflectivity structures, particularly above the melting layer (Smyth and Illingworth, 1998). Based on this, the current operational Radarnet VPR scheme (section 2.5.3) assumes homogeneity in the VPR ($Z(h) = Z_s$ for all h below the precipitation top) in non-bright band conditions. This “control” profile is included in figure 6.3 as option *a*, and is equivalent to making no correction for the vertical profile, except at the very longest ranges where the effect is to correct for partial beam filling (figure 1.8). Given its limited effect the explicit use of a homogeneous convective VPR is not widespread, but

is applied by Goudenhoofdt and Delobbe (2016) as part of an end-to-end assessment of their radar rainfall estimation system.

The observations of Smyth and Illingworth (1998) also suggest possible alternative shapes for convective VPRs. The authors examine and categorise bright band and non-bright band profiles according to surface reflectivity (Z_s), rather than normalising with respect to the zero degree isotherm or the reference level. They find that reflectivity is approximately constant with height for average graupel profiles where Z_s is in the range 30-35 dBZ (their figure 7), which is consistent with the profile used in Radarnet. However, more severe convection and higher surface reflectivities are found to be associated with profiles having a constant reflectivity gradient with height of about -2 dB km^{-1} (their figure 8 and section 4c). This profile matches more closely the high resolution observations of convective VPRs from Wardon Hill (figure 6.1). A constant gradient profile has therefore been chosen as the first test shape for an idealised non-bright band VPR (figure 6.3, option *b*).

Two observational studies provide support for the next test profile shape. Both Delrieu et al. (2009) and Kirstetter et al. (2010) perform analyses of data from the Bollène 2002 experiment, studying normalised apparent VPRs in convective and stratiform cases at ranges of up to 60 km. Figures from these papers suggest a non-bright band vertical profile characterised by constant reflectivities in the rain and decreasing reflectivity with height above the freezing level (Delrieu et al. (2009) figures 7, 8 and 10; and Kirstetter et al. (2010) figures 2, 3, 5 and 6). This shape is used as test option *c*.

A final profile option is informed by both observational and modelling studies. Using observations from a vertically-pointing X-band radar, Fabry and Zawadzki (1995) describe a sharp drop in compact ice profile reflectivities across the melting layer due to the change in dielectric factor between rain drops and small ice particles (as explained in section 4.2.4). This is similar to the non-bright band profile shape of Matrosov et al. (2016). Independently of this, the simplified microphysical framework of Kirstetter et al. (2013) yields non-bright band VPRs with constant reflectivity in rain, decreasing reflectivity in the upper part of the melting layer, and a gradient not dissimilar to that of stratiform profiles in the ice (their figures 3, 4 and 7). For this chapter, therefore, test profile *d* retains the constant rain reflectivity and ice gradient from option *c*, but with an additional drop in reflectivity across the top half of the melting layer. This is the shape that most closely resembles the compact ice profiles in the high resolution dataset (figure 6.1).

The four non-bright band VPR shapes for testing are illustrated qualitatively in figure 6.3, where the horizontal dashed line indicates the onset of melting at the model wet bulb freezing level (the 0°C isotherm). The lack of quantitative axes reflects the variability that results from scaling these profiles to the measured reflectivities and ancillary data

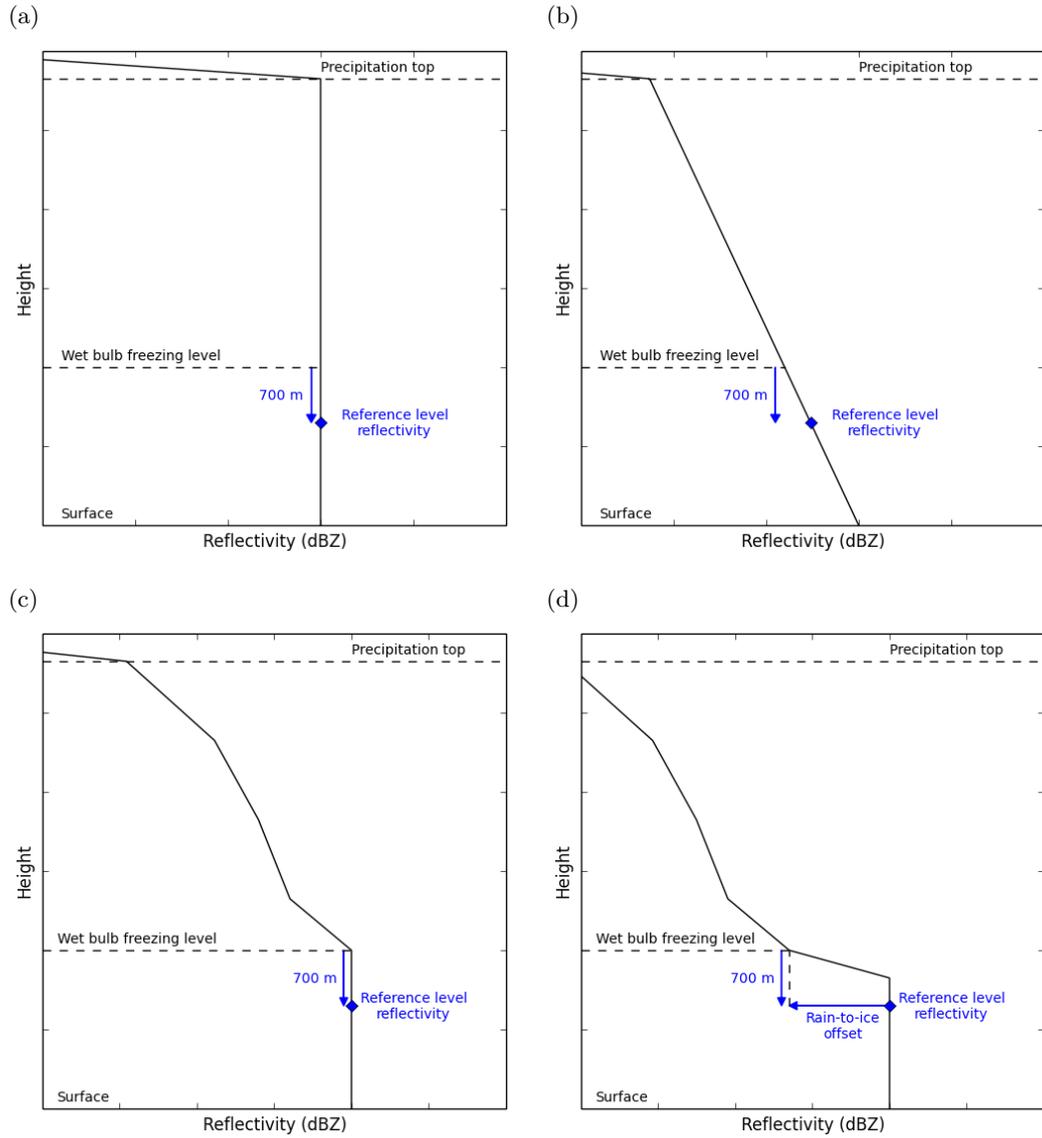


Figure 6.3: Sketches of non-bright band VPR shape options, from top left: (a) constant Z with height, (b) constant Z gradient with height of -2 dB km^{-1} , (c) constant Z in the rain and stratiform gradient in the ice, and (d) as for option *c*, with an additional dB offset across the melting layer. See section 6.2.2 for quantitative details of profiles *c* and *d*.

(cloud top and freezing level heights), following the implementation and climatological constraints of Kitchen (1997). The remaining quantitative parameters required to test these profiles are obtained from the Wardon Hill VPR dataset, as described below.

6.2.2 RHI profile observations

Due to the nature of the existing literature, the potential non-bright band VPR shapes inferred in section 6.2.1 are, of necessity, qualitative. In particular, although the reflec-

tivity gradient above the melting layer is known to be negative, no quantitative values or ranges for non-stratiform situations are available from previous studies. Equally, there is no specific constraint on the offset in profile d between the reflectivities in the rain and ice layers. These parameters were estimated using observations from the high resolution RHI dataset.

Given the substantial variability of the observed VPRs, the size of the non-bright band sample from Wardon Hill (1080 profiles) was not considered sufficient to split into a “training” and “verification” dataset. This presents a challenge in terms of avoiding circular reasoning. Since the whole “training” dataset will later be used in simulation studies to select the most promising candidate VPR, it is important that no VPR shape is based too quantitatively on these observations. It would not be appropriate, for example, to calculate an average gradient above the melting layer from these measurements, and then to test that gradient for VPR determination on the same profiles from which it had been calculated. However, since the most promising candidate shape will then be further tested on independent radar observations and gauge readings, it is safe to derive some limited quantitative constraints on profile shapes from the RHI dataset.

Ice reflectivity gradients

The average shapes of VPRs of each of the three different types are shown in chapter 4, figure 4.6. The differences between the three profile types are largely confined to the region around the freezing level and below. Both non-bright band profiles have an average reflectivity offset of 6-8 dB between the freezing level and the ground surface, as opposed to the stratiform profile, which has an offset of around 2 dB. However, above the melting layer, all three profiles show a similar decrease in reflectivity with height. For this reason it was decided that rather than modelling a new gradient profile above the melting layer, test profiles c and d would use the same ice reflectivity gradient as the Kitchen et al. (1994) stratiform profile. The gradient is dictated according to the cloud top height, approximated by satellite cloud top measurements (Kitchen, 1997), at which the reflectivity (in linear units) is defined to be zero. Reflectivity values are defined at height levels 1, 2, 3, and 4 km above the wet bulb freezing level, using climatological reflectivity ratios derived by Kitchen et al. (1994) for each of these four ice cloud depth scenarios. In the unusual scenario where the cloud top is more than 4 km above the wet bulb freezing level, the 4 km ice depth profile is assumed.

Rain-to-ice reflectivity offset

The option *d* non-bright band profile in figure 6.3 has a reflectivity offset across the melting layer which has not yet been quantified. In the absence of clear predictive correlations within the dataset (eg between surface reflectivity and rain-to-ice offset), a fixed climatological offset was determined using properties of the 1080 high resolution non-bright band profiles.

The distribution of rain-to-ice offsets in the non-bright band VPR dataset is shown in figure 6.4. The grey dashed vertical lines show the 5th, 25th, 50th, 75th and 95th data percentiles (henceforth denoted p_x). The data fit well to a normal distribution with a mean $\mu = 6.5$ dB and standard deviation $\sigma = 4.6$ dB. This value of μ is consistent with the independent microphysical explanation for compact ice profiles (Fabry and Zawadzki, 1995), which predicts an increase in reflectivity with melting of 6-7 dB. Consistency with previous literature provides increased confidence in the suggested offset value.

The parameters of this normal distribution were not fitted automatically, but were calculated using the inherent symmetry of the distribution, and the fact that 68% of normally-distributed data lie within 1σ of the mean. The mean and standard distribution were therefore estimated as:

$$\sigma = 0.5 \times (p_{84} - p_{16})$$

$$\mu = p_{25} + 0.5 \times (p_{75} - p_{25})$$

Use of the inner part of the distribution rather than a statistical “best fit” is appropriate here, because the lower tail of the distribution is truncated due to the methods used for finding peaks in the high resolution VPR sample - notably discarding profiles for which an unambiguous melting layer could not be found (section 4.2.1). If fitted automatically, this truncation would tend to skew the fit towards higher μ and slightly lower σ . The calculated μ value using the percentiles method in this case is almost identical to the median p_{50} value (box plot, figure 6.4).

In using the constant mean of this distribution for an operational profile correction, it is worth considering the errors on this value, and the resulting uncertainty in surface rainfall estimates. While the spread of these data, at 4.6 dB, is large, it is smaller than the mean

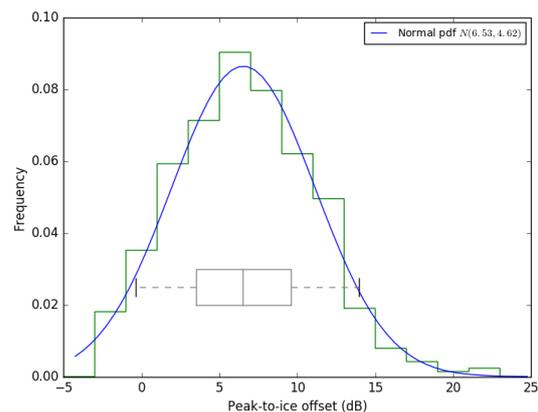


Figure 6.4: Histogram and box plot of rain-to-ice values from the RHI VPR dataset fitted to a normal distribution. The box whiskers span the 5th to 95th percentiles. The central value of 6.5 dB was used in profile option *d*.

offset value itself. A residual random error of approximately a factor of 3 in reflectivity (or a factor of 2 in rain rate, assuming a Marshall-Palmer relation $Z = 200R^{1.6}$) is smaller than the *a priori* underestimation bias of almost a factor of 5 (2.5 in rain rate). Applying this constant mean offset should therefore reduce the overall rain rate bias, but may not significantly reduce the random error component.

6.3 Simulation studies using high resolution VPRs

In chapter 5 a simulation framework based on Kitchen and Jackson (1993) was developed to compare the performance of LDR and high-level reflectivity in classifying VPRs. Measurements were simulated for a 0.5° elevation beam at different ranges, and corrected using the two operational Radarnet VPR shapes via the Kitchen et al. (1994) iterative convergence scheme. The results were evaluated by comparing the resulting surface QPEs with the true surface rain rates from each profile (section 5.3).

In this section, the same simulation framework is applied to evaluate the different non-bright band profile shapes in figure 6.3. Each of the four profile shapes is used to estimate surface QPEs from simulated measurements from the 1080 non-bright band profiles from the Wardon Hill dataset. The accuracy of these QPEs is compared initially between the four profiles, and then with reference to a wider selection of “controls”, including QPEs obtained using the current Radarnet stratiform VPR. Further analysis of QPEs from the chosen profile is then presented in terms of the overall distribution of errors, and the trend in uncertainties with range from the (simulated) radar.

6.3.1 Profile selection

To choose the best of the test profile shapes for a Radarnet implementation, the simulation method described above was applied to all non-bright band profiles in the climatological VPR dataset. Reflectivity measurements from a 0.5° elevation beam were simulated from each VPR at 20 km range intervals, and were corrected using the iterative deconvolution method of Kitchen et al. (1994) using each of the four profile shapes in figure 6.3. The errors in derived surface reflectivity at each range, converted to rain rate units for ease of interpretation, are shown in figure 6.5.

The different candidate profile shapes show different strengths and weaknesses at different ranges. At short range (< 50 km) the simulated radar beam is below the melting layer for most of the VPRs in the dataset, so there is very little difference in the performance of the four idealised profiles. The only profile with non-zero gradient in the rain layer is option *b*, which has a constant -2 dB km^{-1} reflectivity gradient throughout. In terms

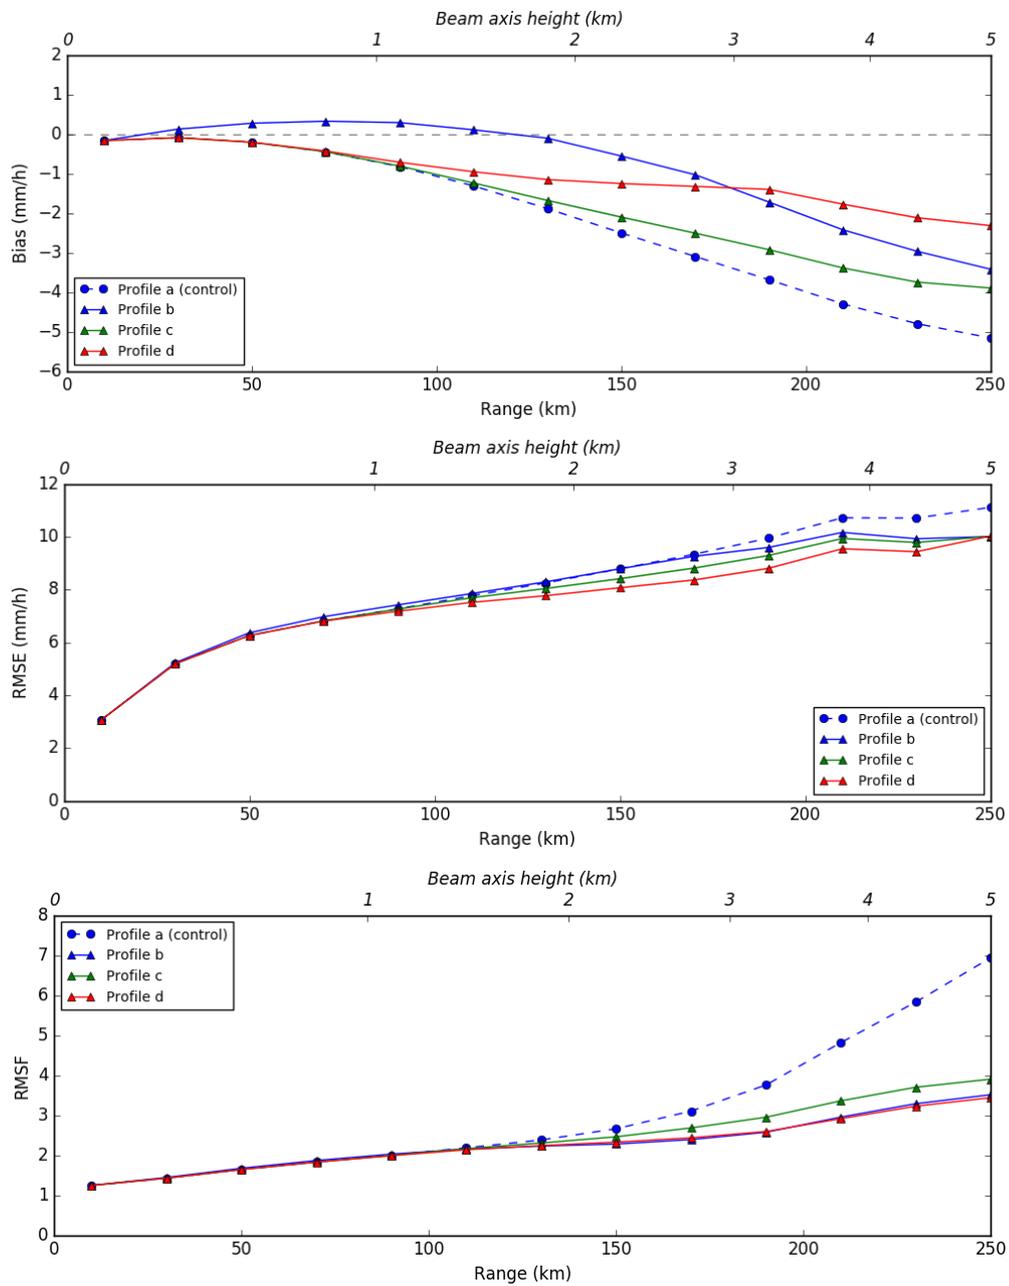


Figure 6.5: Bias, RMSE and RMSF (appendix B) with range of surface rain rates calculated, using the control (option *a*) and test profiles from figure 6.3 to correct reflectivities from a 0.5° elevation beam. Options *b* and *d* perform equally well in terms of bias and RMSF, but option *d* is best for RMSE. All test profiles outperform the control.

of bias, this profile performs better than the others between 70 and 170 km from the radar. However, the overall behaviour of the constant gradient profile with range is not consistent with a good fit. The profile overestimates rain rates slightly at short to medium range, but then declines sharply, showing increasing underestimation at long range. This implies that a gradient of -2 dB km^{-1} is on average too steep below the melting layer, and too shallow above it. If the profile had either over- or underestimated consistently at all ranges, this might have been attributable to an inappropriate gradient rather than an unsuitable shape, and the performance might have been improved by tuning. However, given the variation, it is clear that non-bright band VPRs are not well described by a constant reflectivity gradient. The absolute random error (RMSE, appendix B) for option *b* is also the highest of all the test profiles.

Apart from option *b*, correcting for VPR using all the other test profiles generates underestimated QPEs throughout, but to differing degrees. Profile options *c* and *d* both outperform the control ($Z(h) = Z_s$) beyond 70 km in terms of bias, and perform equally well at shorter ranges, where the radar beam is below the melting layer. This is consistent with observations (throughout this thesis and in the literature) of decreasing reflectivity gradients above the melting layer for almost all types of VPR. The performance of option *d* beyond 130 km is particularly notable, where it reduces underestimation by more than 50%. This profile - which has a constant reflectivity below the melting layer, stratiform reflectivity gradient above, and an offset of 6.5 dB between the rain and ice layers - also performs best in terms of RMSE and RMSF (appendix B). RMSE is considered the more significant factor here, since RMSF as a ratio is weighted towards light rainfall events, while RMSE is more influenced by the larger absolute errors seen at higher rainfall intensities.

In choosing the best candidate profile, the reduction in bias achieved by option *d* at long range must be weighed against the smaller bias at intermediate range of the constant gradient profile (option *b*). The overall performance of these profiles is quantitatively similar in terms of bias, as well as RMSF, and within 50 km of the radar both profiles perform equally well in estimating surface QPEs. At intermediate ranges, profile *b* appears better in terms of bias, but performs worst of all the tested profiles in RMSE. However, at longer range, the profiles of bias, RMSE and RMSF from option *b* suggest that the upper part of this profile is not representative of the cases of most interest. In particular, the pattern of higher absolute and lower proportional errors suggests that this profile is less suitable for higher rain rate events, which are a priority for accurate QPE. Given the likely weighting of the profile *b* errors towards these hydrologically significant events, and the otherwise comparable performance at all ranges from the radar, option *d* was chosen as the most promising candidate for continued testing.

6.3.2 Quality of fit: an illustrative comparison

As an initial illustration of the suitability of the proposed non-bright band profile (figure 6.3, option 3), this shape was fitted to the true VPR from figures 4.1 and 6.2 using the iterative convergence method of Kitchen et al. (1994). Figure 6.6 shows the much improved fit achieved using the new profile shape. Surface rain rate estimates are still slightly underestimated, but by only 10% at 110 km in this case (compared to 57% and 33% respectively in figures 4.1 and 6.2), and by 20% at 150 km (66 / 69% for the Radarnet profiles, table 6.1). This single case illustrates the potential for significant improvements on the current operational treatment of non-bright band profiles.

