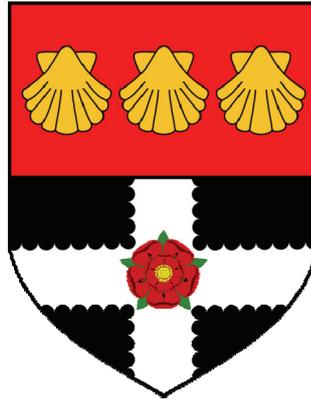


THE UNIVERSITY OF READING



**Profiling Temperature and Humidity by
Ground-based Microwave Radiometers**

Tim J. Hewison

A thesis submitted for the degree of Doctor of Philosophy

Department of Meteorology

September 2006

Abstract

Ground-based microwave radiometry now offers an opportunity to automate upper air observations by providing information on temperature and humidity profiles to Numerical Weather Prediction (NWP) with high time-resolution, but poor vertical resolution. Channels near the 60 GHz oxygen complex are used to sound temperature, while channels near the 22.235 GHz water vapour line are used for humidity, but are also sensitive to the integrated liquid water content.

A one-dimensional variational (1D-VAR) technique is developed to retrieve temperature and total water profiles by combining radiometer and surface observations with NWP background. This provides more accurate *a priori* information than the statistical climatology used by a neural network retrieval technique. The variational retrieval requires specification of a forward model and error covariances of the observations and background. The radiometric noise and calibration stability of a commercial system, the Radiometrics TP/WVP-3000, are analysed.

Case studies show this instrument can resolve basic features of the convective boundary layer and monitor the evolution of near-surface inversions, but not detailed features. They also illustrate the constraints imposed by background data: neural network retrievals being biased when applied outside the training data set's range, and variational retrievals' inability to move poorly forecast inversions.

Temperature profiles retrieved by 1D-VAR have an accuracy of <0.5 K below 1 km and <1.0 K below 4 km, providing ≈ 3 *Degrees of Freedom for Signal* (DFS), which can be improved to ≈ 5 by adding observations at more angles and reducing their noise. However, the humidity profiles' accuracy ($<20\%$ below 1 km, $<40\%$ below 4 km) falls short of the requirements for regional NWP and only provide ≈ 2 or 3 DFS in clear or cloudy conditions, respectively – little more than surface and GPS sensors in clear skies. This shows radiometer observations can improve NWP temperature background to the required accuracy below ~ 3 km.

Declaration

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

© Crown Copyright 2007. Production of this thesis was funded by the Met Office.

"Never trust anything that can think for itself if you can't see where it keeps its brain."

[Rowling, 1999]

"Clouds give you a headache"

– a quote from a child [Limb and Spellman, 2001]

Acknowledgements

A great many people have had a positive impact on this work. I would like to thank everyone who has contributed no matter how indirectly, but the following deserve a special mention.

Firstly, I am also grateful to my manager, John Nash and the Met Office for supporting this research and allowing me to present my work for them as part of my thesis. Other colleagues at the Met Office also deserve special thanks, including those at Camborne and Cardington, who provided excellent support. Catherine Gaffard has offered a wealth of advice on everything from maths to mushrooms and retrievals to renovations. Thanks also to Pete Clark for having the foresight to archive profiles from the mesoscale model over several years, which provided an essential data set used in this study and to Jeremy Price for his immense knowledge of boundary layer meteorology. Other colleagues who abandoned us for the *EI Dorado* of Exeter have also provided support, commitment and advice including Mark Smees, Richard Smout, Tim Oakley, Jonathan Jones, Emma Walker and, especially, John Eyre – I hope I have assimilated a little of his extensive knowledge in his role as supervisor.

I would also like to thank Alan O'Neill in his role as academic supervisor for constantly reminding me to write clearly (hopefully this is more legible than his handwriting!), as well as the other members of my thesis committee for their helpful advice: Ian James, Janet Barlow and John Thuburn. The rest of the department have helped to broaden my knowledge of all things meteorological and provided a wonderful working environment and great entertainment.

This work has also involved a great deal of liaison with the outside world and many people have helped me in industry, conferences and international meetings. The people at Radiometrics Corporation have helped understand and improve the instrument – particularly Mike Exner, Randolph Ware and Fred Solheim. They have generously provided much support and commercially sensitive details of the instrument and its retrievals. My participation in the COST Action 720 generated very useful contacts and good friends, especially Nico Cimini, Christian Mätzler, Dominique Ruffieux, Jürgen Güdner, Susanne Crewell and Ulrich Löhnert.

Finally I am eternally indebted to my loving family for their emotional and technical support. Emily (6) taught me how to use Paint and Jessie (8) taught me all about special fonts in *Word*. (Unfortunately the animated sparkly ones don't print well, so I couldn't use them in my thesis.) Lastly to my wife, Ali, without whose love, support and patience none of this would have been possible: Thank You!

Table of Contents

Chapter 1 Introduction	1
1.1 Motivation	1
1.1.1 Numerical Weather Prediction	1
1.1.2 Nowcasting	2
1.1.3 Other applications	2
1.1.4 The Current Upper Air network.....	3
1.1.5 User Requirements	4
1.2 The Principles of microwave radiometry	6
1.2.1 Why use microwaves?	6
1.2.2 Emission, absorption and scattering mechanisms	6
1.2.3 Radiative Transfer Modelling	9
1.3 Retrieval Techniques	9
1.3.1 Integrated Water Vapour and Liquid Water Retrievals	9
1.3.2 Empirical statistical regression	11
1.3.3 Neural Network	11
1.3.4 Variational Retrievals	11
1.3.5 Integrated Profiling Technique	12
1.3.6 Data Assimilation	12
1.4 Introducing the Radiometrics TP/WVP-3000	14
1.5 The data sets used in this study	14
1.6 Thesis Outline	17
Chapter 2 Radiative Transfer Modelling	18
2.1 Introduction to Radiative Transfer	18
2.1.1 Definition of Brightness Temperature	19
2.1.2 The Radiative Transfer Equation	19
2.1.3 Cosmic Microwave Background	20
2.2 Gaseous absorption models	20
2.2.1 Line broadening	21
2.2.2 Water Vapour	22
2.2.3 Oxygen	22
2.2.4 Other species	22
2.2.5 Example Brightness Temperatures	22
2.2.6 Absorption models available	23
2.2.7 Forward Model Spectroscopic Error	24
2.3 Cloud extinction models	26
2.3.1 Liquid water cloud	26

2.3.2	Supercooled Liquid Water Cloud	27
2.3.3	Ice cloud.....	27
2.4	Precipitation	27
2.4.1	Extinction dependence on drop size and rain rate	27
2.4.2	Drizzle	28
2.4.3	Practical considerations.....	29
2.4.4	Maximum Integrated Liquid Water (ILW).....	29
2.4.5	Data availability due to precipitation	29
2.4.6	Rain Water Content	29
2.4.7	Implications for retrievals	30
2.5	Speeding up Radiative Transfer Calculations	30
2.5.1	Is a fast model needed?	30
2.5.2	Representation of channels' finite bandwidth	31
2.5.3	Fast Absorption Predictor (FAP).....	32
2.5.4	RTTOV (Radiative Transfer for TOVS).....	33
2.6	Profile Representation	34
2.6.1	Radiosonde profiles	34
2.6.2	Mesoscale model levels.....	35
2.6.3	Impact of profile discretisation	36
2.6.4	Atmospheric propagation path.....	37
2.6.5	<i>Topping-up</i> profiles	37
2.6.6	Representing cloud in profiles	37
2.7	Validation of Radiative Transfer Models.....	38
2.7.1	Laboratory measurements for model validation	38
2.7.2	Validation by ground-based radiometers and radiosondes	39
2.7.3	Airborne radiometer/spectrometer and in situ sampling.....	40
2.7.4	Satellite measurements with <i>in situ</i> data	41
2.7.5	Validation with Numerical Weather Prediction analysis.....	41
2.7.6	Choice of clear air absorption model	42
2.8	Infrared Radiative Transfer Model.....	42
2.9	Summary of Radiative Transfer Modelling	44
Chapter 3 Microwave Radiometer Hardware		45
3.1	Radiometer Design.....	45
3.1.1	Microwave Window	46
3.1.2	Antenna System.....	46
3.1.3	Receiver	46
3.1.4	Alternative receiver designs.....	47
3.1.5	Surface sensors	47
3.1.6	Infrared radiometer	47

3.2	Observations	48
3.2.1	Elevation scan.....	48
3.2.2	Observing Cycle.....	49
3.3	Calibration of Microwave Radiometer	50
3.3.1	Noise Diodes.....	50
3.3.2	Internal black body reference target	50
3.3.3	System Equations	51
3.3.4	Tip curve calibrations.....	51
3.3.5	Liquid Nitrogen Calibrations.....	53
3.4	Calibration Bias	54
3.4.1	Modelled calibration bias	55
3.4.2	Validation of observations/model.....	55
3.5	Radiometric Noise	57
3.5.1	Modelled radiometric noise.....	57
3.5.2	Measured radiometric noise	59
3.5.3	Radiometric noise covariance.....	60
3.5.4	Effect of calibration drift	60
3.6	Radio Frequency Interference (RFI).....	62
3.6.1	Interference Mitigation	62
3.7	Future hardware developments.....	63
3.7.1	Simplifying the radiometer for weather station integration	64
3.8	Summary of Radiometer Hardware.....	65
Chapter 4 Variational Retrievals (1D-VAR).....		66
4.1	Theory	67
4.2	Defining State Space.....	68
4.2.1	Choice of state space – level heights and range.....	68
4.2.2	Choice of control variables	68
4.3	Background	70
4.3.1	Mesoscale Model.....	71
4.3.2	Background Error Covariance Matrix	71
4.4	Observation Vector.....	73
4.4.1	Observation Error Covariance	73
4.5	Jacobians	79
4.5.1	Calculation of Jacobians.....	79
4.5.2	Truncating Jacobian calculations at higher levels	81
4.6	Error Analysis	81
4.6.1	Analysis Error Covariance	81
4.6.2	Vertical Resolution.....	83
4.6.3	Validation of B matrices by O-B statistics.....	84

4.7	Cost function minimisation	85
4.7.1	The Linear Case	86
4.7.2	Gauss-Newton Minimisation	86
4.7.3	Levenberg-Marquardt Minimisation	87
4.7.4	Convergence Testing.....	88
4.7.5	Testing for Correct convergence	88
4.8	Pre-processing Classification and Quality Control	89
4.8.1	Cloud Classification	89
4.8.2	Modifying the first guess to linearise problem	90
4.8.3	Rain Detection	90
4.8.4	Additional terms in cost function for physical constraints	91
4.9	Results of 1D-VAR retrieval experiments.....	91
4.9.1	1D-VAR Results with synthetic observations and background	92
4.9.2	Synthetic observations with mesoscale background	95
4.9.3	Results with modified first guess	97
4.9.4	Results with real observations.....	97
4.9.5	Results Dynamic Representativeness.....	99
4.9.6	1D-VAR retrievals of Integrated Water Vapour (IWV)	100
4.9.7	Example retrievals with a misplaced background inversion	100
4.9.8	Summary of 1D-VAR Retrieval Experiment Results.....	102
Chapter 5 Neural Network Retrievals.....		103
5.1	Introduction to Neural Networks	103
5.2	Training Neural Networks	104
5.2.1	State space definition for neural network	104
5.2.2	Training data-sets	105
5.2.3	Training statistics	106
5.3	Neural Network Retrievals from Zenith Observations	108
5.3.1	Bias in neural network retrievals.....	109
5.3.2	Neural network's sensitivity to noise.....	111
5.3.3	Linearity of neural network retrievals.....	112
5.3.4	Application outside range of training data-set	112
5.4	Elevation Scanning in Neural Network Retrievals.....	114
5.4.1	Asymmetry in elevation scans	115
5.4.2	Results of elevation scanning retrievals	116
5.4.3	Technical limitations of elevation scans	117
5.5	Integrated Water Vapour Retrievals by Neural Network	118
5.5.1	Atmospheric variability of IWV	120
5.5.2	Comparing IWV from radiometer and GPS in precipitation.....	121
5.6	Integrated Liquid Water retrievals by Neural Network.....	122

5.6.1	Validation of ILW retrievals and NWP models	122
5.6.2	Problems with ILW retrievals	123
5.6.3	Linearity of cloud fraction averaging	124
5.7	Summary of Performance of Neural Network Retrievals	125
5.7.1	Pros and Cons of neural network retrievals.....	126
5.7.2	Future Improvements to Neural Network.....	126
Chapter 6 Case Studies		128
6.1	Retrievals in Convective Conditions.....	128
6.1.1	Example of Convection – Cardington 8 April 2003.....	128
6.1.2	Static stability	130
6.1.3	Resolving thermals	132
6.2	Retrievals of cloud and fog	134
6.3	Ability to retrieve temperature inversions	136
6.4	Average layer temperatures	138
6.5	Performance in rain	139
6.6	Comparison with model forecast profiles	140
6.6.1	Clear Sky CSIP case – 12 July 2005.....	141
6.6.2	Cumulus capped boundary layer CSIP case – 13 July 2005	143
6.7	Summary of Case Studies.....	145
Chapter 7 Optimising Observations for Operations.....		146
7.1	Information Content – Degrees of Freedom for Signal	146
7.1.1	Instantaneous Observations at Zenith.....	146
7.1.2	Averaging Observations to Reduce Radiometric Noise	147
7.1.3	Extra information from elevation scanning	148
7.1.4	Channel selection	150
7.1.5	Trade-off channel and angle selection	152
7.1.6	Recommended observing sequence	155
7.2	Integrating data from other instruments	155
7.2.1	Ceilometer.....	156
7.2.2	Cloud Radar.....	157
7.2.3	Integrated Water Vapour column from GPS.....	157
7.2.4	Wind profiling radar signal	157
7.3	Applications	158
7.3.1	Research applications	158
7.3.2	Demonstrating impact on NWP	159
7.3.3	Other retrieval techniques.....	160
7.3.4	Filling-in observing networks	161
7.3.5	Real-time Retrieval Visualisation	162

7.4	Exploiting High Time-Resolution Observations.....	162
7.5	Summary of Optimisations for Operational Use.....	164
Chapter 8 Conclusions		165
8.1	Hardware.....	165
8.2	Retrieval Techniques.....	166
8.3	Benefits and Limitations	167
8.4	Information Content of Observations.....	168
8.5	Towards Operational Use.....	168
Acronyms and Abbreviations.....		170
Symbols.....		171
References		172

Chapter 1 Introduction

Can ground-based microwave radiometers provide useful data on atmospheric temperature and humidity profiles?

That is the basic question asked in this thesis. To begin to answer it, this introduction will first consider the requirements for temperature and humidity profiles for different applications and how these are currently being addressed. The microwave radiometer will then be introduced, along with the theoretical basis for deriving profiles from its observations. This thesis assesses the accuracy of the radiometer's observations and the profiles retrieved from them.

1.1 Motivation

Knowledge of the vertical profile of temperature, humidity and cloud is key in determining the basic thermodynamic state of the atmosphere, which, in turn, is key to producing accurate weather forecasts. Together these variables define the static stability and convective inhibition introduced by temperature inversions, which determine the vertical transport of heat, moisture and momentum between the surface, boundary layer and free troposphere. These processes are of primary interest on meso- and convective- scales (2-2000 km and <2 km, respectively). Observed profiles are used in a number of ways in weather forecasting, both in Numerical Weather Prediction and nowcasting applications.

1.1.1 Numerical Weather Prediction

Numerical Weather Prediction (NWP) is used to forecast weather, based on a 4-dimensional computer model of the state of the atmosphere. The model represents atmospheric variables on a discrete grid and applies a set of equations to determine how these evolve over discrete time steps. NWP models often have biases, which require human intervention and interpretation to remove these from forecasts prior to dissemination. Improved horizontal resolution allows smaller scale atmospheric process to be resolved and better representation of orographic influences, although they can only accurately represent the propagation of features with wavelengths 3-4 times their grid spacing. However, it requires shorter steps and much greater computer time, which imposes a practical restriction on the models' domain size and, consequently, the time range of its forecasts.