6.3.3 Further simulation statistics

Following the initial simulations, the performance of the chosen trial profile (figure 6.3d) was investigated with respect to a wider variety of controls. While the initial simulations assume a correct diagnosis of non-bright band precipitation, it is acknowledged that in reality most non-bright band profiles will currently be corrected in Radarnet using the stratiform bright band profile. This Kitchen et al. (1994) bright band VPR was therefore added to the simulation options.

As an additional test, correction using the trial VPR was compared with the simplistic assumption of a fixed surface rain rate, using no local observations or profile information. The *a priori* rain rate was set to 3 mm h^{-1} , which represents approximately the peak of the surface reflectivity distribution in the non-bright band VPR sample (figure 4.7). Anywhere this constant rain rate statistically outperforms rain rates from the idealised profile shapes, it can be concluded that these profiles have no skill in determining and correcting for VPR. This is not strictly equivalent to saying the radar measurement itself has no skill; simply that knowledge of the reflectivity profile at these heights is not yet well enough developed to extrapolate reliable surface reflectivities from measurements at these ranges.

Figure 6.7 shows the performance of the trial VPR shape against the Radarnet stratiform and convective profiles, and compared to the fixed *a priori* rain rate. As with the comparison of the four profile options, there is little difference in performance at short range ($< 50 \text{ km}$), where the simulated radar beam is in the rain layer. Beyond this, the trial profile matches or outperforms both Radarnet profiles by all statistical measures (bias, RMSE and RMSF) at all ranges from the radar. This provides clear evidence in favour of the new trial profile.

It is interesting to note that whilst all three profiles perform better than a fixed rain rate in terms of RMSE, “ 3 mm h^{-1} everywhere” produces a lower bias and RMSF than either

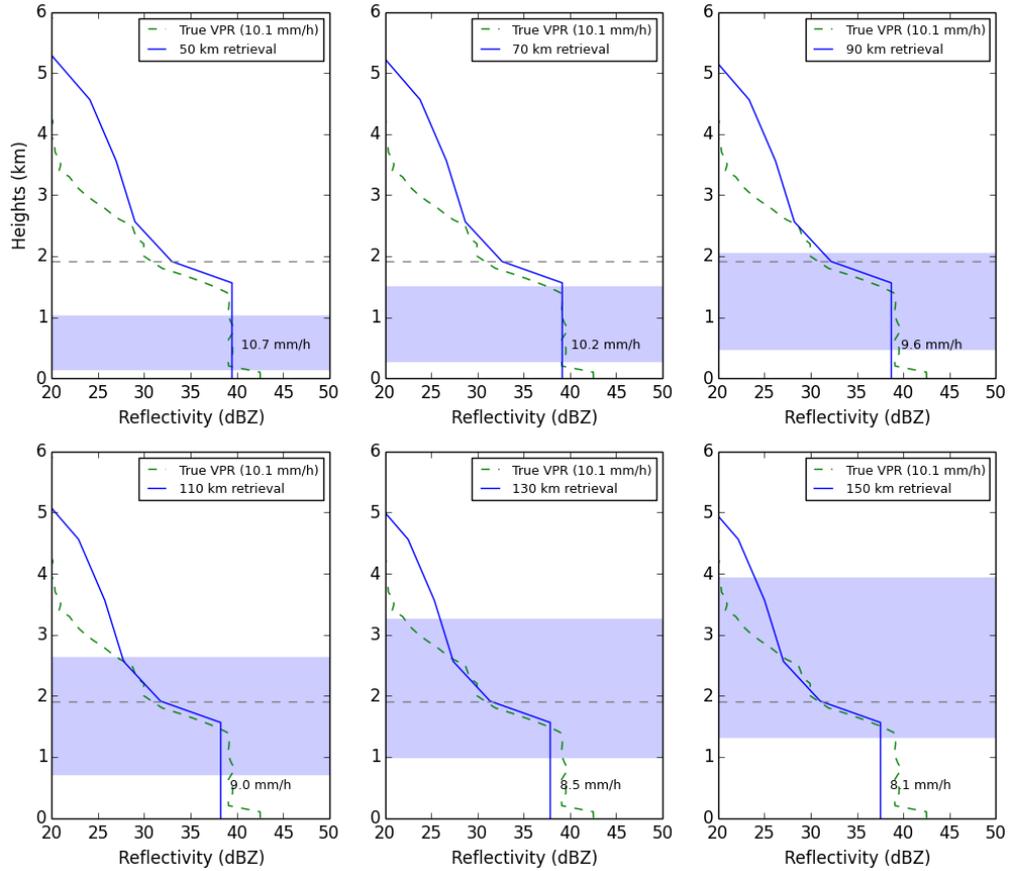


Figure 6.6: The trial non-bright band VPR shape (figure 6.3d) fitted to the non-bright band atmospheric VPR shown in figures 4.1 and 6.2. Underestimation in and above the melting layer region is greatly reduced, with a maximum underestimation at 150 km of 20%, compared to almost 70% underestimation at this range for the two Radarnet profiles.

Range (km)	Stratiform VPR		Convective VPR		Trial VPR	
	Err. (mm)	Err. (%)	Err. (mm)	Err. (%)	Err. (mm)	Err. (%)
50	+0.5	+5	+0.5	+5	+0.5	+5
70	-2.5	-25	+0.1	-1	+0.1	-1
90	-4.8	-48	-1.3	-13	-0.5	-5
110	-5.8	-57	-3.4	-34	-1.1	-11
130	-6.2	-61	-5.5	-54	-1.6	-16
150	-6.7	-66	-7.0	-69	-2.0	-20

Table 6.1: Error in surface QPE resulting from the use of the Radarnet stratiform and convective profiles and the trial non-bright band shape for VPR determination in a single non-bright band profile case, as shown in figures 4.1, 6.2 and 6.6. The “true” surface rain rate from this profile (obtained from the reflectivity above clutter level using $Z = 200R^{1.6}$) is 10.1 mm h^{-1} .

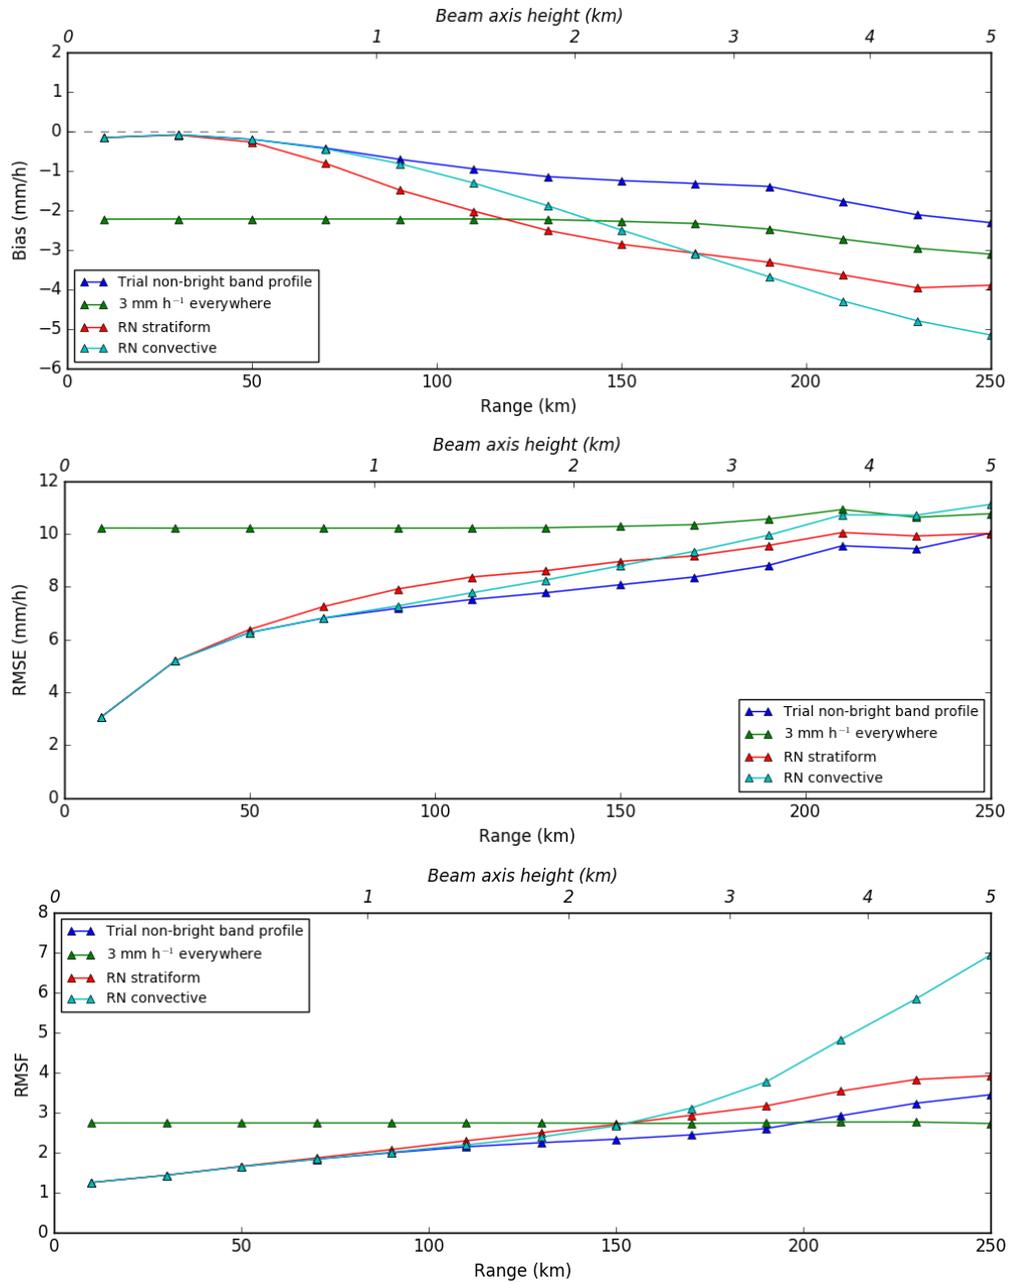


Figure 6.7: Bias, RMSE and RMSF with range of surface rain rates calculated using test profile option *d* (from figure 6.3) against more stringent control scenarios. “RN convective” is the control non-bright band profile (option *a*), “RN stratiform” is the Kitchen et al. (1994) bright band profile, and the 3 mm h⁻¹ is a hard test of “skill against chance” using a representative fixed rain rate for the entire profile sample. Note the colour scale here is **not** the same as for figure 6.5.

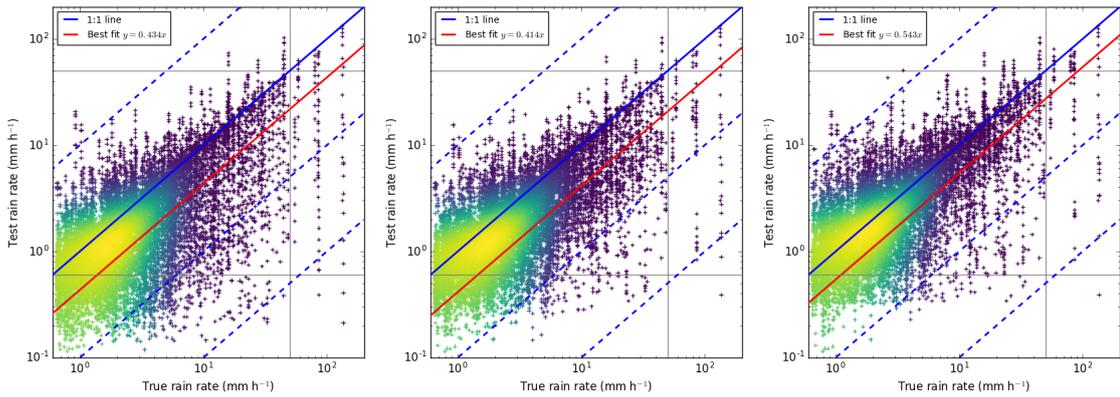


Figure 6.8: Density scatterplots of VPR-corrected vs “true” surface rain rates for the Radarnet convective profile (left), Radarnet stratiform profile (middle) and trial non-bright band VPR (right), including data from all ranges as used in figure 6.7. The plot is colour-coded according to point density, with lighter-coloured points indicating more densely populated regions. Dashed blue lines show the 1:10, 10:1 and 100:1 x:y lines, corresponding to extreme overestimation (10x) and underestimation (10x and 100x) of surface QPEs. Grid lines are plotted in grey at 0.6 and 50 mm h⁻¹.

of the Radarnet profile at ranges beyond 150 km. A fixed “climatological” rain rate also outperforms the trial profile in RMSF beyond 200 km range. The minimal improvements over an *a priori* value beyond 200 km support the assertion of Smyth and Illingworth (1998) that there is inherently minimal skill in VPR determination and correction for measurements above a certain height. However, it is clear from these plots that the trial profile delivers benefits above and beyond the current Radarnet profiles, and has more skill than “climatology” over most, if not all, of the radar domain.

6.3.4 Rain rate scatterplots

Given the high visibility of intense rainfall estimates, and particularly their impact in flood forecasting, it is important for an operational scheme to check not just the statistics but the overall distribution of outcomes. The high rainfall rates that can cause flooding events are uncommon, but their individual uncertainties can be significant, both quantitatively and in terms of visible impacts. These details can be obscured when calculating statistics over a range of intensity bands. For the simulations described in section 6.3.3, therefore, in addition to statistics, the performance of each VPR scheme was evaluated in terms of individual rain rate errors.

Figure 6.8 shows log-log plots of VPR-corrected vs “true” rain rates for the two Radarnet profiles and the trial scheme. Each scatterplot includes rain rate estimates at all ranges, from all 1080 profiles in the sample. Each plot is also annotated with best fit x-y gradient value, which represents the average degree of underestimation in surface rain rates

Profile	Bias (mm h ⁻¹)	RMSE (mm h ⁻¹)
Radarnet convective	-60.9	71.6
Radarnet stratiform	-61.8	71.7
Non-bright band	-51.3	65.8

Table 6.2: Bias and RMSE of rain rate estimates from the three VPR shapes where the “true” rain rate was greater than 50 mm h⁻¹ (mean value 88.4 mm h⁻¹).

obtained from a particular VPR. Both scatterplots generated from the current Radarnet profiles have best fit gradients of approximately 0.4, corresponding to underestimation by factors of 2.3 (convective) and 2.4 (stratiform) respectively. The trial non-bright band profile has a best fit gradient of 0.54. This is consistent with the overall underestimation biases shown in figure 6.7, and indicates that the trial profile reduces underestimation by more than 20% compared to the Radarnet profiles.

In terms of individual QPEs, there is a clear and visible shift of points towards the 1:1 line in the trial profile scatterplot, particularly with respect to the Radarnet convective profile. The highest point density for the trial profile distribution lies on the 1:1 line, rather than slightly below for the Radarnet profiles. In particular, the trial profile has the fewest points in the region between 10 and 100 times underestimation, which is heavily populated for the Radarnet graupel profile. Improvements over the stratiform profile are clearly visible at “true” rain rates exceeding 50 mm h⁻¹, where the trial VPR reduces underestimation bias by 17% and RMSE by 8% (table 6.2).

Although much of this chapter has focused on underestimation, overestimation of high impact rainfall events can be equally damaging. False alarms of flood producing rainfall that lead to unnecessary mitigating action can cause significant costs to be incurred. Problems of extreme overestimation, particularly those involving rain rates in excess of 100 mm h⁻¹, are therefore considered extremely undesirable.

Overestimation at high rainfall intensities (figure 6.8) is slightly increased by the trial VPR. There are a number of cases in which already atypical rain rates (rates exceeding 10 mm h⁻¹ are uncommon in the UK (Thompson, 2007)) are overestimated, sometimes by as much as a factor of 4. This is an inevitable side effect, despite the reduction in random uncertainty, of the shift towards a more centred error distribution. However, extreme overestimation (by an order of magnitude or more) is still extremely rare, occurring at only a handful of points, and is much less common than extreme underestimation even using the trial VPR. It can be concluded that the slight increase in overestimation cases is not sufficient to outweigh the benefits of more accurate overall QPEs.

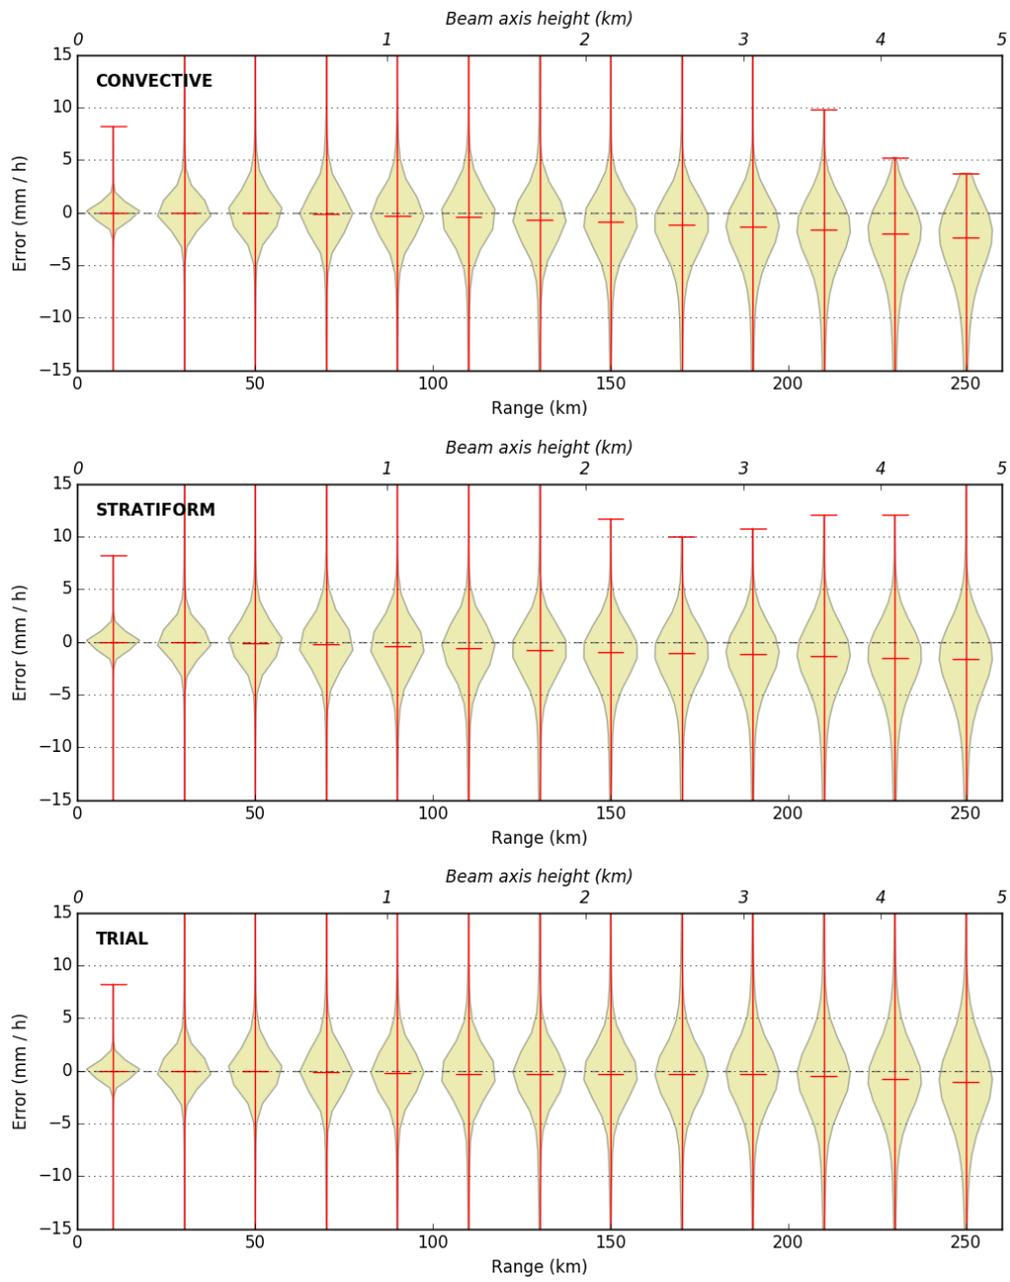


Figure 6.9: Distribution of rain rate bias with range for the 1080 non-bright band profiles when corrected for VPR using the Radarnet convective (top) and stratiform (middle) profile shapes, and the new trial non-bright band profile (bottom).

6.3.5 Distributions of error with range

To complement individual rain rate figures, it is useful to present the distributions of QPE uncertainty with range for each of the trial VPRs. Figure 6.9 shows the distribution of rain rate biases obtained at different measurement ranges using the two Radarnet VPR shapes and the trial profile. This is the same data as in the first panel of figure 6.7; but where figure 6.7 shows the mean bias, figure 6.9 shows the full distribution of rain rate errors at each range, with the median (rather than the mean) as the central value. Most of the distributions in figure 6.9 have significant outliers, shown by the “zero-width” tails of the violin plots that extend beyond the y-axis range. The vast majority of points, however, fit within the “body” of the violins, which illustrate clearly the general properties of the derived error distributions.

After VPR correction, all three profile shapes produce unbiased rain rate error distributions within 100 km of the (simulated) radar. The spread of the distributions increases with range, reflecting the random uncertainty inherent in VPR estimation. Rain rates obtained from the Radarnet convective profile show a clearly increasing negative bias beyond 100 km range. The Radarnet stratiform profile also gives negatively biased QPEs, although succeeds in representing some higher parts of the error distribution at ranges beyond 200 km. The trial non-bright band profile gives distributions beyond 100 km which are more symmetrical and centred closer to zero than either of the Radarnet profiles. Although there are more positive outliers (as in figure 6.8), the overall behaviour of the trial VPR shows less bias and more consistency with range. This provides further evidence that the trial VPR shape better represents the structure of non-bright band precipitation than either of the current Radarnet profiles.

6.4 Case study evaluation with LDR-based classification

Of the candidate profiles constructed in section 6.2, a VPR with constant reflectivity in the rain layer, a drop of 6.5 dB across the melting layer and a negative reflectivity gradient in the ice (figure 6.3d) has been shown to achieve significant improvements in VPR correction over the current Radarnet options. Having demonstrated clear benefits within a simulation framework, this section assesses the ability of the trial profile to generate improved surface QPEs from PPI data. The trial profile, implemented within the Radarnet development framework, is evaluated with respect two controls: one using LDR to identify non-bright band profiles, and the other using the current operational criterion ($Z_1 > 30$ dBZ, see chapter 4 for details).

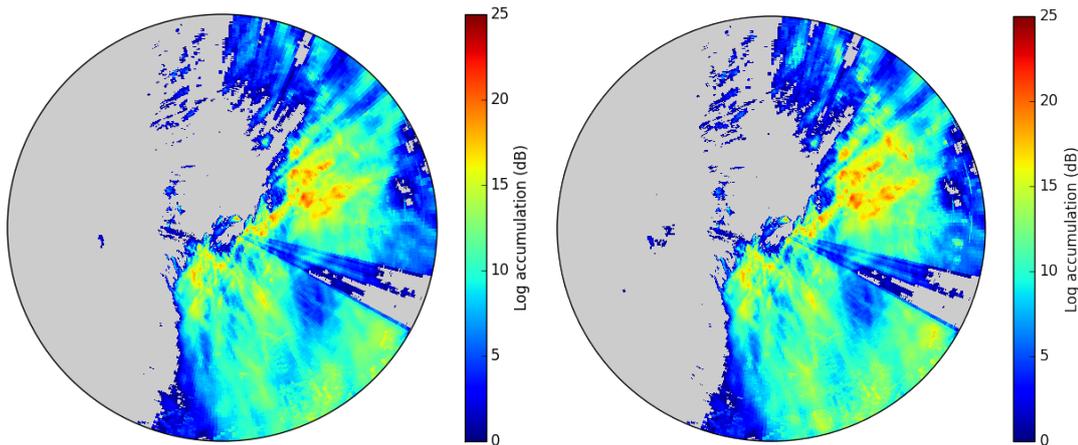


Figure 6.10: Total event accumulations in dB units ($R_{dB} = 10\log(R_{mm})$) for control (left) and trial (right) non-bright band profile shapes, both using LDR to diagnose non-bright band profiles. The maximum range is 255 km. There is visible evidence of significantly higher accumulations at long range using the trial non-bright band VPR, particularly to the North-East of the radar. The arc-shaped artefact at long range to the East of the radar is due to residual non-meteorological echoes, which appeared in a limited number of scans between 01:00 and 04:00 UTC.

6.4.1 Description of case

To assess and compare the benefits of LDR-based classification and the trial non-bright band profile shape, the operational case study from chapter 5 was reused for this chapter. The 24 hour event (from 16:00 on 15th September 2016) was convective and slow moving, generating local accumulations that exceeded 100 mm in some regions. Total event radar accumulations, shown in figure 6.10, give a sense of the magnitude of the event; and the position of major rainfall with respect to Wardon Hill, coupled with a freezing level of around 3 km, provided ample sampling of the melting layer. The radar accumulations from this event, reprocessed using the three VPR methods described, are evaluated quantitatively with respect to “ground truth” provided by colocated tipping bucket rain gauges (as in section 5.5).

6.4.2 Results and discussion

Figures 6.10 and 6.11 show the spatial distribution of radar event accumulations and the differences observed between each trial. As discussed in chapter 5, the use of LDR to classify VPRs results in increased accumulations throughout most of the melting layer region, causing changes of up to 30 mm in 24 hours (figure 6.11, left hand panel). This increase in rain rates is achieved solely through avoiding inappropriate corrections for bright band.

In chapter 5 it was observed that the use of LDR can decrease rainfall accumulations

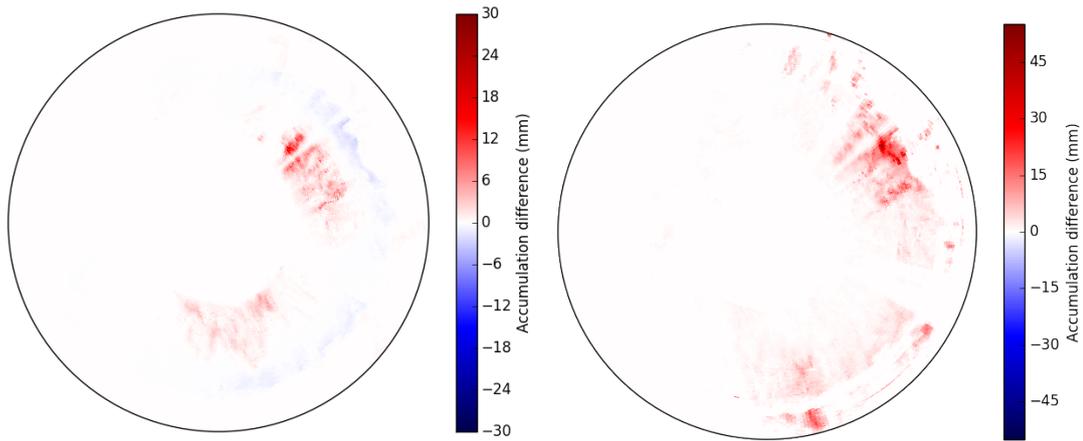


Figure 6.11: Change in accumulation achieved by using LDR to diagnose non-bright band VPRs with the current operational profile shape (left), and then by adding the trial VPR shape to the LDR-based diagnosis scheme (right). The left hand panel reiterates the result of the case study in chapter 5, for comparison with the changes produced by a new VPR shape. Note the difference in colour scale between the two panels, with the new VPR shape generating larger increases in total accumulation.

at long ranges. This occurs where LDR diagnoses non-bright band conditions, but the Radarnet convective VPR does not account for decreasing reflectivity above the melting layer. This effect motivated the investigations of this chapter to determine a more appropriate shape for correcting non-bright band VPRs. Incorporating a trial non-bright band shape into the Radarnet processing (figure 6.11, right hand panel) significantly increases the accumulations in the previously underestimated regions near the top of the melting layer, by properly accounting for the decreasing reflectivity gradient. The pattern of increased overall accumulations can also be seen by comparing the left and right hand panels of figure 6.10, in particular the regions of highest accumulation, which are extended to greater ranges by the trial non-bright band VPR.

To evaluate whether these changes in rainfall accumulation are positive, the radar data are compared with colocated rain gauge totals. Figure 6.12 shows scatterplots of data points where the radar accumulation differs between any of the three reprocessing trials. The left hand panel shows total event accumulations from the radar at rain gauges where at least 22 of the 24 hourly accumulation values were available. (Given that automated gauges often do not report zero accumulations, it was not possible to find gauges for which the complete 24 hour timeseries was available.) The use of LDR alone in classifying VPRs produces a small increase in radar accumulation at most of these points, with another, typically larger increase achieved when the trial VPR shape is added. These changes move the trial accumulations almost universally towards the 1:1 line. The clearest improvements in radar QPEs occur at total gauge accumulations of between 20 and 40 mm over the 24 hour period. The final gauge-radar correlation

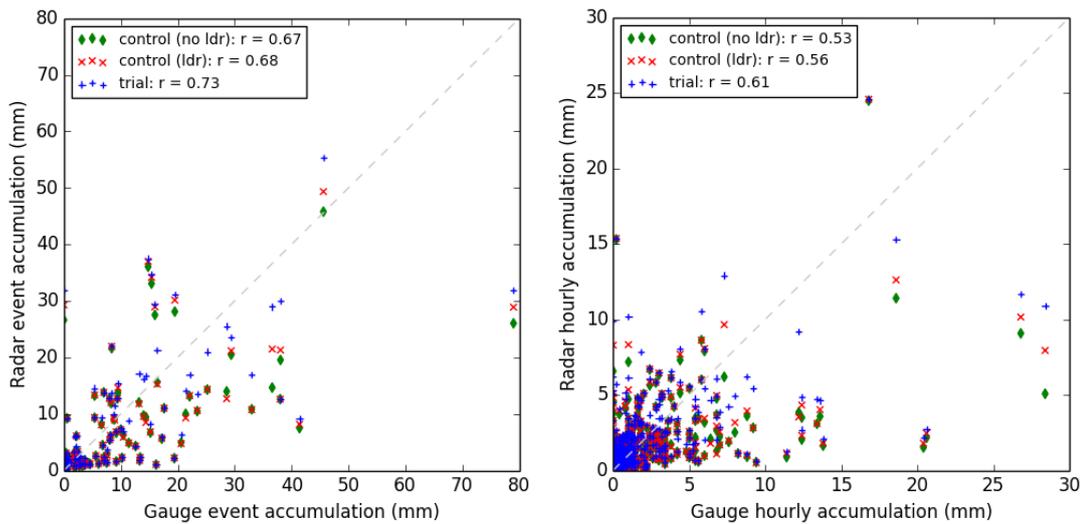


Figure 6.12: Left: total event accumulations at each rain gauge for which a complete timeseries (more than 22 of 24 hours) was available, for each of the two control runs (with and without the LDR scheme) and the new trial profile. Right: hourly accumulations for each run at the points where radar accumulations differed between schemes.

coefficient of 0.73 represents an improvement of 9% over the operational scheme, and of 7% over LDR-based classification alone.

The total event accumulations show a few points at which the radar overestimates. Some of these cases are worsened by the developments, through both LDR-based classification and the new VPR shape. However, for the majority of these points the changes are small, amounting to only a few millimetres over the 24 hour period (eg the cluster of overestimated points at gauge event accumulations of around 20 mm). There is only one point (at an overall gauge accumulation of 45 mm) for which the changes in treatment of non-bright band VPRs are noticeably detrimental to the total radar accumulation.

While total event accumulations are a useful way to characterise overall impacts, the effects of changes to instantaneous rain rates can be obscured by averaging. The right hand panel of figure 6.12 compares radar and rain gauge accumulations from each of the three trials on an hourly timescale. Hourly accumulations from all points where the radar outputs differ are plotted against the corresponding gauge accumulations. At this finer temporal resolution, it is more clearly apparent that the benefits vary with rain rate. The changes to the treatment of non-bright band VPRs yield improvements to all hourly radar accumulations where the true (gauge) accumulation is greater than 10 mm. The performance at lower rainfall intensities is more random, in terms of whether the systematic increase in accumulation improves or degrades the radar QPE. However, the overall gauge-radar agreement is clearly improved by the developments in both classification and determination of the VPR. Over the points where the radar accumulation is

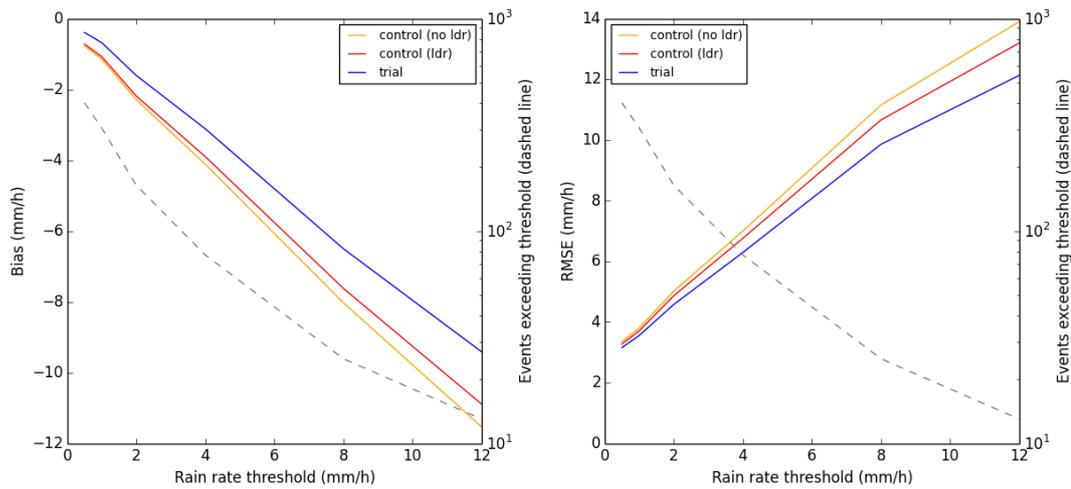


Figure 6.13: Bias (left) and RMSE (right) of radar hourly accumulations for gauge accumulations exceeding the specified threshold. Improvement is evident with the introduction of LDR, with further improvement gained on addition of the trial non-bright band profile.

changed, the gauge-radar correlation coefficient is increased by 15% over the operational algorithm. The larger part of this improvement is achieved by use of a suitable idealised shape to correct for non-bright band VPRs.

The overall error statistics of this event by hourly rain accumulation are shown in figure 6.13. As illustrated throughout this section, the introduction of both LDR-based classification and subsequently the trial VPR shape result in increasing radar accumulations and decreasing underestimation (negative bias) at each stage. RMSE is also decreased by the use of LDR for profile classification, and further again by the trial non-bright band VPR. These improvements tend to be greater for higher rainfall intensities, with a notable 2 mm h^{-1} reduction in RMSE for hourly radar accumulations above 12 mm h^{-1} . This is consistent with the scatterplots in figure 6.12. It is clear from these results that the introduction of a suitable vertical profile shape for correction of non-bright band rainfall can achieve significant benefits for QPE.

6.5 Conclusions

Chapter 5 demonstrated the benefits to QPEs achievable through the application of LDR-based VPR classification to radar PPIs, through avoiding inappropriate corrections for bright band on a local scale. However, there is little information available in the literature to inform actual corrections for non-bright band VPRs. Given the mixed outcomes for QPE resulting from LDR-based classification at longer range, improving VPR determination in non-bright band conditions appears to be a necessary step towards

fully exploiting the potential of real time LDR measurements.

This chapter developed a simple shape for the correction of non-bright band VPRs (figure 6.3d) based on inferences from the existing literature, with additional details informed by the high resolution Wardon Hill dataset. This profile was applied to VPR determination in non-bright band cases using the iterative convergence scheme developed by Kitchen et al. (1994). The behaviour of the new profile shape was initially evaluated using simulation studies. Simulations of bias, RMSE and RMSF with range showed improvements in surface QPEs over both the Radarnet stratiform and convective VPRs. The benefits were particularly marked for intense rainfall, with reductions of 17% in underestimation bias and 8% in RMSE for rain rates greater than 50 mm h^{-1} (table 6.2). These data suggested that the new VPR shape could significantly improve surface reflectivity and rainfall estimates in non-bright band conditions.

The potential benefits of the trial VPR shape for real time QPEs were illustrated using a high impact case study. The case study from chapter 5 was extended to use the new VPR shape wherever melting layer LDR diagnosed non-bright band precipitation, in place of the Radarnet convective profile. The resulting QPEs were assessed with respect to colocated rain gauge accumulations. Comparing results from first the operational convective diagnosis criterion ($Z_1 > 30 \text{ dBZ}$), with the impact of LDR-based classification, and finally the introduction of an improved non-bright band VPR shape, showed incremental improvements throughout as a result of each development. As for the simulation study, the benefits observed in this case were particularly targetted towards the higher rainfall intensities which are of interest in an operational environment: for example the 17% reduction in bias for rain rates greater than 12 mm h^{-1} . It can be concluded that non-stratiform QPEs, when skilfully identified (in this case by LDR), can be measurably improved through the introduction of a suitable idealised VPR.

Although clear benefits have been achieved here, particularly for high intensity QPEs, the investigations of this chapter provide only a preliminary indication of the best treatment of non-bright band VPRs. In section 6.2, several assumptions were made and parallels drawn in order not to define the new VPR shape more precisely than the available evidence could support. For this reason there are a number of aspects of non-bright band VPRs that might benefit from further investigation and tuning.

The main candidate for investigation should be the rain-to-ice reflectivity difference across the melting layer. The magnitude of this offset translates into effectively a bulk bias correction of all reflectivities above the bright band, whether partially or fully sampled by the radar beam; getting this right is therefore extremely important for accurate QPE. The data on non-bright band precipitation from the Wardon Hill sample is relatively limited, both in terms of the sample size and the number of parameters available (only reflectivity and LDR). For this reason only a climatological average value for the rain-to-

ice offset could be robustly derived. The drop in reflectivity across the radar melting layer is likely to be related to microphysical parameters of the precipitation, such as the precise drop size distribution. It is therefore very possible that further observational research, including parameters such as the differential reflectivity, might be able to find a functional form for the rain-to-ice reflectivity offset, more closely correlated with individual meteorological situations than a climatological average value.

Another area that may benefit from future research is the gradient of non-bright band reflectivity profiles above the melting layer. This was observed in chapter 4 to be similar, but not identical to that of stratiform VPRs. Microphysical modelling approaches such as adopted by Kirstetter et al. (2013) could provide a method for exploring suitable profile shapes above the melting layer. However this is likely to have much less impact on QPEs than the magnitude of the rain-to-ice offset, given the smaller proportion of rainfall estimates obtained from measurements at this height in the atmosphere.

Evaluation of more detailed non-bright band profile shapes would benefit from the availability of LDR at 5 minute frequency from a larger number of radars. The need for 5 minute LDR scans as a minimum for VPR classification is discussed in chapter 5, section 5.4. The accurate identification non-bright band cases is crucial to realising the benefits of correcting for this profile shape. However, for this study 5 minute LDR data were only available for a single radar. Evaluation using a larger dataset would be beneficial to explore any regional variations in behaviour, and essential to investigate the impact on multi-radar composites. The benefits shown in this chapter using only simple approximations provide the motivation for future research in the area of non-bright band VPR corrections.

Chapter 7

Improving operational correction for bright band VPRs

7.1 Introduction

The Radarnet operational vertical reflectivity profile correction (VPR) is designed to remove the effects of bright band and compensate for the increasing underestimation at long range seen in stratiform QPEs. The scheme uses an idealised stratiform profile shape (figure 2.3) with a single variable parameter - the reflectivity at the reference level, Z_b , where the reference level is located in the rain immediately below the melting layer (section 2.2). This profile (henceforth “K94”) is fitted iteratively to the lowest usable radar scan reflectivity measurement, using a known beam power profile, to estimate reflectivity and rain rate at the surface. The operational method is described in detail in section 2.5.3.

The K94 VPR method was implemented in Radarnet in the mid-1990’s. In the several years following its implementation some minor refinements were made to the scheme, including the use of satellite (Meteosat / MSG) cloud top measurements to estimate the profile top height, and the inclusion of a model-estimated orographic enhancement correction (Kitchen, 1997; Lewis et al., 2007). Since then, no changes have been made to the operational method of correcting for stratiform VPR.

A significant amount of research literature on the VPR has accumulated since the initial Radarnet implementation. Shortly after the original K94 paper, Fabry and Zawadzki (1995) published an observational study using vertically-pointing radar which described five distinct types of VPR. Their observations of stratiform profiles suggested a typical decrease in reflectivity across the melting layer of 1-2 dB. This contrasts with the K94 VPR shape, in which reflectivities at the top and bottom of the bright band are equal.

The finding of Fabry and Zawadzki (1995) is consistent with a number of later studies involving PPI measurements (eg Andrieu and Creutin, 1995b; Kirstetter et al., 2013), and with microphysically-based simulations of stratiform rain (eg Hardaker et al., 1995), which also describe a drop in reflectivity across the bright band. Recent investigations into bright band and VPR corrections (eg Rico-Ramirez et al., 2005; Matrosov et al., 2007) have built on this growing understanding that stratiform profiles should have different reflectivities at the onset and completion of melting. Yet this knowledge has not yet been applied in the UK.

The vertical profile dataset supporting this thesis provides additional VPR data specific to the UK climate, which could be used to refine and improve the shape of the K94 idealised bright band profile. The dataset, collected using high resolution RHI scans from the Met Office Wardon Hill research radar (chapter 3), contains 5330 bright band VPRs from all UK seasons over a period of almost a year. In most aspects the average bright band profile obtained from this dataset (figure 4.6) matches well with the K94 idealised profile. However, as consistently observed throughout recent radar literature, the average Wardon Hill profile shows an offset in reflectivity between the top and bottom of the melting layer. The median offset in this sample is approximately 2 dB, which matches the observations of Fabry and Zawadzki (1995). These data suggest that the K94 profile shape could be modified to improve the fit with typical stratiform conditions, and by extension to improve VPR correction in the UK Radarnet system.

This chapter investigates the possibility of improving VPR corrections in the UK by adding a 2 dB offset to reflectivities above the bright band in the idealised vertical profile (figure 7.1). The first step is to determine whether there is any residual bias in existing QPEs that could be attributed to the shape of the K94 profile. This question is addressed in section 7.3. Section 7.4 investigates the impact of changing the profile shape on surface reflectivities obtained from the Wardon Hill VPR dataset, using the simulation method developed in chapter 5. The updated VPR shape is tested on real time radar PPIs in section 7.5, and the results discussed in terms of both gauge-radar statistics and range-dependent biases in the final QPEs. Conclusions are presented in section 7.6.