The Met Office run a series of versions of its *Unified Model* (UM) routinely for different applications. The *Global* version is run every 6 hr with a horizontal resolution of ~40 km and forecast out to 5 days (T+120 hr) every 12 hr. The *Global* model provides boundary conditions for regional versions, such as the *North Atlantic-European* mesoscale model, which covers a limited area at higher resolution (~12 km), but to a shorter range (T+36 hr). This study uses the *Mesoscale* version of the UM as it ran operationally in 2002-2005. This covered an area of ~1740 x 2172 km around the British Isles at 12 km resolution and was run every 6 hr to

forecast to T+36. (It has since been superseded by a 4 km version of the model.) Versions of the UM with coarser resolution are run less frequently for longer-range and seasonal forecasts and climate modelling. All versions used 38 vertical levels during this period.

NWP models are initialised by constructing a best estimate of the true state of the atmosphere to fit all available observations, known as an *analysis*. Modern data assimilation schemes continuously adjust NWP models as a compromise between the *a priori* model state and the observations to allow asynchronous observations to be exploited optimally. These techniques are described in more detail below.

Ground-based microwave radiometers are not currently used operationally in NWP – their potential is investigated in this thesis.

1.1.2 Nowcasting

The application of NWP is limited at short time ranges because of the finite time needed to initiate the model. On time scales less than ~6 hr, features in high resolution observations can be exploited by *Nowcasting* techniques. These typically extrapolate the current situation based on the expected trend over the next few hours. For example, frontal rainbands may be advected based on their motion observed by weather radars and model forecast winds. Nonlinear processes are more complex – forecasters may monitor the development of profiles to adjust the timing of predictions of convection onset or fog clearances.

Troitsky *et al.* [1993] illustrated a simple example of a nowcasting application of ground-based microwave radiometer: by monitoring the temperature profiles retrieved from observations at high time-resolution, the user is able to issue warnings when inversions are developing to assess the risk of pollution being trapped at the surface because of inefficient dispersion.

1.1.3 Other applications

Observed profiles are also used in atmospheric research to improve the understanding of atmospheric processes. In turn, these studies can improve the representation of these processes in NWP models – especially the parametrization of sub-gridscale processes. For example, time series of temperature profiles retrieved from a ground-based microwave radiometer have been used to study the influence of an urban heat island on the boundary layer by Golitsyn *et al.* [2002]. They found nocturnal inversions formed less frequently in the centre of Moscow, were weaker, formed much later and dissipated earlier than in nearby rural areas. The boundary layer was usually found to be warmer up to 500-600 m in the city centre than the suburbs, although occasionally a relatively cool layer formed here above 300 m.

Longer time series of observations can be used to monitor and understand climate change.

1.1.4 The Current Upper Air network

Microwave radiometers have the potential to provide data on temperature and humidity profiles for the above applications, as part of a network of upper air observations. The components of the current operational network are reviewed below. The desire to reduce the ongoing costs and manpower requirements of the upper air network makes it attractive to investigate the potential of remote sensing systems that can be automated and require little manual intervention or consumable items.

Historically, radiosondes have formed the backbone of the world's upper air observations. Typically launched 1-4 times a day, they are capable of measuring profiles of pressure, temperature, humidity and horizontal wind vector from the surface to 20-30 km at high vertical resolution (~10 m). However, there are considerable costs involved with maintaining the radiosonde network – both in consumables and manpower to launch and quality control their data. To reduce these costs, the operational radiosonde network in the UK was recently reduced to 6 stations, launching balloons twice daily. This includes four *autosonde* stations, which only require restocking every 24 launches, but still require some manual quality control. This network represents an average spacing between radiosondes of ~250 km (based on the UK land area of 244,820 km²), with an average interval between soundings of 12 hr.

The radiosonde network is supplemented by 6 wind profiling radars, which currently provide profiles of horizontal wind vectors. The potential to exploit their information on the humidity structure is an active area of current research (see §7.2.4).

Commercial aircraft of participating fleets now also provide temperature profiles on ascent or decent into certain airports. Their distribution is limited in space, and particularly, time, due to limitations on night flying. Currently typical horizontal sampling is comparable to the current UK radiosonde network, but with more frequent profiles available during daytime (~6/day) albeit with slightly lower accuracy and vertical resolution [Stringer, 2006]. Although the current E-AMDAR systems do not include humidity sensors, these are under development [Petersen and Moninger, 2006].

Satellite sounders have come to be one of the most important sources of upper air observations for NWP [Böttger *et al.*, 2004]. Although polar-orbiting satellites can provide global coverage, each satellite typically only samples any one location twice a day. There are currently 3 operational polar-orbiting satellites carrying sounding instruments to increase the coverage to ~6/day. Instruments on geostationary satellites could provide much more frequent observations of a more limited area, but with a more limited horizontal resolution, which further degrades at high latitudes. Although nadir-viewing microwave sounding instruments on polar-orbiting satellites can provide information on the vertical profiles of temperature, humidity and other species, their vertical resolution is typically rather coarse

(~3 km) and their application over land is limited by its high and variable emissivity [Hewison, 1999], the large uncertainty in the land surface temperature and limited resolution in the boundary layer. The new generation of infrared sounders have thousands of channels and offer improved vertical resolution (~1 km), but are limited to cloud-free conditions, or application above the cloud tops. Application of data from satellite sounding instruments in the boundary layer can be further complicated by the surface temperature and emissivity.

Ground-based remote-sensing systems can complement their satellite-based counterparts. Microwave radiometers' data can be used to retrieve profiles of temperature and humidity in clear or cloudy conditions at high temporal resolution, but only with moderate resolution in the boundary layer, which degrades further at higher levels. Microwave radiometers are one of the technologies being considered for deployment as part of a proposed network of *Integrated Profiling Systems*, comprising different ground-based remote-sensing instruments operating at the same site [Monna *et al.*, 2006]. In this concept, observations from different instruments are combined in a common retrieval which exploits the benefits of the various observation types to provide a best estimate of the true atmospheric state at the site. It is hoped that this will improve the vertical resolution of the profiles that can be retrieved by microwave radiometry alone. These retrievals are then available for assimilation into NWP models or the data could be assimilated directly. Their observations are expected to have most impact on the highest resolution NWP models because information concentrated in the boundary layer. These models have the shortest range (6-36 hr) because of the limit imposed by their boundary conditions.

1.1.5 User Requirements

The stated requirements of all the Met Office users of upper air observations were reviewed by Stringer [2006]. These are expressed in terms of the accuracy, vertical and horizontal resolution, observing cycle and delay required of temperature in the boundary layer and lower troposphere, humidity in the lower troposphere and column integrated water vapour for climate monitoring, global and regional NWP, synoptic, aeronautical and nowcasting applications. The requirements are defined in terms of the *minimum* threshold for observations to have any impact on each application, the *breakthrough* threshold at which the observations could provide a significant advance in forecast capability (relative to that currently available), and the *maximum* threshold, above which no significant benefit will be felt. All observing systems have strengths and weaknesses – none meet the breakthrough levels for all aspects (accuracy, vertical and horizontal resolution, observation cycle and delay). The best that can be expected is to achieve this level of performance from a combination of systems. The user requirements are shown in Table 1-1 for regional NWP. It can be seen that this method of specifying user requirements for observations gives a very broad range between *minimum* and *maximum* thresholds, which makes it difficult to compare the benefits of different observing systems.