7.2 Proposed change to the Kitchen et al. (1994) stratiform profile

The main feature of the K94 stratiform profile currently used by the Met Office is the symmetrical triangular bright band, which spans a depth of 700 m below the model wet bulb freezing level (Brown et al., 2012). Bright band intensity, in terms of the difference

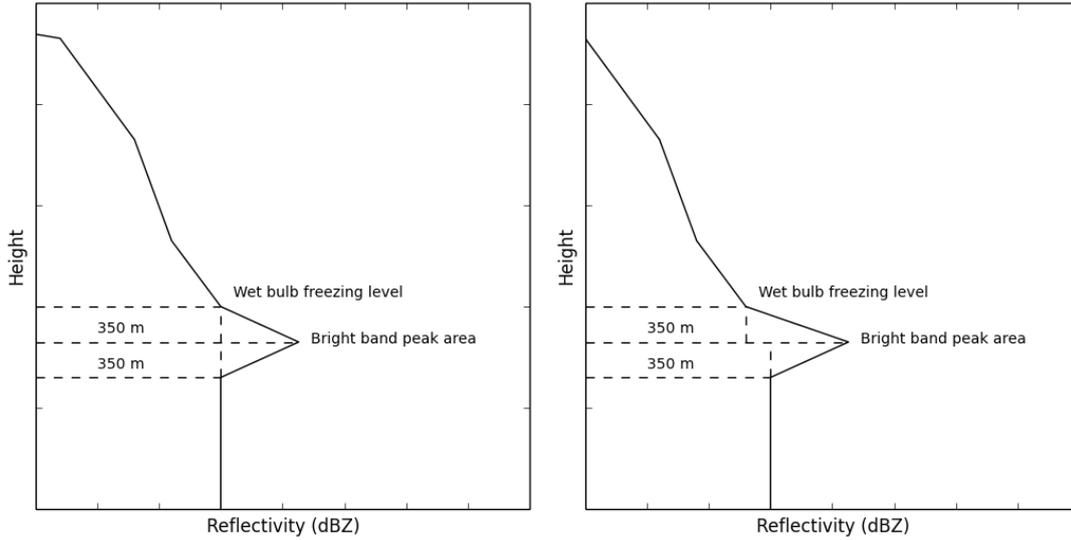


Figure 7.1: Left: the idealised stratiform VPR shape, derived by Kitchen et al. (1994), which is used operationally in the Met Office centralised radar processing system (Radarnet). Right: proposed update to the VPR shape, introducing a 2 dB reflectivity offset across the bright band.

between peak Z_{bb} and rain Z_b reflectivities, is related to the background reflectivity by:

$$A_{bb} = 1.42 \log(Z_b) + 2.1 \quad (7.1)$$

$$\text{where } A_{bb} = 0.5 d_{ml} \times (Z_{bb} - Z_b) \quad (7.2)$$

Quantities A_{bb} and both reflectivities are in linear units ($\text{mm}^6 \text{m}^{-2}$ and $\text{mm}^6 \text{m}^{-3}$ respectively), and the melting layer depth $d_{ml} = 700$ m. The reflectivity below the bright band is constant (except for orographic effects, Kitchen (1997)), and the reflectivity above the bright band decreases exponentially from Z_b according to the cloud depth.

The proposed change to the idealised stratiform VPR for this investigation is illustrated in figure 7.1. This consists of a simple adjustment to reflectivities above the melting layer, to reduce their intensity relative to Z_b by 2 dB. This change does not affect the calculation of bright band intensity, which is based on an empirical best fit relationship between bright band area A_{bb} and Z_b . In deriving equation 7.1, Kitchen et al. (1994) acknowledge that there is a difference between reflectivities at the top and bottom of the melting layer, with the reflectivity below the bright band typically being the larger of the two. A_{bb} is therefore calculated using the reflectivity at the base of the bright band. The change to be investigated here affects only reflectivities above the bright band, and therefore has no impact on the bright band area calculation or the empirical relationship between Z_{bb} and Z_b .

7.3 Evidence of residual bias in stratiform QPEs

The literature and observations discussed in section 7.1 suggest that the K94 idealised VPR shape may overestimate the intensity of reflectivities above the bright band with respect to the rain layer Z_b , by as much as 2 dB. If true, corrections derived using this shape are likely to underestimate the true correction required for reflectivity measurements at long range. The corresponding surface reflectivities, and therefore QPEs, would then be systematically underestimated for measurements where the radar beam is above the melting layer.

Since 2010 the Met Office has maintained an archive of selected operational radar observations and products. This archive contains not only reflectivity PPIs, but processed outputs from the operational Radarnet system. In particular, records of colocated radar and rain gauge accumulations are available for each month since October 2010. These records provide the opportunity to evaluate past operational data for any systematic trends in bias that could be attributed to the VPR.

7.3.1 Selecting cases from the radar archive

The Met Office archive contains radar accumulations derived from single site Cartesian-gridded QPEs, alongside colocated rain gauge accumulations, indexed by range from the radar. This allows biases to be investigated as a function of range. However, for this investigation, we are interested in how QPE errors vary with measurement height, and in particular the region above the melting layer.

The height of the radar beam axis above the ground in normal propagation conditions is determined using the 4/3 Earth approximation (equation 1.15, duplicated here):

$$h(r, \phi, \theta) = h_{rad} + \sqrt{r^2 + R'^2 + 2rR'\sin\theta} \quad (7.3)$$

$$\text{where } R' = \frac{4}{3}R_E \text{ and } R_E = 6374 \text{ km (Earth radius)} \quad (7.4)$$

and ϕ and θ are the azimuth and elevation angles respectively. This height h increases monotonically with range r from the radar for a fixed elevation angle. The height of the beam above the melting layer, however, is:

$$\Delta h(r, \phi, \theta, t) = h_{rad} + \sqrt{r^2 + R'^2 + 2rR'\sin\theta} - h_{fl}(r, \phi, t) \quad (7.5)$$

This relationship varies with the freezing level height h_{fl} , which is not constant, but varies both with spatial location and over time.

Equation 7.5 complicates the analysis of the archived accumulations. Since Δh is not a

monotonically increasing function of range, analysis by range may not represent clearly the effects of VPR due to the presence of other effects that could interfere with any trend. However, by carefully selecting data designed to minimise non-VPR effects, it may be possible to observe underestimation biases with range in the archived radar QPEs.

In choosing an archived study period, it was desirable to try to identify times where Δh should correlate well with r , as well as periods when a high proportion of QPEs would have been collected from heights where $\Delta h \gg 0$. This led to the following data selection considerations:

- The study period should be long enough for statistical validation, but not so long that temperatures (and therefore bright band height) are extremely variable. A single meteorological season (3 months) was selected as a compromise.
- The study period should be from a cool season, but not too cold. This balances the requirement for good sampling above the melting layer with the need for rain at the surface (rain gauge accumulations are unreliable in snow).
- The dominant rain type over the study period should be stratiform.
- A high total rainfall accumulation is desirable, to maximise the amount and variety (in intensity, proxied by hourly rainfall accumulation) of the validation data available.

Using the National Climate Information Centre monthly climate summaries (<http://www.ncic/series/maps.html>), the winter (DJF) of 2013-14 was identified as a suitable case study. This season featured an almost continuous series of Atlantic lows tracking West to East across the UK, resulting in record-breaking 3 month accumulations from low to moderate intensity stratiform rain (Lewis et al., 2015). Typical surface temperatures were also relatively consistent around 4-6°C (appendix D). This season therefore provided a good sample of rainfall estimates obtained from measurements above the wet bulb freezing level.

7.3.2 Choosing suitable radars

To preserve the clearest possible mapping between Δh and r , the choice of radar is as important as the study period. Having controlled as far as possible for variations in h_{fl} , the mapping between range and height above freezing level for a single radar is a function of elevation θ (equation 7.5). Any effect leading to the selection of higher elevation data for QPE therefore introduces noise into the relationship, and confuses the interpretation of trends with range in terms of Δh . Beam blockages in particular have a significant effect on the range-height mapping, as whole rays and sectors of the radar image can be

affected. For a study of VPR through a range proxy, therefore, the radars used should not be affected by any significant beam blockages. This selection criterion disqualified data from over half of the radars in the UK network.

An additional consideration in simplifying VPR effects is orography (see section 2.2.2). The impact of low level orographic enhancement on surface rainfall estimates can be extremely significant (Kitchen, 1997; Lewis et al., 2007), and could easily obscure the signal from simpler range effects. The requirement to avoid significant orography reduced the number of suitable radars to two. The radar at Ingham (appendix A) is affected by neither orography nor beam blockages. The radar at Dean Hill (appendix A) has no orography, but one significant sector blockage to the East. However this blockage extends into the highest elevation scan, so that no QPEs are calculated in this region, and the range-height mapping at the locations of valid QPEs is preserved.

Selecting and blacklisting radars may seem an unnecessarily wasteful approach. An alternative approach - blacklisting rain gauges according to their location with respect to blockages and orography - could potentially retain more valid accumulation data and allow for a more robust inspection for VPR effects. However, the locations at which beam blockage and orographic enhancement occur are not static over seasonal time periods. Orographic enhancement intensity and extent is dependent not only on terrain height but on meteorological conditions; and beam blockages can vary seasonally, for example with the growth and die-back of trees close to the radar. Since a signal for VPR effects in range will already be difficult to retrieve, to minimise the noise on this signal, it was decided to accept a smaller amount of data which more stringently adhered to the necessary quality requirements.

7.3.3 Discussion

To evaluate the nature of residual biases in QPEs with range, hourly radar accumulations from Ingham and Dean Hill were compared with colocated rain gauge accumulations. The radar products were Cartesian, gridded at 5 km spatial resolution, and extended to a maximum range of 255 km from each radar.

Figure 7.2 shows the ratios of radar to gauge hourly accumulations for Ingham and Dean Hill over the period from 1st December 2013 to 28th February 2014. The data are binned by range in 20 km steps, and plotted on a log scale. Before conversion all “zeros” (hourly accumulations of less than 0.01 mm from either radar or gauge) were removed, to ensure as far as possible that differences between radar and gauge accumulations were due to VPR effects, rather than missed detections through beam overshoot or misplaced echoes due to wind drift below the radar beam (Mittermaier et al., 2004; Sandford, 2015).

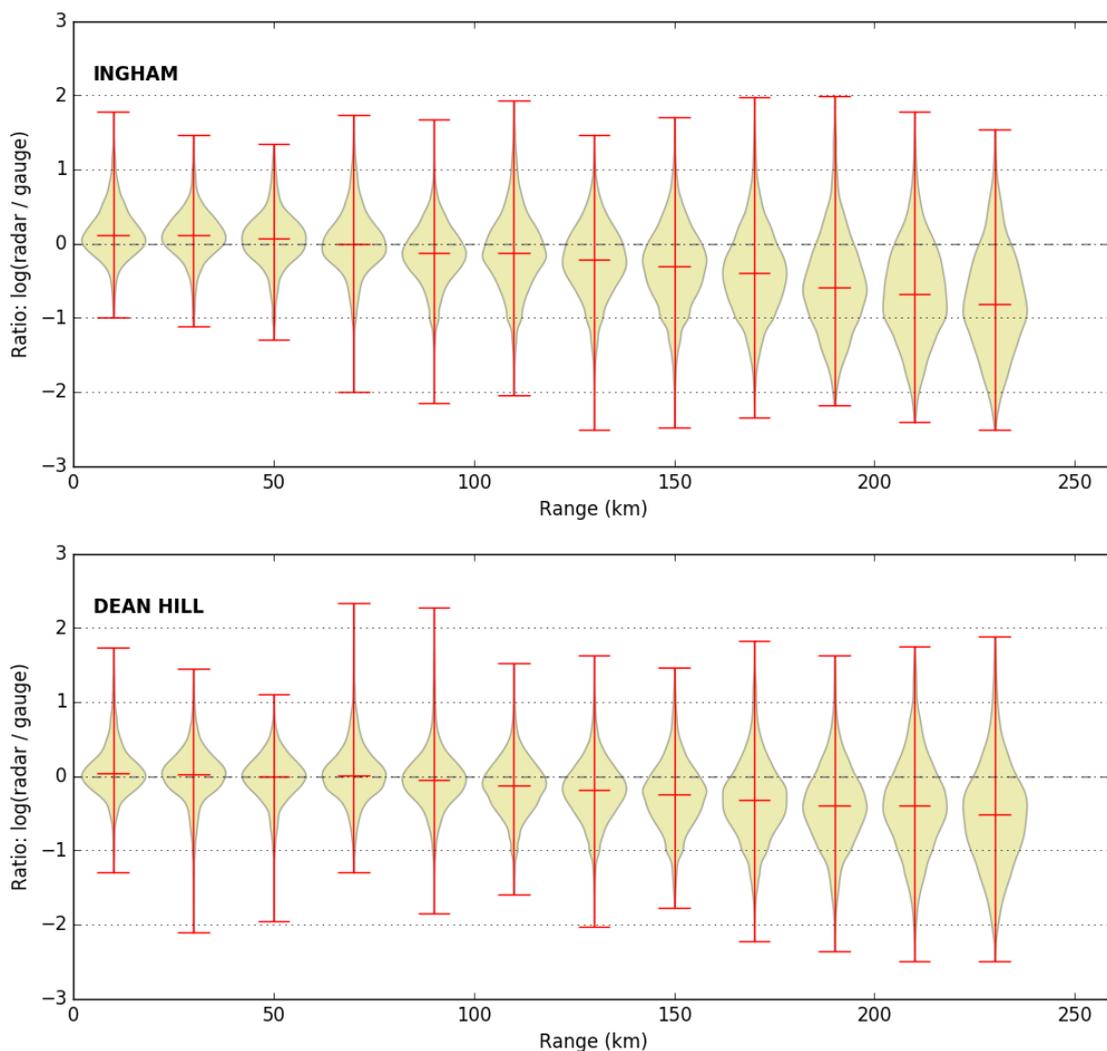


Figure 7.2: Log hourly radar-gauge ratios for archived operational accumulations from Ingham (top) and Dean Hill (bottom) over the winter (DJF) of 2013-14. Red errorbars show the full extent of the data, which consist of 85929 and 62049 non-zero gauge-radar pairs for Ingham and Dean Hill respectively. Hourly accumulation “points” are sorted into 20 km bins by range from the radar.

The plots in figure 7.2 highlight a clear trend towards underestimation with increasing range. This begins at around 100 km from each radar, and becomes a pronounced negative bias by 150 km. Beyond 150 km range, over this season, more than 75% of hourly radar accumulations were underestimated with respect to colocated rain gauges. The median radar-gauge ratio at the longest range bin (at 220-240 km) is 0.30 at Dean Hill and 0.15 at Ingham.

Although affected by variability in the r - Δh relationship over the study period (equation 7.5), this retrospective analysis shows a clear trend towards underestimation in stratiform conditions for long range radar QPEs. This is consistent with the K94 VPR underes-

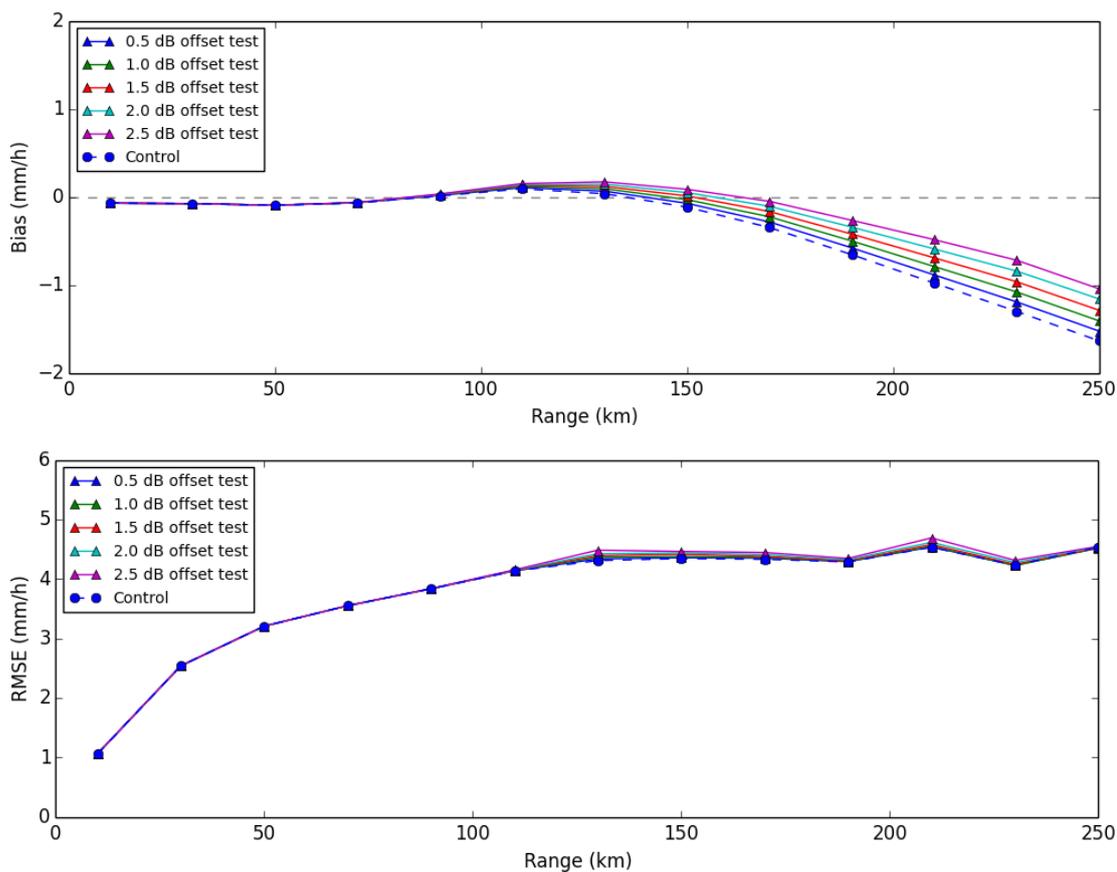


Figure 7.3: Bias (top) and RMSE (bottom) with range of surface rain rate estimates obtained from the Wardon Hill stratiform VPR dataset, by fitting idealised VPR shapes with different bright band offsets. The control is the K94 profile, with equal reflectivities at the top and bottom of the bright band.

timating the correction necessary for reflectivity measurements above the bright band. A finding of underestimation at long range adds weight to the hypothesis discussed in section 7.1, that there could be benefits to adding an offset to the K94 idealised profile, to account for the change in atmospheric reflectivity across the bright band.

7.4 Investigating offset values

The results of section 7.3 suggest that the updates to the stratiform VPR shown in figure 7.1 could be beneficial for QPE. However, the magnitude of the result in terms of the most suitable reflectivity offset across the bright band is not clear. To investigate this further, and to quantify the possible impacts of a bright band offset on QPEs, a simulation study was conducted using the high resolution VPR dataset from Wardon Hill (described in chapter 3).

The impact of a bright band offset on the accuracy of surface reflectivity estimates was tested using the 5330 stratiform “bright band” profiles from the Wardon Hill dataset (section 4.2). The simulation method mirrors closely that of chapter 6, in which reflectivity measurements were simulated from each profile, corrected for VPR, and the outcomes compared with “true” surface reflectivities from the original profile. Measurements were simulated for a radar beam at 0.5° elevation, at ranges from 10 km to 250 km at 20 km intervals. Simulations were run for a range of different dB offsets of similar order to the 1-2 dB observed by Fabry and Zawadzki (1995).

The results of the simulation study are shown in figure 7.3. Accounting for the drop in reflectivity across the bright band has a clear positive impact on long range QPE bias (appendix B), with impacts visible at all ranges beyond 130 km. The effect of different offsets is predictable, with larger offsets achieving greater reductions in bias at the longest ranges. The RMSE (appendix B) however is unaffected by changes in the VPR over this sample, and shows very little sensitivity to the size of the bright band offset.

Given the insensitivity of random error to the bright band offset, the “best” value may be chosen by other means. The choice of 2 dB was informed largely by the literature, which suggests a suitable range of 1-2 dB (Fabry and Zawadzki, 1995; Rico-Ramirez et al., 2005). This is also consistent with the average stratiform profile from the Wardon Hill VPR dataset.

7.5 Impact of a 2 dB offset for real time rainfall estimation

Section 7.4 illustrates the potential for benefits to stratiform QPEs by the addition of a 2 dB reflectivity offset across the melting layer in the idealised VPR. This section evaluates the impacts of a 2 dB melting layer offset for real time QPEs. In section 7.5.1, a subset of the predominantly stratiform archived dataset from Ingham and Dean Hill is reprocessed with the updated profile and compared with the K94 control. The evaluation is extended in section 7.5.2 to a real time trial including more varied weather conditions.

7.5.1 Reducing long range biases in stratiform QPEs

The impact of long range underestimation due to VPR was illustrated using single site data from the radars at Ingham and Dean Hill, over the winter of 2013-14 (figure 7.2). One month of this data was used to evaluate directly the impact of an updated stratiform VPR. The reduction of the trial period from three months to one was due to the computational expense of running the Radarnet system in reprocessing mode for an extended

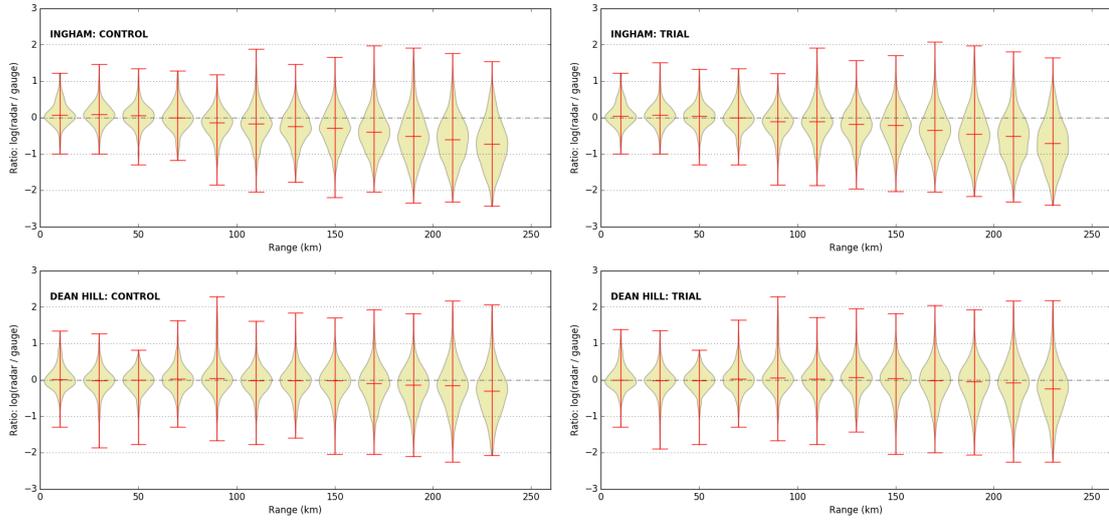


Figure 7.4: Reprocessed hourly radar-gauge ratios with range for Ingham (top) and Dean Hill (bottom) over January 2014. 17992 and 14409 non-zero gauge-radar pairs were available for Ingham and Dean Hill respectively. Left: with operational processing; right: with the new trial VPR.

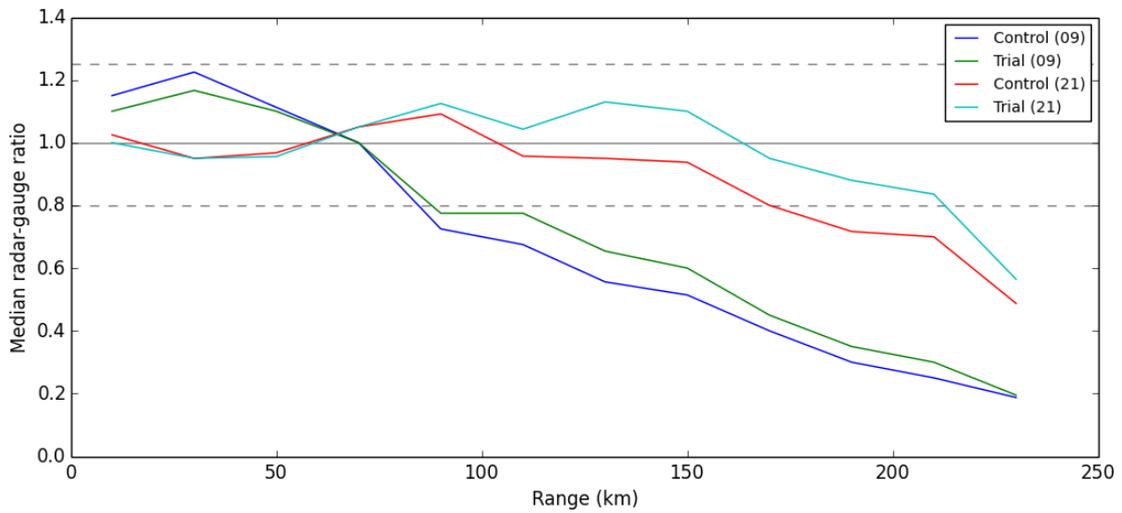


Figure 7.5: Median hourly radar-gauge ratios from figure 7.4, for the control (K94) and trial (2 dB offset) stratiform VPRs applied to January 2014 data. Ingham is radar 09, and Dean Hill is radar 21. The “perfect” value is 1, shown by the grey line, with values less (more) than 1 representing radar underestimation (overestimation) with respect to ground truth. Dashed grey lines delineate the range from 0.8 to 1.25 used by Figueras I Ventura and Tabary (2013) to define “high quality” radar QPEs.

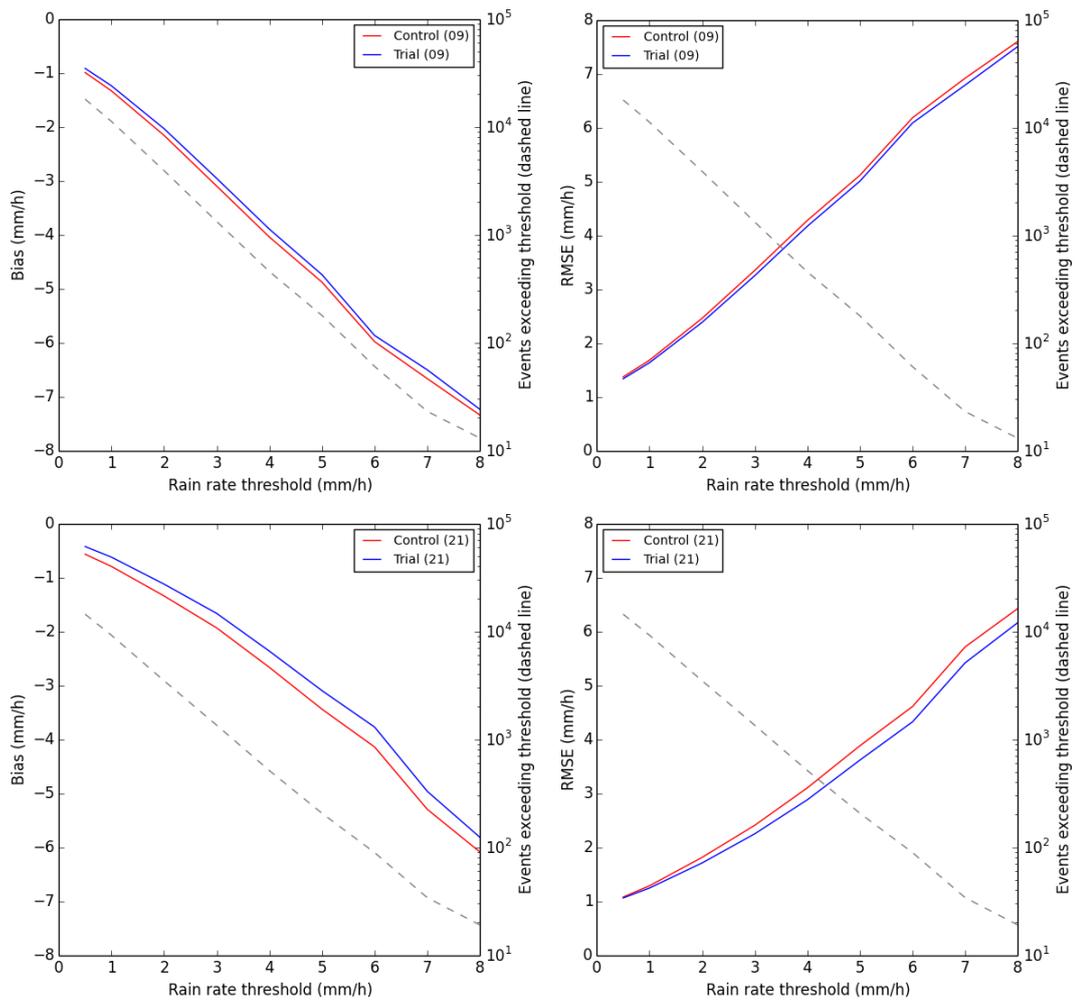


Figure 7.6: Hourly accumulation bias and RMSE with rain rate for the radars at Ingham (09, top) and Dean Hill (21, bottom) over January 2014.

period. The month of January 2014 was selected for this evaluation, as this had the highest number of rainfall events. Reflectivities from Ingham and Dean Hill were reprocessed using both the operational K94 VPR shape, and the updated shape, with a 2 dB offset across the bright band. Hourly single site radar accumulations from both trials were then compared with colocated hourly tipping-bucket rain gauge accumulations.

Figure 7.4 shows the impact of updating the VPR on gauge-radar bias with range. The ratio of radar to rain gauge measurements is plotted in 20 km range intervals up to 250 km (for comparison with archived data plots, figure 7.2). Although the plots are similar, small differences in distribution locations can be seen at ranges beyond 100 km.

To highlight the changes more clearly, figure 7.5 shows the median radar-gauge ratio with range for control and trial data at each radar. This plot also places in context the magnitude of this ratio as an objective quality measure, using the limits defined by

Figueras I Ventura and Tabary (2013) to identify high quality QPEs from the French operational radar network. Figure 7.5 shows a clear reduction in underestimation beyond 80 km range at Ingham when the new profile is used, although the improved ratios still lie outside the “high quality” bounds. The change in range behaviour at Dean Hill shows some overestimation occurring at intermediate ranges (60-150 km); however, the proportion of the radar domain where the median ratio lies between 0.8 and 1.25 is extended by almost 50 km in range. Beyond 150 km, average underestimation is reduced for both radars with respect to the K94 profile.

Figure 7.6 shows the gauge-radar statistics with rain intensity at each radar, for the K94 and test profiles. There is improved agreement at both radars using the new VPR, with a small but visible reduction in both bias and RMSE for the 2 dB offset profile. For quantitative statistics the improvement is greater at Dean Hill, with changes of order 0.2 mm h^{-1} exceeding the quantisation error in tipping bucket rain gauge measurements. This suggests that on average, the reduction in long range underestimation outweighs the overestimation observed at intermediate ranges.

7.5.2 Real time trialling over March 2017

Given the potentially widespread impacts of this change, and the rather small magnitude of the improvements demonstrated, further supporting evidence was required to consider an operational implementation. “Live trialling” of a change to the Radarnet system involves running a version of the code containing the change in parallel with the latest “trunk” version of the shared development system. The Radarnet trunk code runs continuously in real time on a development server, as part of the resilience and integration testing of scientific and technical changes before operational implementation. This version of the code was stable during March 2017, and could therefore be used to generate a set of control results against which to evaluate the performance of the updated VPR.

A live trial of the change to VPR was begun on 16th March 2017. From the beginning of this trial, immediate and significant differences were observed between the old and new VPR schemes. Cold temperatures for the season meant the 0°C isotherm height was low, so that the radar beam was at least partially above the melting layer over most of the composite domain. A 2 dB increase in corrected ice-level reflectivities therefore had a significant and widespread impact on composite rain rates. Given the magnitude of the changes and the outcomes of ongoing analysis of the results, the trial was terminated on 31st March 2017. The results from this two week live trial are discussed below.

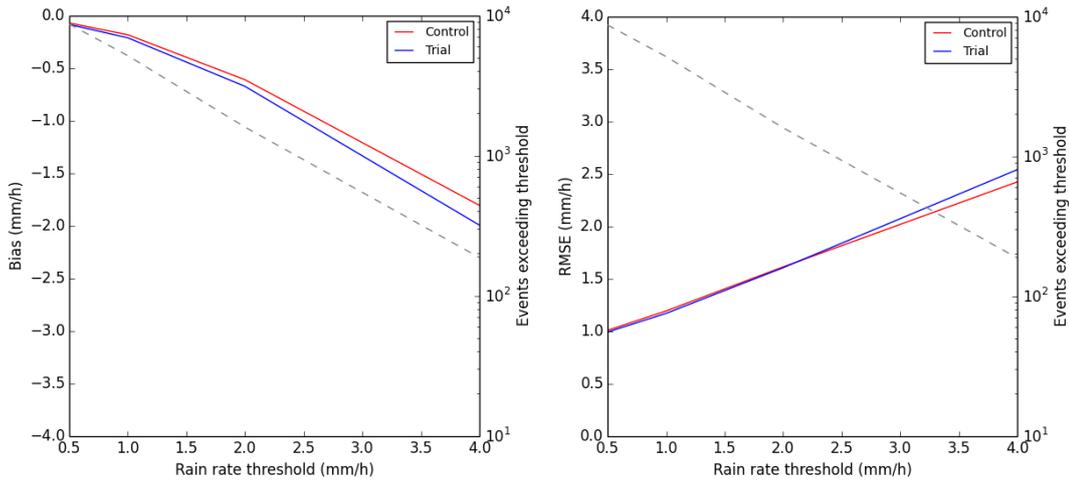


Figure 7.7: Composite hourly accumulation bias and RMSE with rain rate over the live trial period (16th-31st March 2017).

Composite gauge comparisons

Gauge-radar statistics from the 1 km composite for the trial vs the current VPR implementation are shown in figure 7.7. The changes are small, but in the context of a final check before operational implementation, they provide an immediate and unexpected cause for concern. Although the impact of this change to VPR can **only** be to increase radar rainfall estimates, these results are in direct conflict, showing composite QPEs that are increasingly underestimated using the new scheme. The gauge-radar RMSE is increased for rain rates above 2 mm h^{-1} - precisely the reverse of the benefits observed in single site rain rates.

Part of the justification for live trialling is that sometimes unexpected interactions occur between different parts of the Radarnet processing chain. The reduction in rainfall rates observed in the composite cannot be a consequence of the changes to VPR, and therefore must be due to an interaction with some other process within Radarnet. The order of processes applied to generate operational QPEs (section 1.5) leaves only one obvious downstream process that could be causing this effect: the calculation and application of gauge adjustment factors.

Gauge adjustment factors

The use of rain gauge data in radar processing is designed to compensate for residual mean field biases in radar QPEs by calibrating against “ground truth” from gauge accumulations. Residual biases can be the result of miscalibration of the radar reflectivity measurement, or systematic errors caused by the use of an inappropriate reflectivity-rain

rate (ZR) relation (section 1.2.2).

In Radarnet, a mean field bias gauge adjustment factor (GAF) is calculated to minimise the long term gauge-radar root mean squared factor (RMSF, see appendix B). The RMSF represents an average of the mean gauge-radar ratio that is weighted towards lower rain rate values. The GAF is updated every hour using colocated radar and gauge accumulations from within 100 km of the radar location. This range limit aims to minimise the impacts of radar detection failures due to beam overshoot (section 1.5.3), which artificially biases a ratio-based adjustment towards higher values. A Kalman filter evolution with weighting factors based on the total accumulation effectively enforces a varying time window for calculation, so that the GAF updates more rapidly in response to intense precipitation.

The mechanism for gauge adjustment factors to adversely affect composite rain rates is straightforward. The use of an updated stratiform VPR increases QPEs from radar measurements taken above the wet bulb freezing level. This would tend over time to decrease the GAF: the factor by which the (now increased) radar QPE should be adjusted to optimise agreement with gauges. However, the reality of radar errors is that they are not well captured by a mean field bias ratio. In particular, higher rain rates are subject to greater underestimation by radar, due to the dependence of the empirically-derived ZR relation on drop size distribution (section 1.2.2). Reducing the gauge adjustment factor might therefore cause further underestimation of these higher rain rates, negating the benefits of the improved VPR correction.

As a first step in testing this hypothesis, gauge adjustment factors from the trial period were retrieved and are shown in figure 7.8. These plots confirm that the majority of gauge adjustment factors for the updated VPR were significantly reduced with respect to those in the live Radarnet trunk. It is clear, therefore, that the gauge adjustment process acted to reduce the impact of the VPR changes in reducing the bias of long range rainfall estimates.

To verify that mean field gauge adjustment was the only cause of degradation in the live trial statistics, single site gauge adjusted and unadjusted data were analysed separately, to isolate the impact of GAFs on both the individual and aggregated radar data. These results are discussed below.

Single site gauge comparisons

Single site radar accumulations from the live and control trials were analysed using 5 km-gridded Cartesian products, whose range extends 255 km from the radar location. The radars at Shannon, Dublin and the research radar at Wardon Hill were excluded

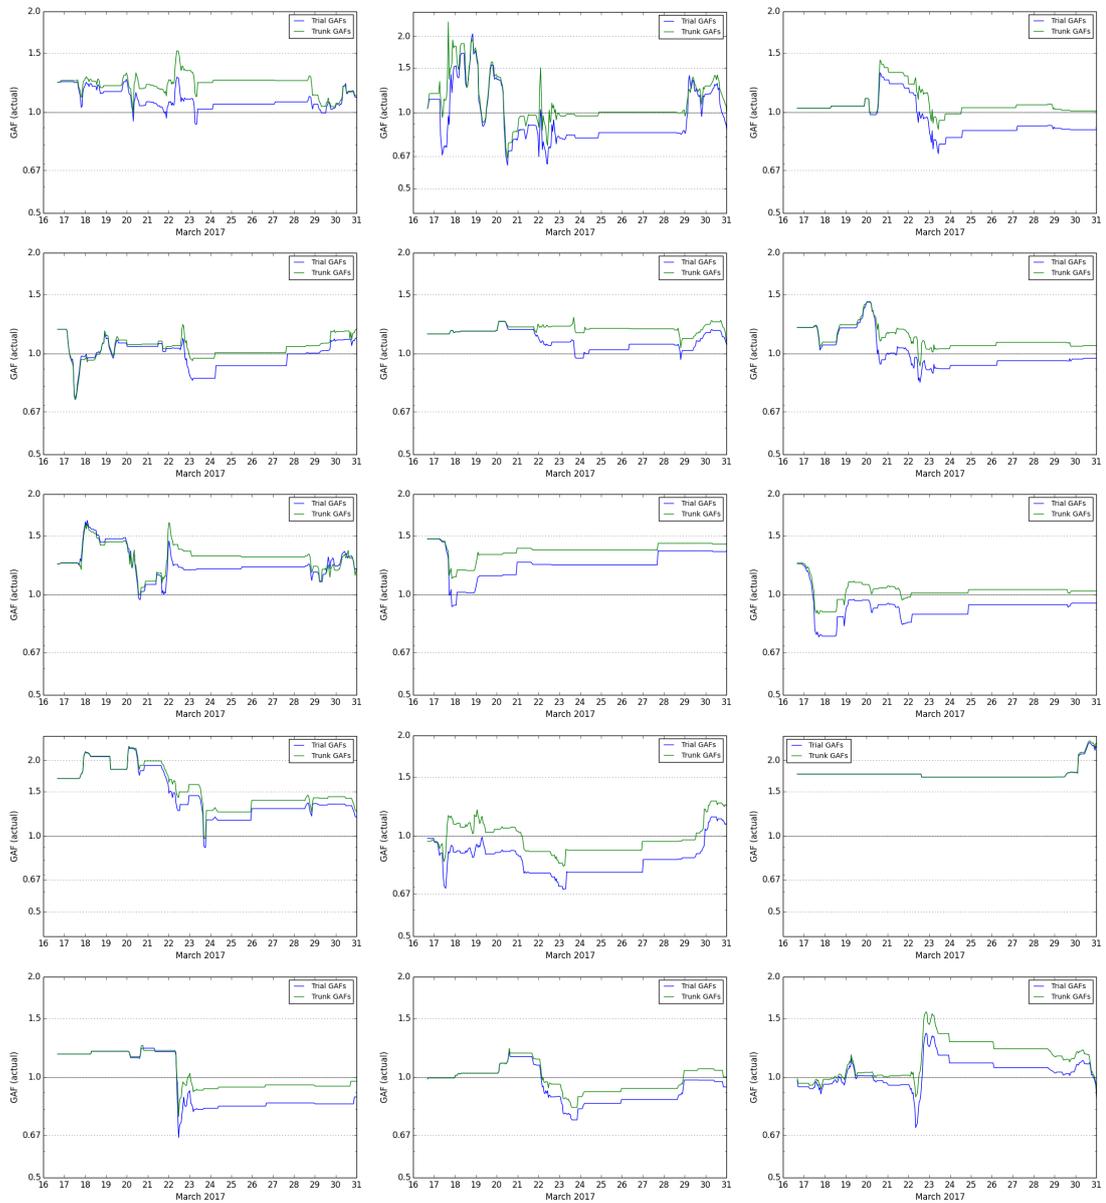


Figure 7.8: Applied gauge adjustment factor timeseries for each radar during the live trial. Green lines show the operational gauge adjustment factors, while blue lines show those derived and applied in the trial setup. From top left: Clee Hill, Hameldon, Chenies (row 1), Castor Bay, Predannack, Ingham (row 2), Crugy, Dudwick, Lewis (row 3), Cobbacombe, Holehead, Munduff (row 4), Thurnham, Dean Hill and High Moorsley (row 5). Details of these radars can be found in appendix A.

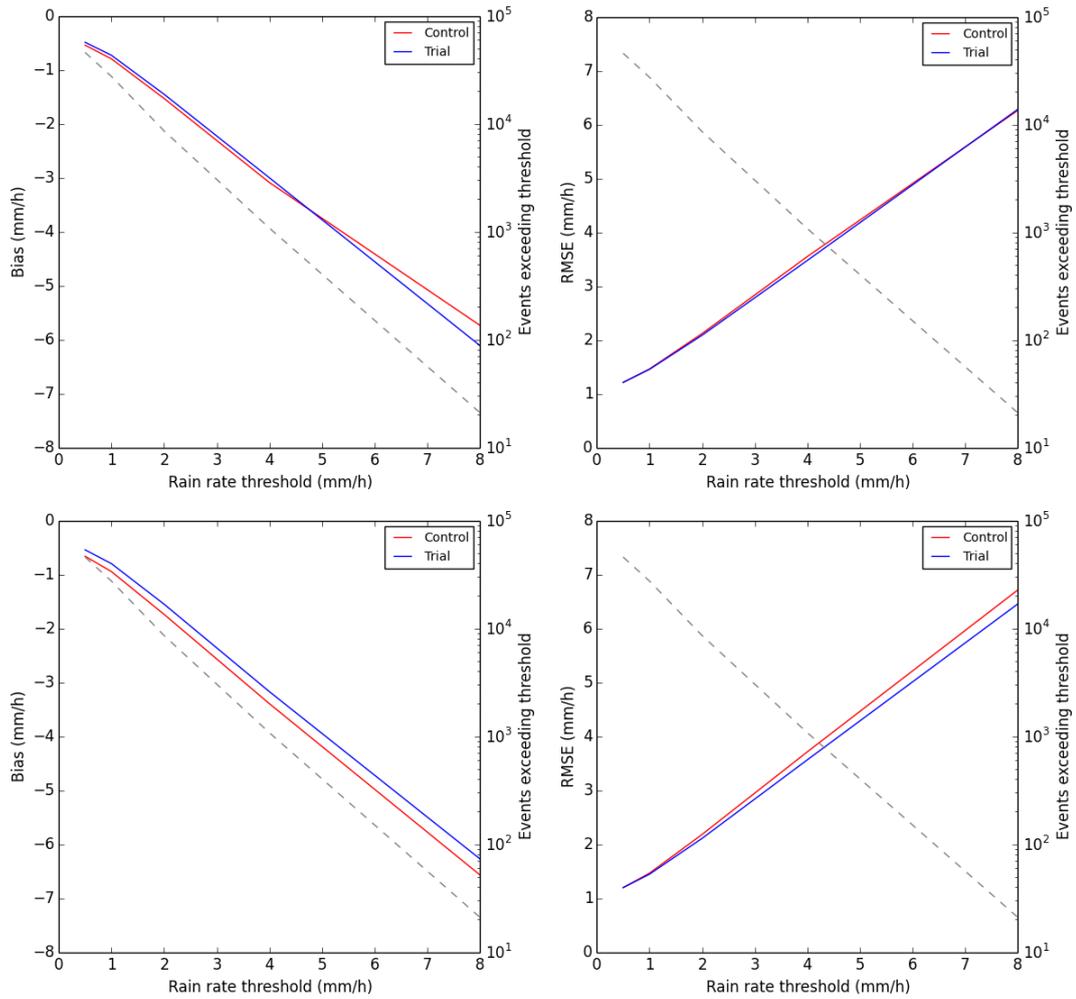


Figure 7.9: Aggregated single site gauge-adjusted (top) and unadjusted (bottom) hourly accumulation bias and RMSE with rain rate over the live trial period (16th-31st March 2017).

from this analysis, due to differences in hardware and the resulting data quality from these radars. Aggregated single site statistics for all other radars are shown in figure 7.9.