Table 1-1 User requirements of temperature and humidity profiles for Regional NWP – minimum, breakthrough and maximum thresholds

NWP Regional	Temperature (K) Boundary Layer			Temperature (K) Lower Troposphere			Relative Humidity (%) Lower Troposphere			Integrated Water Vapour (kg/m ²)		
	Min.	Brk.	Max.	Min.	Brk.	Max.	Min.	Brk.	Max.	Min.	Brk.	Max.
Accuracy	1.5		0.5	1.5		0.5	10		5	5		1
Vertical Resolution (km)	0.5	0.3	0.01	2	1	0.1	2	1	0.1	N/A	N/A	N/A
Horizontal Resolution (km)	50	10	1	200	30	3	200	30	3	100		10
Observing Cycle (hr)	3	1	0.166	12	3	0.5	12	3	0.5	1		0.5
Delay in Availability (hr)	3		0.083	5		0.25	5		0.25	0.5		0.1

Regional NWP needs observations with better accuracy, vertical and horizontal resolution than global NWP. It is here that observations from ground-based microwave radiometers are expected to have most impact, because their information is concentrated in the boundary layer, which often has high temporal variability that may be well captured by their high frequency observations. Better temperature accuracy, vertical and horizontal resolution is required in the boundary layer than the rest of the lower troposphere due to its greater variability in space and time.

However, if the user requirement for the minimum horizontal resolution of boundary layer temperature profiles is taken at face value, then ground-based microwave radiometers will have no impact in regional NWP, unless they can be deployed in a dense network of ~100 in the UK. It would be prohibitively expensive to deploy a network of ground-based instruments that essentially take spot measurements (i.e. do not cover a significant area) capable of exceeding the stated *minimum* threshold for horizontal resolution of regional NWP. This thesis will investigate whether the high time-resolution of their observations can compensate for this deficiency.

Global climate monitoring has similar requirements to regional NWP for accuracy and vertical resolution of temperature and humidity profiles, although the horizontal resolution, observing cycle and delay in availability are much less stringent. However, these requirements do not distinguish between the accuracy of the observations due to short-term noise and long-term drift, which become critical for climate monitoring. As will be shown, the latter can be

significant for microwave radiometers. Climate monitoring also requires a network of very stable instruments sampling in the same configuration over periods >10 years to detect small changes [IPCC, 2001]. The Hadley Centre recommend more stringent criteria for inter-decadal biases with minimum and maximum mean value accuracies of 0.3 °C and 0.1 °C for temperature and 5% and 2% for humidity [Stringer, 2006]. As a new observing system and one for which the absolute calibration is not sufficiently repeatable for this purpose, it is unlikely that data from ground-based microwave radiometers will contribute significantly to the climate monitoring requirements for some time.

1.2 The Principles of Microwave Radiometry

Absorption features of atmospheric molecular species can be exploited to retrieve profiles of their concentration and temperature. Rotational transitions of molecular gases generate spectral lines at microwave frequencies – the dominant species being oxygen and water vapour. Oxygen, being a well-mixed gas, can be used to profile the atmospheric temperature structure. Once this is established, emission from water vapour lines can be used to profile humidity.

1.2.1 Why use microwaves?

As will be shown, microwaves are able to penetrate through non-precipitating clouds with small drop sizes relative to the wavelength. This is a big advantage for profiling over visible and infrared wavelengths, which are strongly absorbed by cloud, which generates a discontinuity in their response to the humidity changes as saturation is reached. Microwave radiometers can be made sufficiently sensitive for atmospheric profiling with channel bandwidths less than the width of absorption features. Each channel is sensitive to a narrower range of opacities than the relatively broad channels of an infrared radiometer. This can improve the vertical resolution of the profiles retrieved from microwave radiometers [Houghton *et al.*, 1984]. Additionally, microwave radiometers are relatively robust instruments, capable of continuous unattended operation for long periods with minimal maintenance – making them suitable for use as part of an automated system. Their measurements are not seriously affected by dust or solar radiation, but are sensitive to liquid water, which can degrade the accuracy of the measurements in rain.

1.2.2 Emission, absorption and scattering mechanisms

Figure 1-1 shows the absorption due to oxygen, water vapour and liquid water cloud in typical conditions near the surface. The oxygen spectrum includes 33 spin-rotational lines between 51.5-67.9 GHz. In the lower troposphere pressure broadening merges these into a single band of strong absorption, referred to as the *60 GHz oxygen complex*. Absorption is very strong near the centre of this band due to the strength and density of the spectral lines. A channel of a ground-based radiometer detecting down-welling radiation near the centre of this

band (~59 GHz) experiences very strong absorption, so the atmosphere has a 1/e penetration depth of only ~300 m [Troitsky *et al.*, 1993]. Frequencies further from the centre experience less absorption, such that a channel at ~51 GHz sees atmospheric emissions integrated over ~10 km. The differential absorption between different channels in this band provides the principle for retrieving temperature profiles.

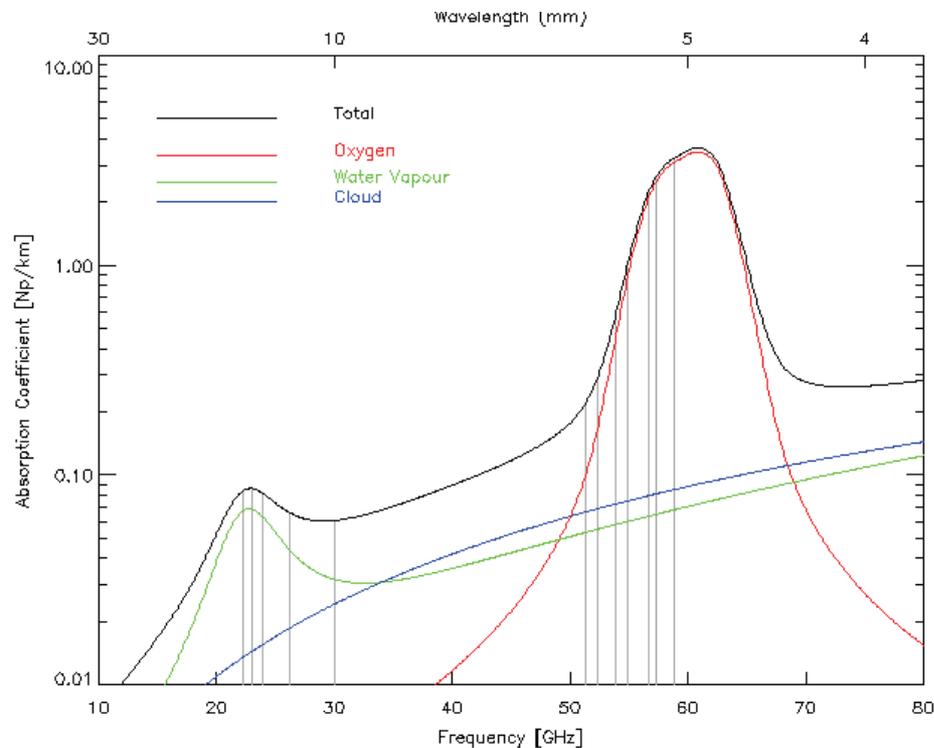


Figure 1-1 Atmospheric absorption spectrum for typical surface conditions: $T=288.15$ K, $p=1013.25$ hPa, $RH=100\%$, $L=0.2$ g/m³ following Rosenkranz [1998]. Lines show total absorption coefficient and contribution from oxygen, water vapour and cloud, coloured according to the legend. Grey vertical bars indicate the centre frequencies of the Radiometrics TP/WVP-3000 microwave radiometer.

In this part of the spectrum water vapour has a single line at 22.235 GHz due to rotational transition of its electric dipole within a single vibrational state. At higher frequencies absorption increases due to the non-resonant *water vapour continuum*. These are relatively weak features, so the atmosphere is optically thin, even at its centre – the down-welling radiance at the surface is a function of the emissions integrated over the whole atmosphere, which in turn depend on the temperature and partial pressure of water vapour at each level. However, channels further from the line centre experience weaker absorption and will be less sensitive to water vapour near the surface. This principle is used to retrieve humidity profiles, once the temperature profile has been established. However, the validity of its assumptions is questioned later in this thesis.

Absorption due to liquid water also follows a non-resonant continuum. Depending on the liquid water content, its emission can dominate the signal in the 30-50 GHz *window* region, whilst still remaining optically thin. So channels here are sensitive to the integrated amount of liquid water, but provide little information on its vertical distribution.