The gauge-adjusted single site radar accumulations in the top panel of figure 7.9 show a similar signal to the composite data. The modified VPR produces small benefit at low rain rates over the current scheme, and a degradation towards higher accumulations. In the unadjusted data, however, there is an improvement in both bias and RMSE at all rain intensities. This is consistent with data obtained from the more limited stratiform reprocessing trials (figure 7.6), and supports the inference that changes in gauge adjustment factors are responsible for the overall degradation in the composite.

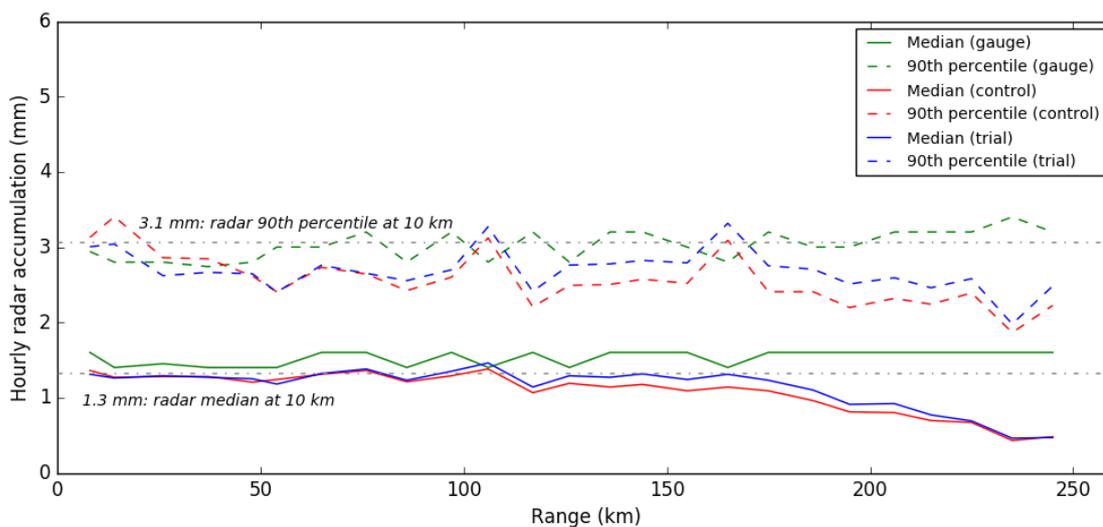


Figure 7.10: Radar and gauge rainfall distribution parameters with range for live trial data, 16th-31st March 2017. The median and 90th percentile of hourly accumulations measured at different ranges are plotted for the subset of data where the hourly rain gauge accumulation was greater than 1 mm.

Comparing radar and gauge rainfall distributions with range

A further difference between composite and single site gauge comparisons is the distribution of data available with range. Given the better quality and representativity of measurements close to the radar, composite accumulations are heavily skewed towards QPEs from shorter ranges. Single site accumulations, however, represent a broader distribution of ranges, and provide the opportunity to investigate further the impact of changes to the VPR. In particular, single site unadjusted data can be used to illustrate the distribution of radar QPE values with range, and how this is affected by the new vertical profile.

Figure 7.10 illustrates the changes in rain rate distribution with range measured by radar for the current and trial VPR shape. With perfect correction for VPR, and neglecting the impacts of beam overshoot, the radar measured distribution should not vary with range. Accumulations are filtered to remove points where the gauge accumulation was less than 1 mm, to remove the effect of minimum detectable signal, which due to the extreme skewness of the climatological rainfall distribution introduces an increasing trend in percentiles of the radar measured distribution with range. The median and 90th percentile of the truncated radar distributions are plotted in 10 km range bins, and compared with the same percentiles for the distribution of rain gauge accumulations.

For hourly radar accumulations, the median for both the trial and control (K94) profiles is approximately constant out to a range of 100 km. There is also good agreement with

the rain gauge accumulation median in this range. Beyond 100 km, both radar medians begin to decrease. The median of the trial distribution, however, remains closer to its value at 10 km range (1.3 mm) than the control, out to a range of 220 km. At this point the radar median accumulations are around half the value of the gauge median.

The 90th percentile of rainfall accumulations for both radar and gauge is noisier, particularly in the first 50 km. This likely represents genuine variability in the underlying rainfall distribution, given that the number of events exceeding 1 mm h^{-1} is only of order 10^4 for the entire dataset (before subsetting by range), and the drop off in event count with rate is exponential (see eg figure 7.9). Even with this instability however, in general, the trial percentile values remain more consistent with range and track the gauge distribution more closely than percentiles from the control VPR.

It is clear from figure 7.10, as well as several previous figures in this thesis analysing trends in error with range (figures 5.7, 6.7, 6.9, and 7.2-7.5), that the bias in radar QPEs is far from constant across the radar domain. This chapter in particular highlights that the difference between radar and gauge accumulations has a large component that is strongly dependent on range. These results suggest that mean field bias gauge adjustment is not a sensible approach to calibrating radar QPEs.

7.6 Conclusions

This chapter assessed the impact on QPEs of a minor change to the Kitchen et al. (1994) stratiform VPR to account for the difference in reflectivities immediately above and below the melting layer. Justification for such a change was presented in the context of the existing literature, as well as observations made during the course of this PhD; and the continuing need to account for residual bias in long range QPEs was demonstrated in section 7.3. A proposed offset of 2 dB was tested via simulation study on short range profiles from the Wardon Hill research radar, and was shown to be effective in reducing long range bias with no impact on the root mean squared error (RMSE). This result was then extended to VPR correction in PPIs using a winter trial period dominated by stratiform precipitation. Alongside small but measurable reductions in gauge-radar bias and RMSE, this trial showed a reduction in range-dependence of the errors on radar measured rain rates (figure 7.5). This demonstrates improvements in the underlying representativity of the modified VPR, leading to a more spatially consistent retrieval of rain rate values.

Building on the positive results from both simulation and stratiform controlled trials, a live trial of the updated stratiform VPR was run over the second half of March 2017. The conclusions from this trial were twofold. Firstly, these results in conjunction with

observations from chapter 4 support a conclusion that the updated VPR shape is more representative of atmospheric conditions than the current K94 operational profile. This is reflected in the improved consistency of radar measured rainfall distributions with range (figure 7.10), and in small improvements in gauge-radar statistics for single site unadjusted gauge comparisons. This result is consistent with VPR shapes proposed in previous literature (eg Fabry and Zawadzki, 1995; Rico-Ramirez et al., 2005, see section 2.3 for a detailed review), but is the first time this has been demonstrated by objective comparison of the impact of different VPR shapes on QPEs.

However, in terms of an operational Met Office implementation, the interaction of the new VPR scheme with the gauge adjustment factor scheme is cause for concern. The reduction in gauge adjustment factor values caused by increased radar QPEs leads to a degradation in the statistical performance of the composite, with increasing underestimation at all rain intensities. The impact was on RMSE was more mixed, causing detriment only at moderate rain rates (2-4 mm h⁻¹) over the course of the trial. Further research is needed to establish the reasons for the gauge adjustment factor behaviour and to minimise adverse impacts, particularly at high rain rates, before this modification can be implemented operationally.

Chapter 8

Summary and outlook

8.1 Improving local classification and correction for the vertical reflectivity profile

The vertical profile of reflectivity (VPR) provides a relationship between a meteorological radar measurement at height and the corresponding reflectivity at ground level. Characterisation and correction for VPR is therefore a crucial step in obtaining accurate radar surface rainfall estimates (QPEs). The Met Office “Radarnet” software applies a local correction, using the idealised stratiform profile defined by Kitchen et al. (1994) with refinements from Kitchen (1997) to correct for VPR independently at every point in the radar domain. This method has performed reliably in an operational setting for over twenty years.

The recent Weather Radar Network Renewal project has provided opportunities to review and refine many aspects of QPE as performed in the UK. This thesis aimed to apply the high quality dual polarisation information available from the upgraded UK radar network to improve determination and correction for VPR. The work has focused on the linear depolarisation ratio (LDR), as this has shown skill in distinguishing the large melting snowflakes responsible for stratiform bright band from other hydrometeor types (Smyth and Illingworth, 1998; Illingworth and Thompson, 2011). Two specific hypotheses were tested:

1. That the use of LDR to classify VPRs, alongside better characterisation of VPR shapes, can lead to measurable improvements in radar QPEs at the local scale.
2. That, contrary to the assumptions of some existing literature (eg Steiner et al., 1995), it is worth correcting for VPR in certain non-bright band situations.

To test the hypotheses above, a dataset of high resolution vertical profiles of reflectivity and LDR, large enough to be considered representative of the UK climatology, was collected (chapter 3). It was shown using these profiles that LDR values in the melting layer could skilfully distinguish between VPRs with bright band and those without (chapter 4, published in Sandford et al. (2017)). A high resolution LDR melting layer peak of -20 dB has a 42% probability of detecting non-bright band precipitation, for a false alarm rate of 6%, compared with only a 6% hit rate for the reflectivity-based convective diagnosis criterion currently used in the UK ($Z_1 > 30$ dBZ). These findings were extended via a simulation study in chapter 5 to develop an LDR-based classification algorithm for radar PPIs. This included an illustration of the range-insensitivity of peak LDR values in the melting layer, due to the dominance of peak depolarisation values in convolving the atmospheric LDR with the radar beam power profile. The initial LDR-based criterion yielded reductions in underestimation bias of order 0.5 mm h^{-1} for simulated QPEs from non-bright band profiles (figure 5.7). Issues relating to operational LDR-based classification were also addressed, including the need for a scan strategy to maximise data availability and representativity. The final classification algorithm was shown to improve radar QPEs in a high impact non-stratiform case study, reducing both bias and RMSE by more than 1 mm h^{-1} for the highest hourly accumulation categories ($12\text{-}16 \text{ mm h}^{-1}$).

The second hypothesis was addressed in chapter 6, which made an objective comparison of four different idealised non-bright band VPR shapes informed by observations and previous literature. The application of a simulation framework to compare objectively the performance of different profile shapes using a real profile dataset is unique in the VPR literature. The profile shape developed in this chapter yielded further improvements on the non-stratiform case study of chapter 5, including reductions in underestimation bias and RMSE of over 2 mm h^{-1} for accumulations greater than 12 mm h^{-1} . This demonstrates that it is possible, using only a simple VPR model, to improve upon a “constant reflectivity with height” assumption in cases without a stratiform bright band.

The high resolution profile dataset developed in chapter 3 also presented opportunities to revisit and refine the work of Kitchen et al. (1994) on stratiform VPRs. Chapter 7 proposed a small change to the idealised bright band profile, introducing an offset of 2 dB between reflectivities immediately above and below the bright band. This change was justified on the basis of previous literature and observations, including an illustration of residual range-dependent bias in VPR-corrected QPEs at long range, and was found to reduce the systematic negative bias of long range QPEs.

The results of chapter 7 call into question the usefulness of mean field bias adjustment using rain gauge measurements. The example of VPR illustrates a broader problem, that the benefits of improving the underlying physical processing of radar data can be negated completely by the application of a broad brush correction that does not reflect

the structured nature of the residual errors. While mean field gauge adjustment was originally designed to correct for errors such as miscalibration and inappropriate ZR relationships (Harrison et al., 2000; Figueras I Ventura and Tabary, 2013) - both of which are, on average, spatially invariant - advances in radar hardware and processing methods have reduced the influence of these errors compared to those with more obvious spatial structure, such as attenuation and VPR. Particularly following the introduction of dual polarisation rain rate estimators, which can only be exploited at ranges where the radar beam is located below the melting layer, the influence of range on overall QPE errors has become even more pronounced. A gauge adjustment factor that varies with range may therefore be more suitable for modern radar applications.

A mean field bias correction designed to optimise agreement between radar accumulations and the gauges used to calculate the adjustment will always evaluate more favourably than data which are not bias corrected (as exemplified in figure 7.9). However, further assessment of gauge-adjusted radar data using independent rain gauges would be an interesting objective test of the validity of such schemes.

8.2 Future work

8.2.1 Correcting for the VPR in non-bright band conditions

Chapter 6 developed a highly simplified non-bright band profile shape based primarily on observations, both from the high resolution VPR dataset (chapter 3) and previous literature. This is an initial demonstration of an idealised linear VPR that can improve QPEs, on average, in any non-bright band conditions (except in the presence of hail). The performance achieved using only this first approximation suggests there could be benefits to further research into the behaviour of non-bright band profiles.

The microphysics behind the occurrence of “compact ice” profiles has been understood, at least in concept, since the 1990s. This thesis finds anecdotal evidence in favour of the Fabry and Zawadzki (1995) hypothesis, which is also supported anecdotally by Matrosov et al. (2016). However, to date, there has been no attempt to verify the hypothesis directly, and in particular no consideration of the signature this microphysics may leave in other dual polarisation variables. For example, if the compact ice profile truly results from a population of smaller, more uniform ice particles than stratiform “bright band” precipitation, then the copolar correlation coefficient ρ_{hv} might be expected to be higher in the ice layer than in stratiform precipitation, due to the more uniform mix of high density ice particles. Investigating such signatures could provide opportunities to improve identification of compact ice profiles in regions outside the melting layer, where no useful LDR measurement is available.

A useful side effect of such direct investigations would be the potential to refine or parameterise the shape of the non-bright band profile. Two particular areas of interest are the rain-to-ice reflectivity offset - which was estimated empirically in chapter 6 by averaging values from the Wardon Hill VPR dataset - and the gradient of reflectivity in the ice above the melting layer. The simplicity of the approximations made in this thesis are a reflection of the limited size of the non-bright band profile dataset, which made up only 16% (1080 profiles) of the climatological VPR sample. A larger dataset containing several thousand non-bright band profiles would ideally be collected for more robust investigation of observed profile shapes. However, the comparative rarity of non-bright band precipitation means it may not be possible to characterise these profiles using purely empirical approach such as that adopted by Kitchen et al. (1994). The use of microphysically-based models to inform idealised profile shapes, such as in the method developed by Kirstetter et al. (2013), could be a viable alternative to aid parameterisation of compact ice profiles.

Beyond the shape of the non-bright band profile, some investigation would be beneficial to characterise its behaviour on different spatial and temporal scales. This thesis tested a non-bright band profile within the specific framework of the Kitchen et al. (1994) VPR correction scheme, which by iterative deconvolution fits an idealised profile shape to the measured reflectivity at each pixel in the radar domain. The use of idealised profiles in this way is not unique to the Met Office (eg Rico-Ramirez et al., 2005; Matrosov et al., 2007; Tabary, 2007, see chapter 2). However, the ability of the Kitchen et al. (1994) scheme to resolve and correct for very local features is not typical of the existing VPR literature. Further testing of the impacts of this profile in more global VPR schemes would broaden the applicability of these findings.

8.2.2 Extending LDR findings to operational systems

This thesis focused on LDR due to its demonstrated (although previously unquantified) skill in detecting the presence of bright band, as opposed to non-bright band melting. However, very few radar networks make operational measurements of LDR. This makes it difficult to extend the specific findings on LDR classification skill into the wider operational radar community.

Regardless of specifics, however, this thesis has illustrated in principle the benefits of classifying VPRs on a local level using dual polarisation information. There is potential to investigate other methods of local classification, for example using the outputs of hydrometeor classification algorithms to provide proxy information as to the type of precipitation present. Hall et al. (2015) and Rico-Ramirez et al. (2005) both demonstrate the use of hydrometeor classifiers to provide information on the local VPR - but both

algorithms require LDR as a part of the classification process. Future work to apply this operationally would need to restrict the classifier inputs to SHV mode parameters only.

LDR is not the only source of depolarisation information available from operational radars. Ryzhkov et al. (2017) show how the depolarisation ratio (DR) can be calculated using ρ_{hv} and Z_{DR} from SHV mode measurements. However, DR differs from LDR in that it represents the average depolarising nature of the target hydrometeors for all incident polarisations. This means DR is sensitive to aspect ratio, but not orientation. The clear signal given by LDR in the presence of large, canted melting snowflakes is therefore significantly reduced in DR, which does not respond to canting angle. Further investigation would be needed to determine whether DR (as a function of ρ_{hv} and Z_{DR}) could contribute useful information towards identifying stratiform bright band.

8.2.3 The question of scale: revisited

As discussed in chapter 2, the scale on which adjustments for VPR should be applied is an open question in the literature. The widespread use of global VPR schemes has advantages, largely in addressing the underdetermination problem with no assumptions as to the background microphysics, and no requirement for gridded ancillary data such as the height of the wet bulb freezing level. However, when additional information is available, local schemes have the potential to produce corrections that are much more accurate and representative at the local level. Where a radar image contains a mixture of stratiform and convective rainfall, global averaging results in both over-correcting in convection and under-correcting for bright band in stratiform rainfall; hence the ability to define the VPR at the local scale is crucial in such cases. Even in purely stratiform rain, VPRs averaged over smaller regions have been shown to represent more accurately local variations in bright band shape and intensity, leading to improved surface reflectivity and rainfall estimates (Vignal et al., 2000). Added to this, the assumption of widespread consistency in a global VPR scheme limits the benefits that can be achieved from increasing information present in radar data: specifically the microphysical detail that can be derived from dual polarisation measurements.

This thesis has demonstrated the benefits of using LDR measurements at the kilometre scale to identify the microphysical situation and to choose a VPR shape suitable for the local conditions. Such information can only be exploited by a scheme capable of correcting for VPR on an equally local level. It clearly follows that local schemes have the potential to achieve more accurate QPEs than global schemes, not only through spatial resolution, but through the ability to make effective use of detailed research findings to improve the representativity of the parameterised VPRs typically used in local corrections (section 2.3). However, to date, there has been very little research addressing directly the

question of scale. Only the findings of Vignal et al. (2000) exist to support the hypothesis that relatively local VPR determination and correction can provide measurable benefits over global schemes. Objective comparison of local and global schemes will be key to informing the research direction for VPR corrections, and exploiting the benefits of recent and future radar developments for long range operational QPEs.

Bibliography

- Adams, D., 2015: Using Polarimetric Self-Consistency to Calibrate Reflectivity. Tech. rep., Met Office.
- Anagnostou, E., 2004: A Convective/Stratiform Precipitation Classification Algorithm for Volume Scanning Weather Radar Observations. *Meteorological Applications*, **11**, 291–300.
- Andrieu, H. and J. D. Creutin, 1995a: Identification of Vertical Profiles of Radar Reflectivity for Hydrological Applications Using an Inverse Method. Part I: Formulation. *Journal of Applied Meteorology*, **34**, 225–239.
- Andrieu, H. and J. D. Creutin, 1995b: Identification of Vertical Profiles of Radar Reflectivity for Hydrological Applications Using an Inverse Method. Part II: Sensitivity Analysis and Case Study. *Journal of Applied Meteorology*, **34**, 240–259.
- Battan, L. J., 1973: *Radar Observation of the Atmosphere*. The University of Chicago Press.
- Berne, A. and W. Krajewski, 2013: Radar for hydrology: Unfulfilled promise or unrecognized potential? *Advances in Water Resources*, **51**, 357–366.
- Biggerstaff, M. and S. Listemaa, 2000: An Improved Scheme for Convective/Stratiform Echo Classification Using Radar Reflectivity. *Journal of Applied Meteorology*, **39**, 2129–2150.
- Boodoo, S., D. Hudak, N. Donaldson, and M. Leduc, 2010: Application of Dual-Polarization Radar Melting-Layer Detection Algorithm. *Journal of Applied Meteorology and Climatology*, **49**, 1779–1793.
- Boudevillain, B. and H. Andrieu, 2003: Assessment of Vertically Integrated Liquid (VIL) Water Content Radar Measurement. *Journal of Atmospheric and Oceanic Technology*, **20**, 807–819.
- Bowler, N., C. Pierce, and A. Seed, 2006: STEPS: A probabilistic precipitation forecast-

- ing scheme which merges an extrapolation nowcast with downscaled NWP. *Quarterly Journal of the Royal Meteorological Society*, **132**, 2127–2155.
- Brandes, E., G. Zhang, and J. Vivekanandan, 2002: Experiments in Rainfall Estimation with a Polarimetric Radar in a Subtropical Environment. *Journal of Applied Meteorology*, **41**, 674–685.
- Brandes, E., G. Zhang, and J. Vivekanandan, 2003: An Evaluation of a Drop Distribution-Based Polarimetric Radar Rainfall Estimator. *Journal of Applied Meteorology*, **42**, 652–660.
- Brandes, E. A., G. Zhang, and J. Vivekanandan, 2004: Drop Size Distribution Retrieval with Polarimetric Radar: Model and Application. *Journal of Applied Meteorology*, **43**, 461–475.
- Bringi, V. and V. Chandrasekar, 2001: *Polarimetric Doppler Weather Radar: Principles and Applications*. Cambridge University Press.
- Bringi, V., V. Chandrasekar, N. Balakrishnan, and S. Zrnić, 1990: An Examination of Propagation Effects in Rainfall on Radar Measurements at Microwave Frequencies. *Journal of Atmospheric and Oceanic Technology*, **7**, 829–840.
- Bringi, V., C. Williams, M. Thurai, and P. May, 2009: Using Dual-Polarized Radar and Dual-Frequency Profiler for DSD Characterization: A Case Study from Darwin, Australia. *Journal of Atmospheric and Oceanic Technology*, **26**, 2107–2122.
- Brown, A., S. Milton, M. Cullen, B. Golding, J. Mitchell, and A. Shelly, 2012: Unified Modelling and Prediction of Weather and Climate: A 25-Year Journey. *Bulletin of the American Meteorological Society*, 1865–1877.
- Carey, L., S. Rutledge, D. Ahijevych, and T. Keenan, 2000: Correcting Propagation Effects in C-Band Polarimetric Radar Observations of Tropical Convection Using Differential Propagation Phase. *Journal of Applied Meteorology*, **39**, 1405–1432.
- Carruthers, D. J. and T. W. Chouarton, 1983: A model of the feeder-seeder mechanism of orographic rain including stratification and wind-drift effects. *Quarterly Journal of the Royal Meteorological Society*, **109**, 575–588.
- Darlington, T., 2014: Development, Evaluation and Applications of the Cyclops-DP Weather Radar Processing System. Ph.D. thesis, University of Bath.
- Darlington, T., M. Edwards, J. Sugier, and M. Kitchen, 2004: Opportunities presented by advanced off-the-shelf signal processing technology. *Third European Conference on Radar in Meteorology and Hydrology*, Visby.

- Darlington, T., et al., 2016: Designing an operational C-band radar to realise the benefits of dual-polarization. Tech. rep., Met Office.
- Delrieu, G., B. Boudevillain, J. Nicol, B. Chapon, and P. Kirstetter, 2009: Bollene-2002 Experiment: Radar Quantitative Precipitation Estimation in the Cevennes-Vivarais Region, France. *Journal of Applied Meteorology and Climatology*, **48**, 1422–1447.
- Doviak, R. and D. Zrnić, 1993: *Doppler Radar and Weather Observations*. 2d ed., Academic Press.
- Fabry, F. and W. Szyrmer, 1999: Modeling of the Melting Layer. Part II: Electromagnetic. *Journal of the Atmospheric Sciences*, **56**, 3593–3600.
- Fabry, F. and I. Zawadzki, 1995: Long-Term Radar Observations of the Melting Layer of Precipitation and Their Interpretation. *Journal of the Atmospheric Sciences*, **52**, 838–851.
- Figueras I Ventura, J. and P. Tabary, 2013: The New French Operational Polarimetric Radar Rainfall Rate Product. *Journal of Applied Meteorology and Climatology*, **52**, 1817–1835.
- Frech, M., M. Hagen, and T. Mammen, 2017: Monitoring the Absolute Calibration of a Polarimetric Weather Radar. *Journal of Atmospheric and Oceanic Technology*, **34**, 599–615.
- Gabell, a. M., P. Speirs, U. Hamann, U. Germann, and A. Berne, 2017: Measurement of Precipitation in the Alps Using Dual-Polarization C-Band Ground-Based Radars, the GPM Spaceborne Ku-Band Radar, and Rain Gauges. *Remote Sensing*, **9**, 1147–1165.
- Gabella, M., M. Boscacci, M. Sartori, and U. Germann, 2016: Calibration Accuracy of the Dual-Polarization Receivers of the C-Band Swiss Weather Radar Network. *Atmosphere*, **7**, 76.
- Germann, U. and J. Joss, 2002: Mesobeta Profiles to Extrapolate Radar Precipitation Measurements above the Alps to the Ground Level. *Journal of Applied Meteorology*, **41**, 542–557.
- Giangrande, S., M. Krause, and A. Ryzhkov, 2008: Automatic Designation of the Melting Layer with a Polarimetric Prototype of the WSR-88D Radar. *Journal of Applied Meteorology and Climatology*, **47**, 1354–1364.
- Giangrande, S. and A. Ryzhkov, 2008: Estimation of Rainfall Based on the Results of Polarimetric Echo Classification. *Journal of Applied Meteorology and Climatology*, **47**, 2445–2462.

- Goddard, J. W. F., J. D. Eastment, and J. Tan, 1994: Self-consistent measurements of differential phase and differential reflectivity in rain. *IEEE International Geoscience and Remote Sensing Symposium*, 369–371.
- Goudenhoofdt, E. and L. Delobbe, 2016: Generation and Verification of Rainfall Estimates from 10-Yr Volumetric Weather Radar Measurements. *Journal of Hydrometeorology*, **17**, 1223–1242.
- Gourley, J., A. Illingworth, and P. Tabary, 2009: Absolute Calibration of Radar Reflectivity Using Redundancy of the Polarization Observations and Implied Constraints on Drop Shapes. *Journal of Atmospheric and Oceanic Technology*, **26**, 689–703.
- Gunn, K. and T. East, 1954: The Microwave Properties of Precipitation Particles. *Quarterly Journal of the Royal Meteorological Society*, **80**, 522–545.
- Hall, W., M. Rico-Ramirez, and S. Krämer, 2015: Classification and Correction of the Bright Band using an Operational C-band Polarimetric Radar. *Journal of Hydrology*, **531**, 248–258.
- Hardaker, P., A. Holt, and C. Collier, 1995: A melting-layer model and its use in correcting for the bright band in single-polarization radar echoes. *Quarterly Journal of the Royal Meteorological Society*, **121**, 495–525.
- Harrison, D., S. Driscoll, and M. Kitchen, 2000: Improving Precipitation Estimates from Weather Radar using Quality Control and Correction Techniques. *Meteorological Applications*, **6**, 135–144.
- Harrison, D., K. Norman, C. Pierce, and N. Gaussiat, 2012: Radar products for hydrological applications in the UK. *Proceedings of ICE Water Management*, **165**, 89–103.
- Hazenberg, P., P. Torfs, H. Leijnse, G. Delrieu, and R. Uijlenheot, 2013: Identification and uncertainty estimation of vertical reflectivity profiles using a Lagrangian approach to support quantitative precipitation measurements by weather radar. *Journal of Geophysical Research – Atmospheres*, **118**, 10 243–10 261.
- Helmert, K., et al., 2014: DWDs new radar network and algorithm post-processing chain. *Eighth European Conference on Radar in Meteorology and Hydrology*, Garmisch-Partenkirchen, Germany.
- Herzogh, P. and R. Jameson, 1992: Observing Precipitation through Dual-Polarization Radar Measurements. *Bulletin of the American Meteorological Society*, **73**, 1365–1374.
- Heyraud, C., W. Szyrmer, S. Laroche, and I. Zawadzki, 2008: Modeling of the Melting Layer. Part IV: Brightband Bulk Parameterization. *Journal of the Atmospheric Sciences*, **65**, 1991–2001.

- Hitschfeld, W. and J. Bordan, 1954: Errors Inherent in the Radar Measurement of Rainfall at Attenuating Wavelengths. *Journal of Meteorology*, **11**, 58–67.
- Hubbert, J., M. Dixon, S. Ellis, and G. Meymaris, 2009: Weather Radar Ground Clutter. Part I: Identification, Modeling and Simulation. *Journal of Atmospheric and Oceanic Technology*, **26**, 1165–1180.
- Hubbert, J., S. Ellis, M. Dixon, and G. Meymaris, 2010: Modelling, Error Analysis and Evaluation of Dual-Polarization Variables Obtained from Simultaneous Horizontal and Vertical Polarization Transmit Radar. Part I: Modelling and Antenna Errors. *Journal of Atmospheric and Oceanic Technology*, **27**, 1583–1598.
- Husnoo, N., T. Darlington, R. Thompson, and A. Illingworth, 2018: Evaluation of reflectivity/polarimetry/radiometry-based C-band attenuation correction scheme using neighbouring radars. *Tenth European Conference on Radar in Meteorology and Hydrology*, Ede, Netherlands.
- Illingworth, A., 2004: Improved Precipitation Rates and Data Quality by Using Polarimetric Measurements. *Weather Radar: Principles and Advanced Applications*, P. Meischner, Ed., Springer, chap. 5, 130–166.
- Illingworth, A. and R. Thompson, 2011: Radar Bright Band Correction using the Linear Depolarisation Ratio. *International Symposium on Weather Radar and Hydrology*, Exeter, UK, 64–68.
- Johnson, J., P. MacKeen, A. Witt, E. Mitchell, G. Stumpf, M. Eilts, and K. Thomas, 1998: The Storm Cell Identification and Tracking Algorithm: An Enhanced WSR-88D Algorithm. *Weather and Forecasting*, **13**, 263–276.
- Jones, C. and B. Macpherson, 1997: A latent heat nudging scheme for the assimilation of precipitation data into an operational mesoscale model. *Meteorological Applications*, **4**, 269–277.
- Kalogiros, J., M. Anagnostou, E. Anagnostou, M. Montipoli, E. Picciotti, and F. Marzano, 2013: Correction of Polarimetric Radar Reflectivity Measurements and Rainfall Estimates for Apparent Vertical Profile in Stratiform Rain. *Journal of Applied Meteorology and Climatology*, **52**, 1170–1186.
- Keat, W. J., C. D. Westbrook, and A. J. Illingworth, 2016: High-Precision Measurements of the Copolar Correlation Coefficient: Non-Gaussian Errors and Retrieval of the Dispersion Parameter μ in Rainfall. *Journal of Applied Meteorology and Climatology*, **55**, 1615–1632.
- Kirstetter, P., H. Andrieu, B. Boudevillain, and G. Delrieu, 2013: A Physically Based

- Identification of Vertical Profiles of Reflectivity from Volume Scan Radar Data. *Journal of Applied Meteorology and Climatology*, **52**, 1645–1663.
- Kirstetter, P., H. Andrieu, G. Delrieu, and B. Boudevillain, 2010: Identification of Vertical Profiles of Reflectivity for Correction of Volumetric Radar Data Using Rainfall Classification. *Journal of Applied Meteorology and Climatology: Notes and Correspondence*, **49**, 2167–2180.
- Kitchen, M., 1997: Towards improved radar estimates of surface precipitation rate at long range. *Quarterly Journal of the Royal Meteorological Society*, **123**, 145–163.
- Kitchen, M., R. Brown, and A. Davies, 1994: Real-time correction of weather radar data for the effects of bright band, range and orographic growth in widespread precipitation. *Quarterly Journal of the Royal Meteorological Society*, **120**, 1231–1254.
- Kitchen, M. and A. Illingworth, 2011: From Observations to Forecasts - Part 13: The UK weather radar network - past, present and future. *Weather*, **66**, 291–297, doi: 10.1002/wea.861, URL <http://dx.doi.org/10.1002/wea.861>.
- Kitchen, M. and P. M. Jackson, 1993: Weather Radar Performance at Long Range - Simulated and Observed. *Journal of Applied Meteorology*, **32**, 975–985.
- Koistinen, J. and H. Pohjola, 2014: Estimation of Ground-Level Reflectivity Factor in Operational Weather Radar Networks Using VPR-Based Correction Ensembles. *Journal of Applied Meteorology and Climatology*, **53**, 2394–2411.
- Kumjian, M., 2013: Principles and Applications of Dual-Polarization Weather Radar. Part I: Description of the Polarimetric Radar Variables. *Journal of Operational Meteorology*, **1**, 226–242.
- Kumjian, M., S. Mishra, S. Giangrande, T. Toto, A. Ryzhkov, and A. Bansemmer, 2016: Polarimetric radar and aircraft observations of saggy bright bands during MC3E. *Journal of Geophysical Research – Atmospheres*, **121**, 3584–3607.
- Le Bastard, T., O. Caumont, N. Gaussiat, and F. Karbou, 2018: On the use of NWP model outputs to produce radar VPR corrections. *Tenth European Conference on Radar in Meteorology and Hydrology*, Ede, Netherlands.
- Lewis, H., D. Harrison, and M. Kitchen, 2007: Local vertical profile corrections using data from multiple scan elevations. *33rd AMS Conference on Radar Meteorology*, Cairns, Australia.
- Lewis, H., et al., 2015: From months to minutes – exploring the value of high-resolution rainfall observation and prediction during the UK winter storms of 2013/2014. *Meteorological Applications*, **22**, 90–104.

- Marshall, J. and W. Palmer, 1948: The Distribution of Raindrops with Size. *Journal of Meteorology: Shorter Contributions*, **5**, 165–166.
- Matrosov, S., R. Cifelli, P. Neiman, and A. White, 2016: Radar Rain-Rate Estimators and Their Variability due to Rainfall Type: An Assessment Based on Hydrometeorology Testbed Data from the Southeastern United States. *Journal of Applied Meteorology and Climatology*, **55**, 1345–1358.
- Matrosov, S., K. Clark, and D. Kingsmill, 2007: A Polarimetric Radar Approach to Identify Rain, Melting-Layer, and Snow Regions for Applying Corrections to Vertical Profiles of Reflectivity. *Journal of Applied Meteorology and Climatology*, **46**, 154–166.
- Mittermaier, M. P., R. J. Hogan, and A. J. Illingworth, 2004: Using mesoscale model winds for correcting wind-drift errors in radar estimates of surface rainfall. *Quarterly Journal of the Royal Meteorological Society*, **130**, 2105–2123.
- Mittermaier, M. P. and A. J. Illingworth, 2003: Comparison of model-derived and radar-observed freezing-level heights: Implications for vertical reflectivity profile-correction schemes. *Quarterly Journal of the Royal Meteorological Society*, **129**, 83–95.
- Park, H., A. Ryzhkov, D. Zrnić, and K. Kim, 2009: The Hydrometeor Classification Algorithm for the Polarimetric WSR-88D: Description and Application to an MCS. *Weather and Forecasting*, **24**, 730–748.
- Park, S.-G., M. Maki, K. Iwanami, V. Bringi, and V. Chandrasekar, 2005: Correction of Radar Reflectivity and Differential Reflectivity for Rain Attenuation at X Band. Part II: Evaluation and Application. *Journal of Atmospheric and Oceanic Technology*, **22**, 1633–1655.
- Penide, G., A. Protat, V. Kumar, and P. May, 2013: Comparison of Two Convective/Stratiform Precipitation Classification Techniques: Radar Reflectivity Texture versus Drop Size Distribution-Based Approach. *Journal of Atmospheric and Oceanic Technology*, **30**, 2788–2797.
- Qi, Y. and J. Zhang, 2017: A Physically Based Two-Dimensional Seamless Reflectivity Mosaic for Radar QPE in the MRMS System. *Journal of Hydrometeorology*, **18**, 1327–1340.
- Qi, Y., J. Zhang, and P. Zhang, 2013a: A real-time automated convective and stratiform precipitation segregation algorithm in native radar coordinates. *Quarterly Journal of the Royal Meteorological Society*, **139**, 2233–2240.
- Qi, Y., J. Zhang, P. Zhang, and Q. Cao, 2013b: VPR correction of bright band effects in radar QPEs using polarimetric radar observations. *Journal of Geophysical Research – Atmospheres*, **118**, 3627–3633.

- Rasmussen, R. and A. Heymsfield, 1987: Melting and Shedding of Graupel and Hail. Part I: Model Physics. *Journal of the Atmospheric Sciences*, **44**, 2754–2763.
- Rico-Ramirez, M. and I. Cluckie, 2008: Classification of ground clutter and anomalous propagation using dual-polarization weather radar. *IEEE Transactions on Geoscience and Remote Sensing*, **46**, 1892–1904.
- Rico-Ramirez, M., I. Cluckie, and D. Han, 2005: Correction of the bright band using dual-polarisation radar. *Atmospheric Science Letters*, **6**, 40–46.
- Rigo, T. and M. Llasat, 2004: A methodology for the classification of convective structures using meteorological radar: Application to heavy rainfall events on the Mediterranean coast of the Iberian Peninsula. *Natural Hazards and Earth System Sciences*, **4**, 59–68.
- Rosenfeld, D., E. Amitai, and D. Wolff, 1995: Classification of Rain Regimes by the Three-Dimensional Properties of Reflectivity Fields. *Journal of Applied Meteorology*, **34**, 198–211.
- Ryzhkov, A., et al., 2017: Estimation of Depolarization Ratio Using Weather Radars with Simultaneous Transmission/Reception. *Journal of Applied Meteorology and Climatology*, **56**, 1797–1816.
- Ryzhkov, A. V., T. J. Schuur, D. W. Burgess, P. L. Heinselman, S. E. Giangrande, and D. S. Zrnić, 2005: The Joint Polarization Experiment: Polarimetric Rainfall Measurements and Hydrometeor Classification. *Bulletin of the American Meteorological Society*, **86**, 809–824.
- Ryzhkov, A. V. and D. S. Zrnić, 1995: Comparison of Dual-Polarization Radar Estimators of Rain. *Journal of Atmospheric and Oceanic Technology*, **12**, 249–256.
- Sánchez-Diezma, R., I. Zawadzki, and D. Sempere-Torres, 2000: Identification of the bright band through the analysis of volumetric radar data. *Journal of Geophysical Research – Atmospheres*, **105**, 2225–2236.
- Sandford, C., 2015: Correcting for wind drift in high resolution radar rainfall products: a feasibility study. *Journal of Hydrology*, **531**, 284–295.
- Sandford, C., A. Illingworth, and R. Thompson, 2017: The Potential Use of the Linear Depolarisation Ratio to Distinguish between Convective and Stratiform Rainfall to Improve Radar Rain-Rate Estimates. *Journal of Applied Meteorology and Climatology*, **56**, 2927–2940.
- Sauvageot, H., 1992: *Radar Meteorology*. Artech House Inc.

- Scovell, R., H. Lewis, D. Harrison, and M. Kitchen, 2008: Local vertical profile corrections using data from multiple scan elevations. *Fifth European Conference on Radar in Meteorology and Hydrology*, Helsinki, Finland.
- Sinden, G., 2007: Characteristics of the UK wind resource: Long-term patterns and relationship to electricity demand. *Energy Policy*, **35**, 112–127.
- Smyth, T. J. and A. J. Illingworth, 1998: Radar estimates of rainfall rates at the ground in bright band and non-bright band events. *Quarterly Journal of the Royal Meteorological Society*, **124**, 2417–2434.
- Snow, K., A. Seed, and G. Takacs, 2011: Determining the vertical profile of reflectivity using radar observations at long range. *International Symposium on Weather Radar and Hydrology*, Exeter, UK.
- Steiner, M., R. Houze, and S. Yuter, 1995: Climatological Characterization of Three-Dimensional Storm Structure from Operational Radar and Rain Gauge Data. *Journal of Applied Meteorology*, **34**, 1978–2007.
- Sugier, J., P. Parent du Châtelet, P. Roquain, and A. Smith, 2002: Detection and removal of clutter and anaprop in radar data using a statistical scheme based on echo fluctuation. *Second European Conference on Radar in Meteorology and Hydrology*, Delft, Netherlands.
- Szyrmer, W. and I. Zawadzki, 1999: Modeling of the Melting Layer. Part I: Dynamics and Microphysics. *Journal of the Atmospheric Sciences*, **56**, 3573–3592.
- Tabary, P., 2007: The New French Operational Radar Rainfall Product. Part I: Methodology. *Weather and Forecasting*, **22**, 393–408.
- Tabary, P., J. Desplats, K. Do Khac, F. Eideliman, C. Gueguen, and J.-C. Heinrich, 2007: The New French Operational Radar Rainfall Product. Part II: Validation. *Weather and Forecasting*, **22**, 409–427.
- Tabary, P., A. Le Henaff, G. Vulpiani, J. Parent-du Chatelet, and J. J. Gourley, 2006: Melting layer characterisation and identification with a C-band dual-polarization radar: a long term analysis. *Fourth European Conference on Radar in Meteorology and Hydrology*, Barcelona, Spain.
- Thompson, E., S. Rutledge, B. Dolan, and M. Thurai, 2014: Drop Size Distributions and Radar Observations of Convective and Stratiform Rain over the Equatorial Indian and West Pacific Oceans. *Journal of the Atmospheric Sciences*, **72**, 4091–4125.
- Thompson, R., A. Illingworth, and J. Ovens, 2011: Emission: a simple new technique to

- correct rainfall estimates from attenuation due to both the radome and heavy rainfall. *International Symposium on Weather Radar and Hydrology*, Exeter, UK.
- Thompson, R. J., 2007: Rainfall Estimation Using Polarimetric Weather Radar. Ph.D. thesis, University of Reading.
- Vignal, B., H. Andrieu, and D. Creutin, 1999: Identification of Vertical Profiles of Reflectivity from Volume Scan Radar Data. *Journal of Applied Meteorology*, **38**, 1214–1228.
- Vignal, B., G. Galli, J. Joss, and U. Germann, 2000: Three Methods to Determine Profiles of Reflectivity from Volumetric Radar Data to Correct Precipitation Estimates. *Journal of Applied Meteorology*, **39**, 1715–1726.
- Villarini, G. and W. Krajewski, 2010: Review of the Different Sources of Uncertainty in Single Polarization Radar-Based Estimates of Rainfall. *Surveys in Geophysics*, **31**, 107–129.
- Vulpiani, G., P. Tabary, J. Parent du Chatelet, and F. Marzano, 2008: Comparison of Advanced Radar Polarimetric Techniques for Operational Attenuation Correction at C Band. *Journal of Atmospheric and Oceanic Technology*, **25**, 1118–1135.
- Wen, G., A. Protat, P. May, W. Moran, and M. Dixon, 2016: A Cluster-Based Method for Hydrometeor Classification Using Polarimetric Variables. Part II: Classification. *Journal of Atmospheric and Oceanic Technology*, **33**, 45–60.
- Wood, N., T. L'Ecuyer, A. Heymsfield, and G. Stephens, 2015: Microphysical Constraints on Millimeter-Wavelength Scattering Properties of Snow Particles. *Journal of Applied Meteorology and Climatology*, **54**, 909–931.
- Zawadzki, I., W. Szyrmer, C. Bell, and F. Fabry, 2005: Modeling of the Melting Layer. Part III: The Density Effect. *Journal of the Atmospheric Sciences*, **62**, 3705–3723.
- Zhang, G., J. Vivekanandan, and E. Brandes, 2001: A method for estimating rain rate and drop size distribution from polarimetric radar measurements. *IEEE Transactions on Geoscience and Remote Sensing*, **39**, 830–841.
- Zhang, J., C. Langston, and K. Howard, 2008: Brightband Identification Based on Vertical Profiles of Reflectivity from the WSR-88D. *Journal of Atmospheric and Oceanic Technology*, **25**, 1859–1872.
- Zhang, J. and Y. Qi, 2010: A Real-Time Algorithm for the Correction of Brightband Effects in Radar-Derived QPE. *Journal of Hydrometeorology*, **11**, 1157–1171.