Radiation propagating through a medium can be scattered by inhomogeneities if they have dielectric properties dissimilar to that of the medium. For liquid water, absorption is relatively strong, so scattering only becomes significant as particle sizes approach the Rayleigh scattering criterion, $2\pi r \ll \lambda$. Figure 1-2 shows how the scattering efficiency of water droplets depends on the size of the particles and the frequency of radiation. At the microwave frequencies used in this study, this only becomes significant for airborne hydrometeors larger than drizzle drops ($r >> 0.1$ mm), such as rain drops, snow and hail. Haze and cloud droplets are smaller than this, so scattering can be neglected as it is insignificant compared to absorption.

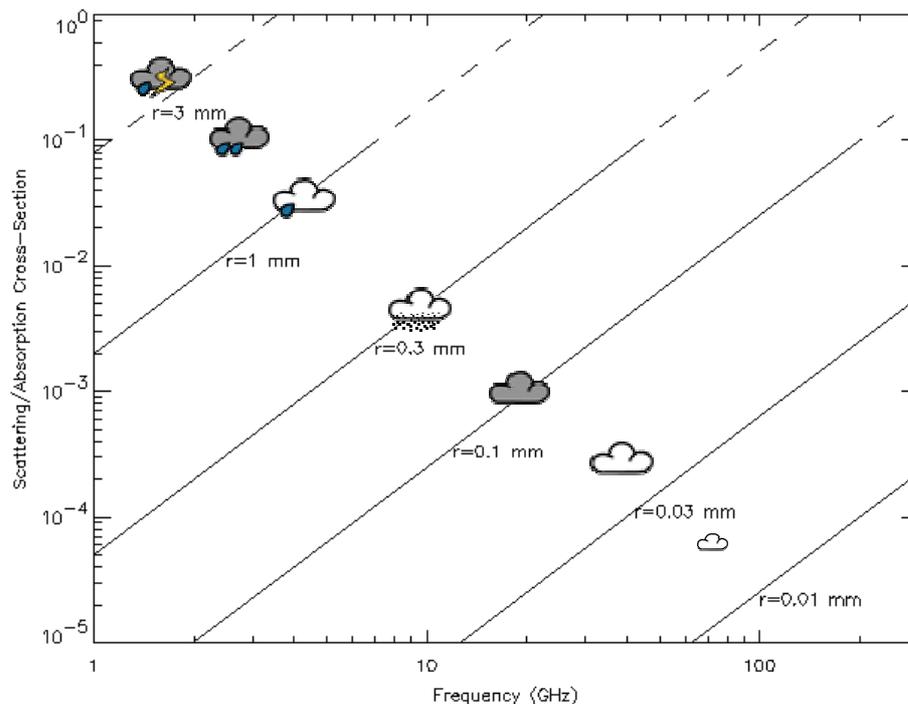


Figure 1-2 Ratio of scattering to absorption cross-sections for water spheres radius, r , at microwave frequencies. Typical drop sizes are indicated ranging from small and large cumulus clouds, drizzle, light rain, heavy and torrential rain. Solid lines indicate where Rayleigh scattering criterion is valid $2\pi r \ll \lambda$. Dashed lines indicate the transition to the Mie scattering regime.

Adapted from Janssen [1993].

1.2.3 Radiative Transfer Modelling

The radiative transfer of microwaves through clear air is modelled by integrating absorption due to each spectral line for both oxygen and water vapour. This is referred to as a *line-by-line* model. Extra absorption is then added to account for the water vapour continuum and liquid water, which are modelled empirically. Scattering is neglected as it is only significant for large particles of liquid water or ice associated with moderate to heavy precipitation, when the accuracy of radiometer observations is likely to be compromised. Down-welling radiances at discrete frequencies are propagated along the atmospheric path by combining the emission from each atmospheric layer with radiation from above after accounting for its absorption within the layer.

The radiative transfer model forms part of the *forward model*, used to translate from *state space* of the profile to *observation space* of the radiometer. Other components of the forward model include the discretisation of the atmospheric profile to a finite number of levels and the partitioning of the variables of state space into the control variables used in the retrievals – for example profiles of temperature, absolute humidity and liquid water averaged over atmospheric layers. Forward models are needed for all *physical* retrieval techniques.

1.3 Retrieval Techniques

Observations at a series of frequencies along an atmospheric absorption feature or a range of elevation angles can be used to retrieve profiles of temperature and humidity (or the partial pressure of the absorbing species). The principle of profiling by frequency scanning is differential absorption, described above. Views of the atmosphere at different elevation angles can also be used to profile temperature in the boundary layer – even with a single channel near the centre of the strongly absorbing 60 GHz oxygen complex [Troitsky *et al.*, 1993]. This relies on the increase of absorption *per vertical kilometre* with zenith angle [§5.4].

However, the retrieval of profiles from radiometer observations is an *ill-posed* problem: no unique solution exists that matches a given set of observations within their errors. This can lead to extreme sensitivity to low levels of observation noise. The basic solution to this problem is to use extra information in the retrieval. This *a priori* information can take the form of a statistical climatology (e.g. of radiosonde profiles), a short-range forecast from an NWP model or other observations, such as a nearby or recent radiosonde profile.

1.3.1 Integrated Water Vapour and Liquid Water Retrievals

For several years microwave radiometers have been used to retrieve the column Integrated Liquid Water (ILW) and Water Vapour (IWV) [e.g. English *et al.*, 1995, Crewell *et al.*, 2001,

Martin *et al.*, 2006]. These terms are also referred to as Liquid Water Path (LWP) and Total Water Vapour, Total Precipitable Water or Precipitable Water Vapour in the literature.

IWV and ILW can be retrieved using only 2 channels: one near the 22.235 GHz line and one in the window region (30-50 GHz) where the atmospheric absorption is dominated by water vapour and liquid water, respectively, as shown in Figure 1-1. Although these frequencies are not sensitive to emission from other species, they have a weak dependence on the temperature of the water vapour/liquid water, which, in turn, depends on their vertical distribution. For this reason, the retrieval of IWV/ILW is still under-constrained and requires *a priori* information to resolve this potential ambiguity. The accuracy of the retrieval depends on the accuracy of the *a priori* temperature profile – so rapid changes can bias the retrieval slightly. IWV and ILW can be retrieved by a simple linear regression of the opacities at these frequencies, which can be calculated from the radiances observed by a 2 channel microwave radiometer [Gaussiat *et al.*, 2006]:

$$\begin{aligned} IWV &= v_1 \tau_1^* + v_2 \tau_2^* \\ ILW &= l_1 \tau_1^* + l_2 \tau_2^* \end{aligned} \quad (1.1)$$

where τ_1^* and τ_2^* are the opacities in each channel after subtracting the dry contribution, τ_d ,
 v_1, v_2, l_1 and l_2 are constants, functions of the path averaged absorption coefficients.

The coefficients, v_1, v_2, l_1 and l_2 , can be determined by linear regression of IWV and ILW calculated from a large ensemble of profiles and the opacities predicted for these by a forward model. In this method the data set of profiles provide the *a priori* information, which determines the vertical distribution of water vapour and liquid water statistically. (However, if this is determined from radiosonde profiles, additional algorithms are needed to generate the liquid water distribution synthetically based on the humidity profile.) It would also be possible to use other *a priori* information, such as NWP model data. However, the *training* data set used to calculate these coefficients must be representative of the conditions at the site of the retrievals to ensure the retrievals are not biased. Despite these ambiguities, microwave radiometry is regarded as the most accurate method of observing IWV and ILW [Crewell *et al.*, 2002], at least in cloudy conditions.

More sophisticated methods to retrieve IWV and ILW from radiometer data have been developed to exploit other *a priori* data, extra microwave channels [Crewell and Löhnert, 2003] and observations from other instruments [Löhnert *et al.*, 2004].