Appendices

Appendix A

UK dual polarisation network parameters

Table A.1 lists the identifiers and locations of the 16 operational UK dual polarisation radars, along with the Met Office research radar at Wardon Hill. The table includes measures of dual polarisation quality in the form of the modal LDR and ρ_{hv} measured in light rain (defined by meteorological reflectivities in the range 20-25 dBZ). These are “system limit” values: the effective minimum LDR and maximum ρ_{hv} measurable at each radar.

	Easting (m)	Northing (m)	Max. ρ_{hv}	Min. LDR (dB)
Predannack	169143	16453	0.9965	-38.3
Wardon Hill	360907	102337	0.9970	-36.0
Hameldon Hill	381030	428750	0.9985	-35.3
Clee Hill	359585	278006	0.9945	-34.8
Cobbacombe Cross	298120	119275	0.9975	-34.5
Chenies	501688	199990	0.9975	-34.1
Dean Hill	424399	125756	0.9975	-34.0
Channel Islands	383800	-80200	0.9975	-34.0
High Moorsley	433875	545519	0.9965	-34.0
Munduff Hill	318742	703248	0.9975	-33.7
Druim-a-Starraig	154460	932393	0.9975	-33.7
Holehead	261880	682795	0.9965	-33.3
Castor Bay	119000	520600	0.9955	-33.0
Hill of Dudwick	397893	837784	0.9965	-32.8
Crugy-y-Gorllwyn	232207	234085	0.9955	-32.5
Thurnham	581670	158290	0.9965	-31.4
Ingham	496027	382949	0.9965	-31.3

Table A.1: Location and dual polarisation quality (LDR and ρ_{hv} system limits) for Met Office / Environment Agency radars. Northing and Easting coordinates are with respect to the UK National Grid.

Appendix B

Definition of statistics

Probability of detection (POD, also known as hit rate HR), false alarm rate (FAR) and Heidke Skill Score (HSS) are defined as follows:

$$\begin{aligned} a &= \text{Hit} ; b = \text{False alarm} \\ c &= \text{Miss} ; d = \text{No detection} \\ \text{POD} &= \frac{a}{a+c} \\ \text{FAR} &= \frac{b}{b+d} \\ \text{HSS} &= \frac{2(ad-bc)}{(a+b)(b+d) + (a+c)(c+d)} \end{aligned}$$

In this thesis correct diagnosis of no bright band by LDR is a “Hit”; false identification a “False alarm”; incorrect default to bright band a “Miss”; and correct default to a bright band profile as “No detection”. POD and FAR range from 0 to 1; HSS from -1 to +1.

Quantitative statistics bias, root mean squared error (RMSE) and root mean squared factor (RMSF) are used to describe the error in radar QPEs (R) with respect to collocated hourly rain gauge accumulations (G):

$$\begin{aligned} \text{Bias} &= \frac{1}{n} \sum_{i=1}^n (R_i - G_i) \\ \text{RMSE} &= \sqrt{\frac{1}{n} \sum_{i=1}^n (R_i - G_i)^2} \\ \text{RMSF} &= \exp \left(\sqrt{\frac{1}{n} \sum_{i=1}^n \left(\log \left(\frac{R_i}{G_i} \right) \right)^2} \right) \end{aligned}$$

Appendix C

LDR calibration factors for the Wardon Hill VPR dataset

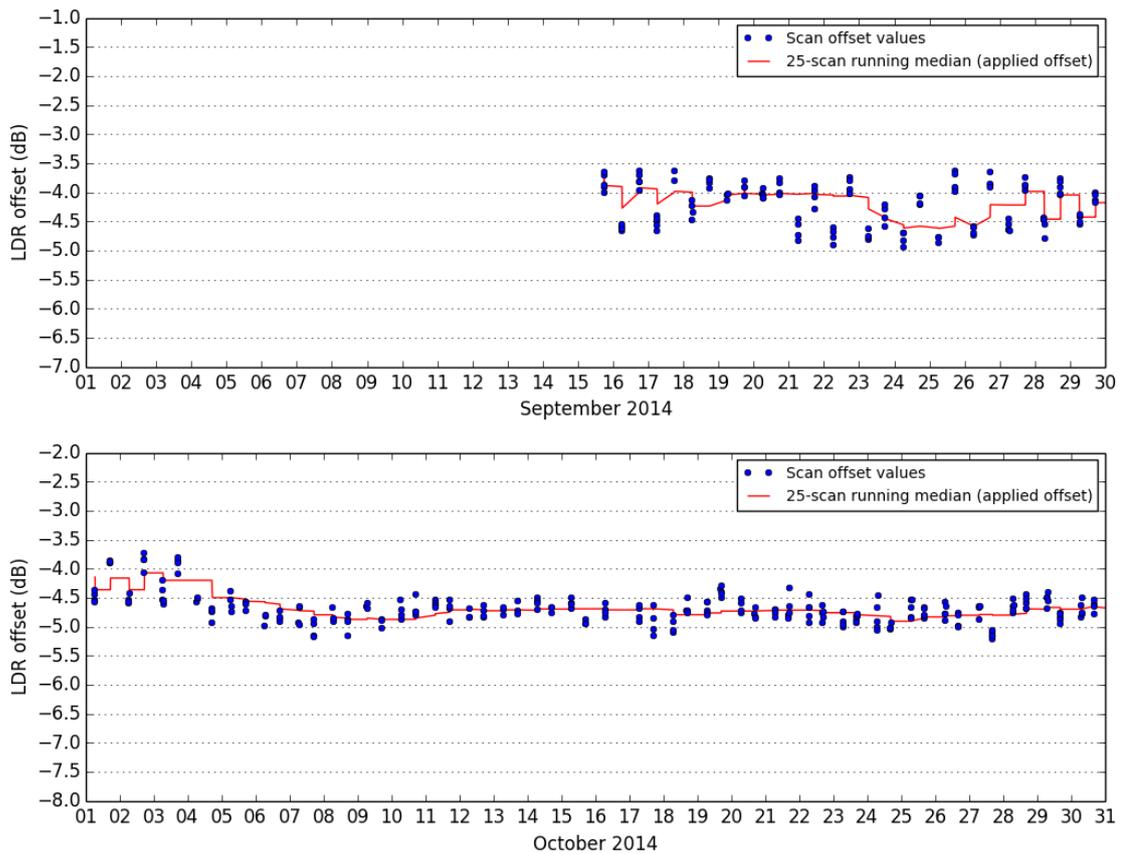


Figure C.1: LDR offset timeseries and applied calibration factors applied to Wardon Hill RHI data (part 1: September-October 2014).

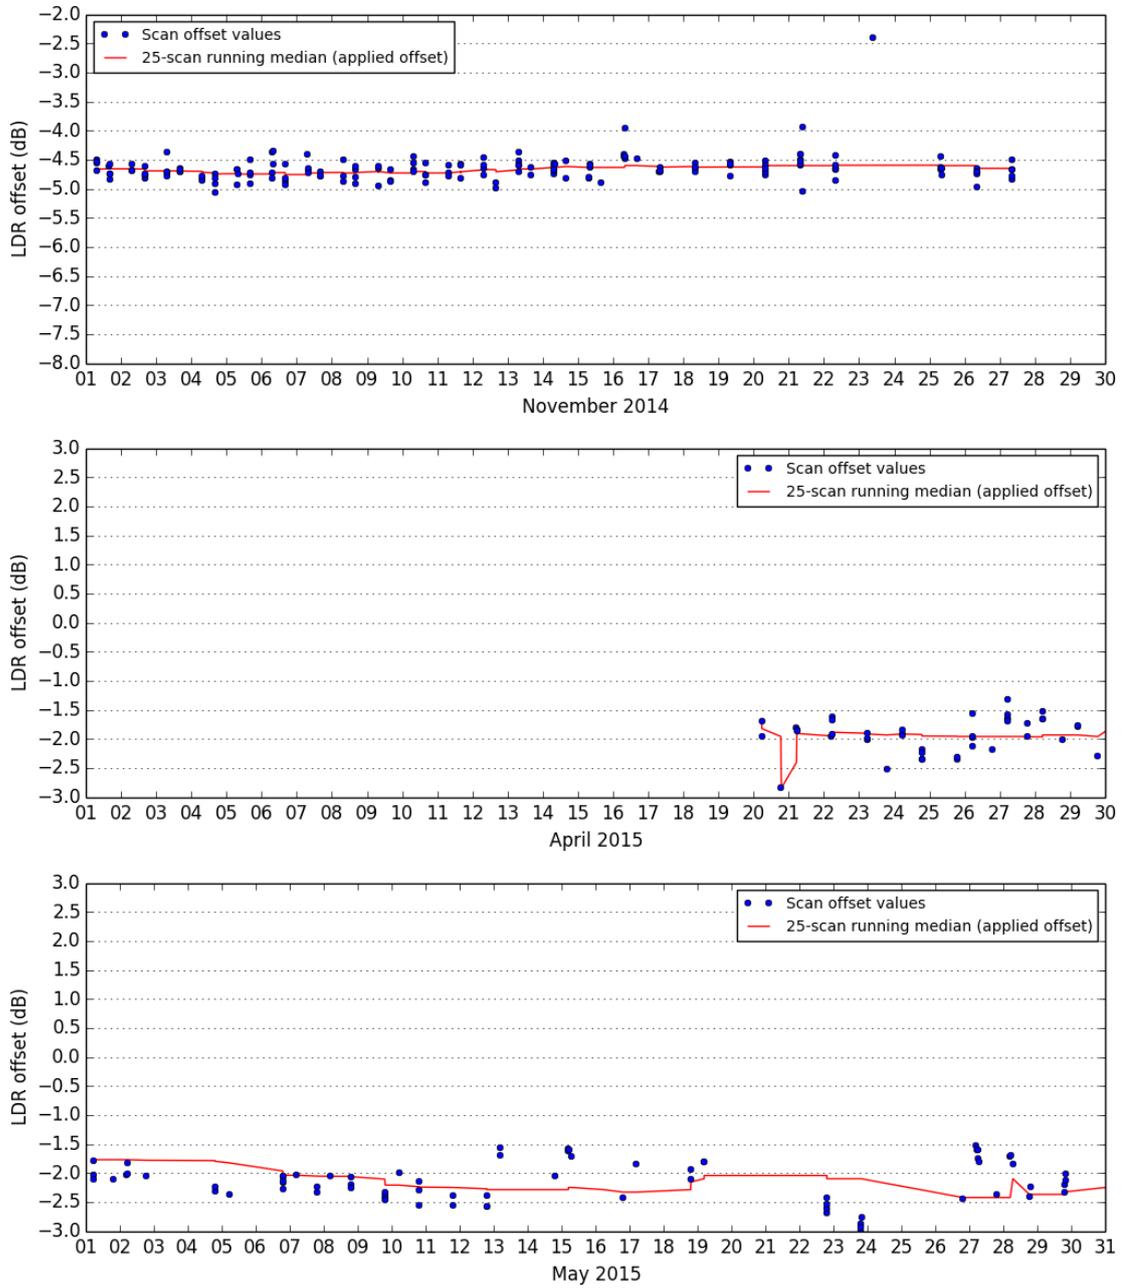


Figure C.2: LDR offset timeseries and applied calibration factors applied to Wardon Hill RHI data (part 2: November 2014; April-May 2015). Note the sparsity of this data is a reflection of the sparsity of useful RHI scans during this time, rather than the stability of the LDR offset calculation, since offsets were only calculated on dates with RHIs contributing to the vertical profile sample.

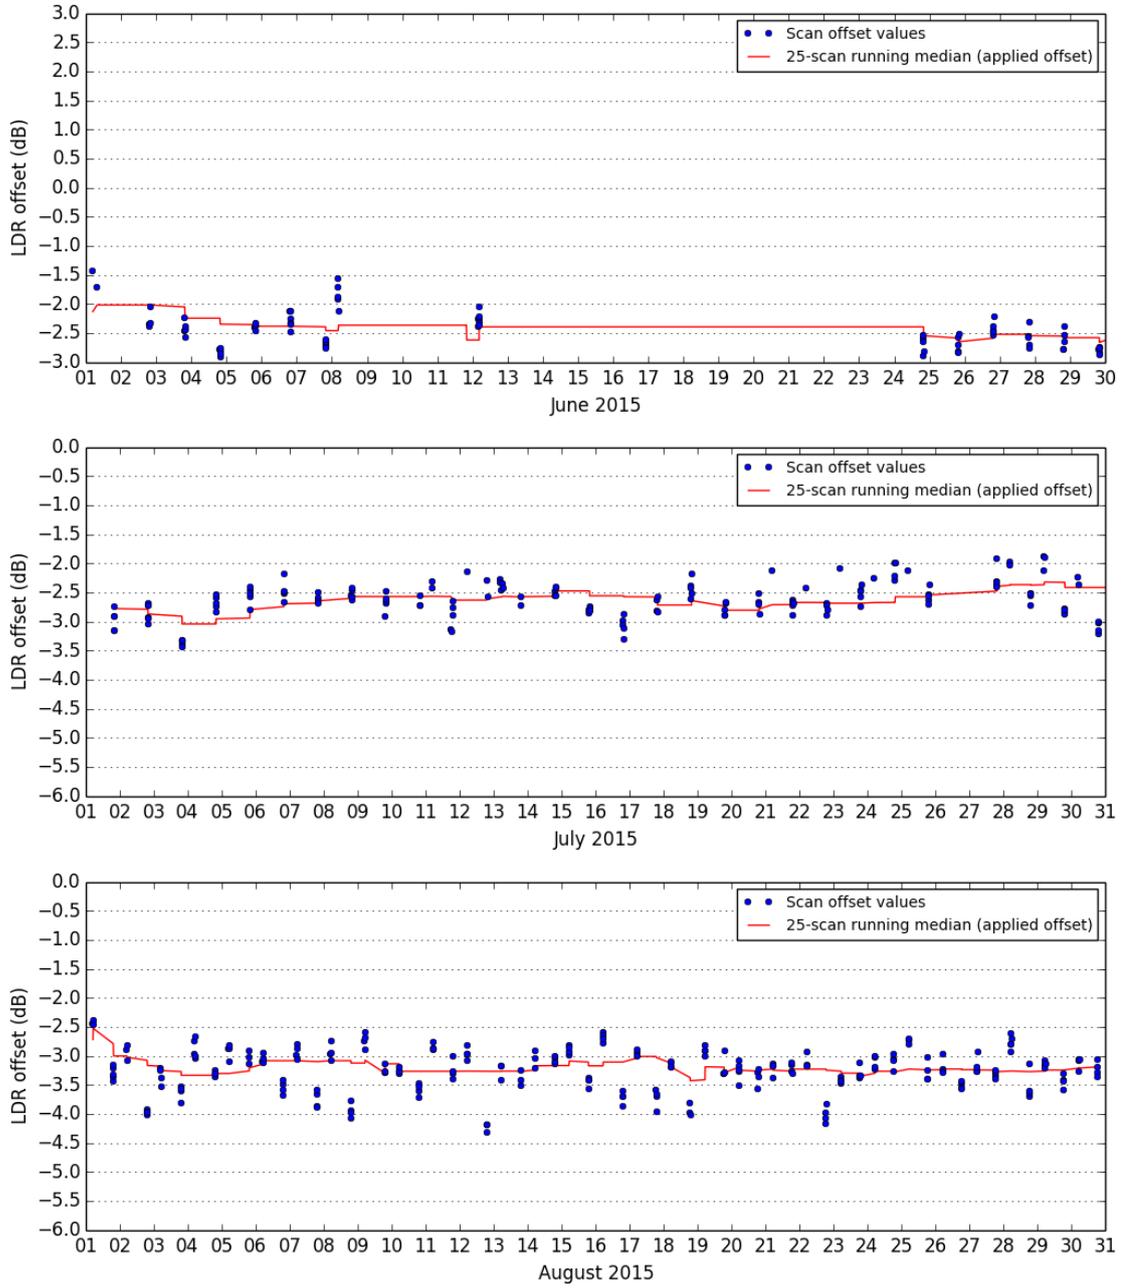


Figure C.3: LDR offset timeseries and applied calibration factors applied to Wardon Hill RHI data (part 3: June-August 2015).

Appendix D

NCIC temperature data for the winter of 2013-14

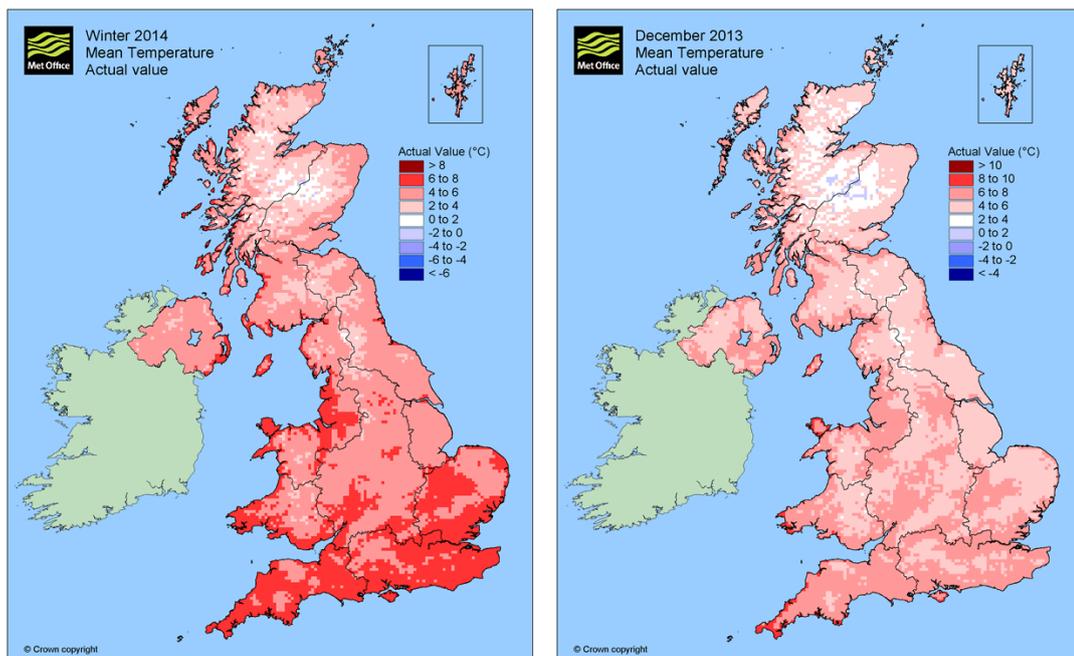


Figure D.1: NCIC monthly temperatures from winter 2013-14 (left) and December 2013 (right). Note scale difference for December map.

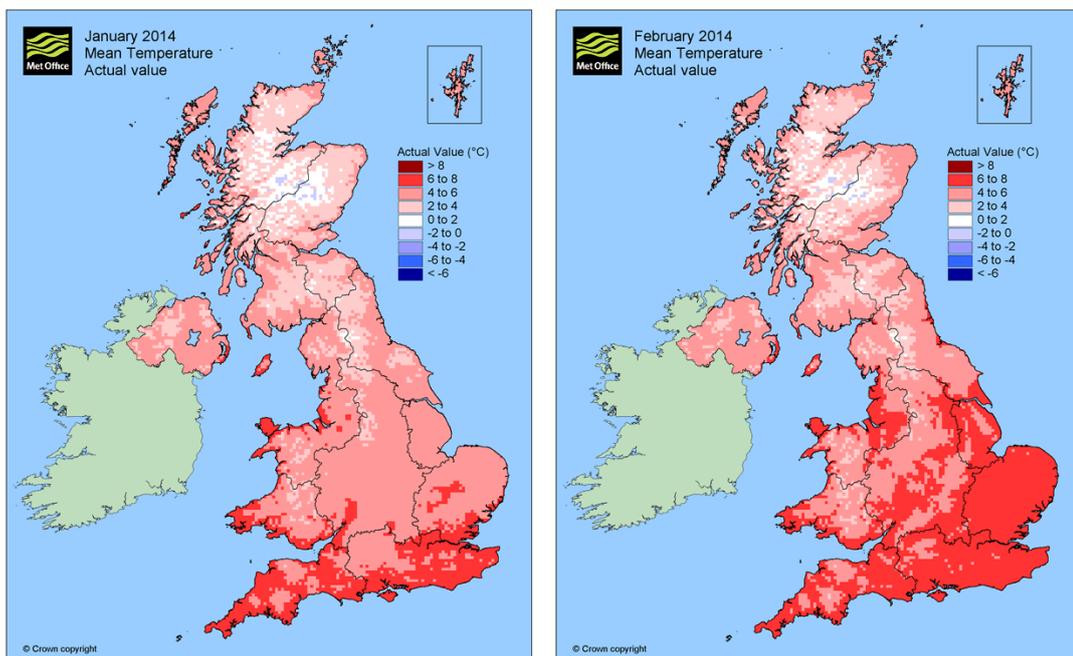


Figure D.2: NCIC monthly temperatures from January and February 2104. Note scale difference from December map.