1.3.2 Empirical statistical regression

Although the retrieval of temperature and humidity profiles is nonlinear, it can also be implemented as an empirical statistical regression between sets of observations and profiles. This technique was demonstrated by Güldner and Spänkuch [2001] using a data set of 18 months' observations from a Radiometrics TP/WVP-3000 and co-located radiosondes. While the application of this method is limited to those few sites where long time-series of co-located radiometer and radiosonde observations are available, it has the advantages of not requiring a forward model and implicitly accounting for any calibration biases in the radiometer (assuming they are constant). The coefficients derived for one location cannot readily be applied to other sites, even if they are at the same height (so pressure broadening is the same), because the *a priori* information used in the retrieval will not be representative of the climate for other sites [Cimini *et al.*, 2006b]. This will bias the retrievals based on them, as they are essentially weighted averages of the observations and the background.

1.3.3 Neural Network

A nonlinear statistical regression can also be implemented as an artificial neural network. Solheim *et al.* [1996] developed this method to retrieve profiles of temperature, humidity and liquid water content from observations of the Radiometrics TP/WVP-3000 microwave radiometer, based on a training data set of thousands of radiosonde profiles covering 10 years (with synthetic cloud added), which were input into a forward model to give synthetic observations. Being based on a nonlinear statistical regression makes neural networks prone to large errors if applied outside the range of their training data set.

1.3.4 Variational Retrievals

Just as NWP models are expected to provide short-range forecasts that are more accurate than climatological averages, they can provide a more accurate background than climatological data sets. These can be used in variational retrievals following a Bayesian framework. This thesis develops a One-Dimensional Variational (1D-VAR) retrieval by minimising a cost function to calculate a best fit between the radiometer observations and this background projected into observation-space. This minimisation requires the calculation of the derivatives of the forward model, which is computationally slow. Furthermore, because the retrieval of temperature and (particularly) humidity profiles is nonlinear, the minimisation requires iterative calculation, following variations of Newton's method. These variational retrievals are *optimal*, in that they combine the observations and background, each weighted according to their expected error characteristics. They can also be readily extended to include observations from other instruments to investigate the potential benefit of the *integrated profiling* approach. However, they require considerable development and careful design in the choice of control variables, forward model operators, background and observation error

covariances, minimisation techniques, convergence testing and validation method. These aspects are addressed in Chapter 4 of this thesis.

1.3.5 Integrated Profiling Technique

During the course of this study, Löhnert *et al.* [2004] independently developed an optimal estimation method to retrieve profiles from ground-based remote sensing observations, similar to the variational method described above. Known as the *Integrated Profiling Technique* (IPT), this was developed to improve estimates of liquid water content by combining radiometer observations with those from a laser ceilometer and cloud radar, taking background information from a nearby radiosonde profile. Because the radiosondes do not measure liquid water, this was synthesised in the background and its error characteristic estimated using a high-resolution cloud resolving model. The IPT used profiles of temperature, absolute humidity and the logarithm of the liquid water content as control variables. (The logarithm is used to prevent the retrieval of unphysical negative cloud.) Löhnert *et al.* [2004] found the IPT retrieved temperature and humidity profiles with an r.m.s. accuracy in the lowest 3 km of 0.63 K and 0.64 g/m³, respectively, compared to 1.0 K and 1.11 g/m³ for the background available from a nearby radiosonde – on average 6 hr and 30 km from the remote sensing observations.

1.3.6 Data Assimilation

Data assimilation describes the process of how the observations are used in NWP models. Observations are compared to the background field in observation-space, interpolated to the observation point in space and time. This allows gross errors to be rejected and small biases to be corrected. The differences, known as *observation increments*, are analysed to generate a series of corrections to the previous short-range forecast, which are added to it to initiate the next forecast run. This analysis uses statistical characteristics of the background and observation errors to minimise a cost function describing their difference in an iterative process. The basic assumption of data assimilation is that short-range NWP provides good forecasts and that small corrections to these are applied based on their difference from a series of observations. This process forms a continuous cycle at each model run and aims to improve the accuracy of its forecasts.

Three-dimensional variational assimilation (3D-VAR) is currently used in the assimilation of observations into the mesoscale version of the Met Office's Unified Model. This can be thought of as an extension of the 1D-VAR retrieval described above, but instead of retrieving a single column profile, the state vector comprises a full 3-D representation of the model state. This allows the model to maintain balance between its thermodynamic fields in 3-D, without which the model would be prone to generate spurious results. In a simple version of 3D-VAR all observations within a relatively short time window of each assimilation cycle are used and assumed to represent the same validity time. This process is shown schematically

in Figure 1-3. An improved version of 3D-VAR is used operationally in the Met Office mesoscale model, which compares observations with the First Guess at Appropriate Time (FGAT), rather than just the nominal analysis time.

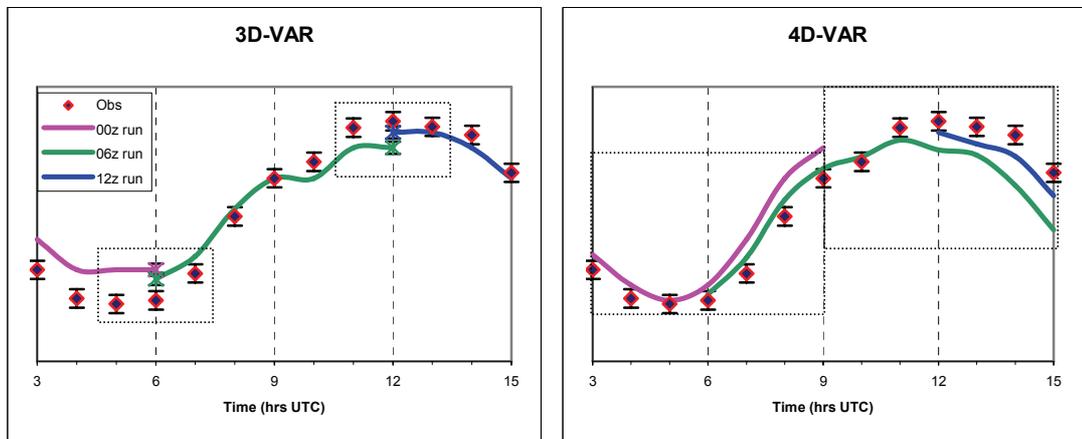


Figure 1-3 Schematic diagrams showing process of assimilating hourly observations (points with error bars) in boxed windows using previous model run as background. 3D-VAR using all observations within ± 1.5 hr assimilation window and single T+6 hr forecast as background to initiate a new forecast run every 6 hr. 4D-VAR uses all observations within ± 3 hr assimilation window and T+3 to T+9 hr forecasts as background to initiate new forecast run every 6 hr.

4D-VAR is an extension of 3D-VAR, where observations are assimilated continuously over the assimilation window of the NWP model runs, as shown schematically in Figure 1-3, for example. This allows the model to build a complete, dynamically consistent structure which is better retained in the forecast. This, in turn, allows the assimilation scheme to better select between conflicting information from different observation types. 4D-VAR also allows greater exploitation of isolated observations – horizontally, vertically or temporally and may allow a series of high time-resolution observations to be interpreted correctly as the atmosphere advects over the observation site. However, 4D-VAR is very computationally expensive as it requires running the NWP model iteratively forward and backward over the assimilation window. For this reason, it is only currently implemented in the global and larger regional versions of the Met Office Unified Model, but is likely to become operational in convective scale models in a few years.

Instead of assimilating observations directly in 4D-VAR, it is also possible to assimilate 1D-VAR retrievals of atmospheric profiles in a process referred to as 1D+4D-VAR. Marécal and Mahfouf [2003] demonstrated this provided more robust convergence than direct assimilation of observations in 4D-VAR due to problems in their particular implementation. However, they noted that the 1D-VAR minimisation did not convergence towards the absolute minimum.

While this method is not optimal, because it uses the background twice, it has practical advantages of allowing profiles to be retrieved independently off-line and for these to be assimilated by the same method used for radiosondes.

1.4 Introducing the Radiometrics TP/WVP-3000

This study concentrates on the application of ground-based microwave radiometers and is based on a commercial system – the Radiometrics TP/WVP-3000, shown in Figure 1-4. This radiometer has 7 channels in the oxygen band from 51-59 GHz for temperature profiling and 5 channels between the water vapour line at 22.235 GHz and the window region at 30 GHz, to provide information on humidity and liquid water cloud. The centre frequencies of its channels are shown on Figure 1-1.



Figure 1-4 Radiometrics TP/WVP-3000 Microwave Radiometer System at Camborne.

This system has participated in several trials and experiments at different locations over 4 years and proved very reliable. During this period extensive liaison with the manufacturers has resulted in a greater understanding of the system and a number of improvements, both hardware – to improve its performance in rain and resilience to corrosion, and software – to speed up the observing cycle, reduce the time between calibrations and the noise introduced by them. Its performance has been evaluated by different tests [Hewison and Gaffard, 2003].

1.5 The data sets used in this study

Table 1-2 summarises the trials and experiments the Radiometrics TP/WVP-3000 has participated in, providing the data used in this study. The radiometer has been based at Camborne for most of the period, where it has been characterised in terms of noise, calibration accuracy and its reliability assessed. Another microwave radiometer – Radiometer

Physics' Humidity And Temperature Profiler (HATPRO) [Rose *et al.*, 2005] – was also operated here from Oct 2004 to Jan 2005, allowing direct comparison of the observations. Profiles retrieved from the radiometers' data have been validated against numerous radiosondes launched from this operational sounding station in the UK upper air network.

Table 1-2 Dates and locations of deployments of Radiometrics TP/WVP-3000 microwave radiometer used in this study

Trial	Dates	Location	Sondes	Other Instruments	Aims of participation	Reference
Trial	2002-2005 except dates listed below	Camborne, UK	~2000	Wind Profiler Cloud Radar HATPRO	Reliability, Noise, Calibration and Retrieval performance	Hewison and Gaffard [2003]
Trial	Spring 2003	Cardington, UK	38	WVR1100 Ceilometer	Stratocumulus studies	Hewison [2003]
CWAVE	Summer 2003	Chilbolton, UK	0	Cloud Radar WV LIDAR	Water Vapour Retrieval validation	Agnew and Hewison [2004]
TUC	Winter 2003/04	Payerne, Switzerland	264	Wind Profiler Cloud Radar ASMUWARA	Retrieval Techniques: ability to retrieve inversions and fog	Ruffieux <i>et al.</i> [2006]
CSIP	Summer 2005	Linkenholt, UK	57	Wind Profiler Ceilometer HATPRO	Convection studies: stability	Browning <i>et al.</i> [2006]

The radiometer was also deployed at Cardington and Chilbolton in 2003 to participate in other trials and its retrievals compared to the integrated water vapour and liquid water derived from a two channel microwave radiometer, the Radiometrics WVR1100 [Hewison, 2003] and humidity profiles from a water vapour lidar [Agnew and Hewison, 2004]. 38 radiosondes were launched at ~hourly intervals on 5 days during the Cardington trial to validate the retrievals.

In the winter of 2003/04, the radiometer participated in the international Temperature hUmidity and Cloud (TUC) campaign [Ruffieux *et al.*, 2006], together with several other instruments, aiming to allow the development of *Integrated Profiling* retrievals. This included the All-Sky Multi-wavelength radiometer (ASMUWARA), whose observations were compared with the Radiometrics TP/WVP-3000 as brightness temperatures [Cimini *et al.*, 2006a], retrieved profiles [Cimini *et al.*, 2006b], IWV retrievals [Martin *et al.*, 2006] and with absorption models [Hewison *et al.*, 2006]. The climate of Payerne, Switzerland (altitude 491 m) is favourable to generate strong inversions and fog during winter because of katabatic drainage flow from the nearby Jura and Prealp mountain ranges. This provided a useful data set for testing the performance of profile retrievals, which could be validated against co-located radiosondes from this operational sounding station.

The Convective Storms Initiation Project (CSIP) [Browning *et al.*, 2006] took place in central, southern England during summer 2005. The Radiometrics TP/WVP-3000 took part on the same site as a wind profiling radar and within 2 km of the HATPRO radiometer to investigate these systems' abilities to resolve the structure of the convective boundary layer by sampling at high temporal resolution (~30 s) [Hewison and Gaffard, 2006]. 57 radiosondes were also launched from this site at ~hourly intervals on 8 days.

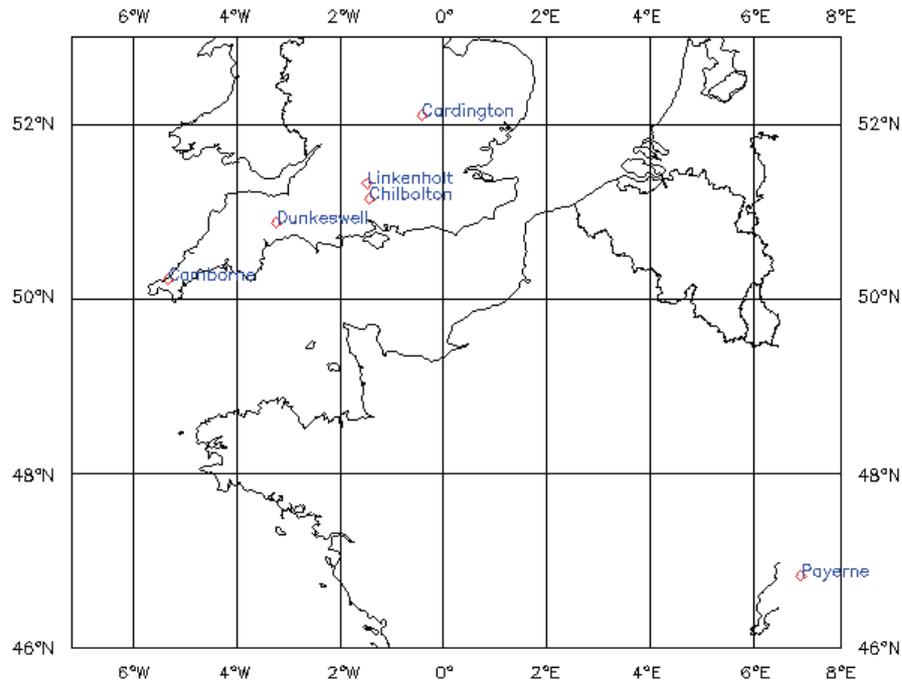


Figure 1-5 Location of Radiometrics TP/WVP-3000 deployments used in this study.

The locations of these sites are shown in Figure 1-5. Profiles interpolated from the mesoscale version of the Met Office Unified Model have been archived at all these sites for the duration of these trials, including the analysis and hourly forecasts from model runs every 6 hours. During this period, the Mesoscale model used a horizontal grid of 12 km with 38 levels in the vertical and covered the whole British Isles, extending southwards to cover the area shown in Figure 1-5. These model profiles have been used to provide background for the 1D-VAR retrievals and compared with the radiometer data in this study.

1.6 Thesis Outline

The remainder of this thesis discusses the following questions in more detail using data from this radiometer obtained in the above trials and experiments during the course of this work:

- What is a microwave radiometer?
- How good are its observations?
- How can it be used to retrieve temperature and humidity profiles?
- How good are the profiles? And why?
- What are the problems with the observations and their retrievals?
- How can they be improved?
- How can they be exploited?

Chapter 2 describes the components of the radiative transfer models needed by the retrievals' forward model to translate *state space* into *observation space*. The uncertainties involved and possible optimisations are also discussed. Chapter 3 describes the Radiometrics TP/WVP-3000 – the microwave radiometer used in this study. Results of various tests are analysed to characterise its performance in terms of calibration bias and radiometric noise, as well as practical issues. A variational technique to retrieve profiles from the radiometer's observations is developed in Chapter 4. This includes a series of experiments conducted to investigate different aspects of the retrievals. Chapter 5 analyses the performance of a different retrieval method: the neural network. These retrievals are used in Chapter 6 to investigate case studies to show the benefits and limitations of the radiometer observations. Finally, Chapter 7 extends the analysis of the previous chapters to draw together some recommendations for how microwave radiometer observations could be optimised for operational use. Each chapter closes with a summary of its key findings, with the conclusions summarised in Chapter 8.

Chapter 2 Radiative Transfer Modelling

Retrievals need a forward model to transform from state space of the temperature, humidity and liquid water profile into observation space of radiances measured by the radiometer. The key component of the forward model is the *Radiative Transfer Model* (RTM), which is described in this chapter. As illustrated in Figure 2-1, the predictions of the RTM depend on the input atmospheric variables, the radiometer frequency, ν , and zenith angle, θ , and boundary conditions given by the cosmic microwave background (§2.1.3).

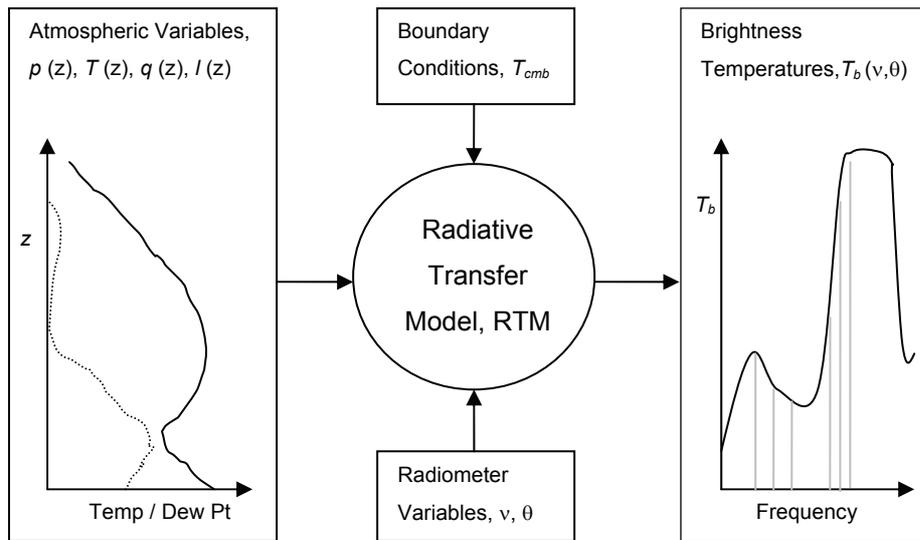


Figure 2-1 Schematic diagram of the *forward problem* of estimating observations from an atmospheric profile using the radiative transfer model used in this study.

2.1 Introduction to Radiative Transfer

The fundamental basis of any remote-sensing application is the *Radiative Transfer Equation*, first developed by Chandrasekhar [1960], which describes the flow of energy to the instrument. This can be derived by considering the physical principles responsible for emission, absorption and scattering of a signal, referred to jointly as *extinction*.

The *specific intensity*, I_ν , describes the power that flows at a given point in a given direction per unit frequency, ν , per unit solid angle. The differential form of the radiative transfer equation (RTE) considers how I_ν changes due to sources and sinks along its path, ds :

$$\frac{dI_\nu}{ds} = -I_\nu \alpha + S \quad (2.1)$$

where α , the absorption coefficient and S , the source terms describe the loss and gain of energy.

If scattering can be neglected (as shown later in this chapter), α is a scalar quantity describing absorption in the medium and S describes its emission. In local thermodynamic equilibrium at temperature, T , emission is given by Kirchhoff's law assuming a balance between emission and absorption (these terms will be used interchangeably in this thesis):

$$S = \alpha B_\nu(T) \quad (2.2)$$

where $B_\nu(T)$ is spectral brightness given by the Planck function:

$$B_\nu(T) = \frac{2h\nu^3}{c^2} \frac{1}{e^{h\nu/kT} - 1} \quad (2.3)$$

where h is Planck's constant, k is the Boltzmann constant, c is the speed of light.

Substituting these expressions for S and $B_\nu(T)$ into (2.1) and integrating gives this form of the RTE:

$$I_\nu(0) = I_\nu(s_0)e^{-\tau(s_1)} + \int_{s_0}^{s_1} B_\nu(T)e^{-\tau(s)}\alpha(s) \cdot ds \quad (2.4)$$

where s_0 and s_1 are the limits of the propagation path – in this case, the surface and space, $\tau(s)$ is the opacity, defined as:

$$\tau(s) = \int_{s_0}^s \alpha(s') \cdot ds' \quad (2.5)$$

2.1.1 Definition of Brightness Temperature

At low frequencies used in microwave radiometers, $h\nu \ll kT$ for tropospheric temperatures. This allows the Planck law to be simplified by the Rayleigh-Jeans approximation, as a linear relationship with temperature:

$$B_\nu(T) \approx \frac{2\nu^2 kT}{c^2} = \frac{2kT}{\lambda^2} \quad (2.6)$$

where λ is the wavelength.

This is used to define the *brightness temperature*, T_b , as an equivalent unit of I_ν or B_ν , as the physical temperature of a black body to emit at $T_b = T$, expressed in units of kelvin (K). Importantly, this is consistent with the definition of T_b used to calibrate the radiometer against the known temperature of reference black body targets. T_b is used as the unit of radiance throughout this study.

2.1.2 The Radiative Transfer Equation

This definition allows us to express the basic form of the RTE used in this thesis:

$$T_b = T_{cmb} e^{-\tau(s_1)} + \int_0^{s_1} \alpha(s) T(s) e^{-\tau(s)} ds \quad (2.7)$$

where T_b is the brightness temperature at the receiver. T_{cmb} is the brightness temperature of the background, s is the propagation path, $\alpha(s)$ is the absorption coefficient, $T(s)$ is the physical temperature, and $\tau(s)$ is the opacity along the path, ds , defined by (2.5).

The calculation of T_b by applying a discrete form of (2.7) is referred to as a *Radiative Transfer Model* (RTM). The discretisation is performed by calculating the average virtual temperature, specific humidity and cloud liquid water content over finite layers as described in §2.6.

2.1.3 Cosmic Microwave Background

For radiation down-welling through the atmosphere, the cold space boundary condition, s_1 , is given by the *Cosmic Microwave Background* (CMB) in the absence of other extra-terrestrial sources, such as the Sun or Moon. This corresponds to the residual radiation field emitted during the recombination of protons and electrons to form hydrogen in the early universe, when the temperature was about 3000 K [Janssen, 1993]. This radiation is red-shifted and now appears to have an almost uniform brightness temperature, $T_c = 2.736$ K. To account for the departure from the Rayleigh-Jeans approximation, the effective brightness temperature of CMB, T_{cmb} is adjusted by a frequency dependent amount following Janssen [1993]:

$$T_{cmb} = \frac{h\nu}{2k} \left(\frac{e^{\frac{h\nu}{kT_c} + 1}}{e^{\frac{h\nu}{kT_c} - 1}} \right) \quad (2.8)$$

These offsets are small at these frequencies – 0.04 K at 22.235 GHz increasing to 0.24 K at 58.8 GHz and become more significant at millimetre wavelengths.

2.2 Gaseous absorption models

Absorption models lie at the heart of the forward model. They describe the physical processes responsible for generating the signal seen by the radiometer and used in the retrieval.

Atomic gases radiate electromagnetic waves at discrete frequencies, known as *line spectra* corresponding to transitions between electronic states. These tend to fall in the visible or ultraviolet parts of the spectrum. Molecular gases have additional emission lines due to vibration and rotational modes of the constituent atoms, which tend to occur at infrared and microwave wavelengths, respectively. The frequency of these lines depends on the difference of quantised energy levels of the transitions, which are determined by quantum theory. A detailed description of these processes and the coupling that can occur between them is given in Rosenkranz [1993].

2.2.1 Line broadening

In practice, spectral lines are not infinitesimally sharp. The constant motion of the molecules/atoms results in their emission being Doppler shifted, resulting in *Thermal Broadening*. This is an important process for microwave propagation in the mesosphere, where it typically broadens spectral lines by 10^{-6} times their frequency. In the lower atmosphere, this is surpassed by *Pressure Broadening* which results from the interaction between molecules. It is by exploiting the pressure broadening of spectral lines that we are able to get information on the species' vertical distribution needed to retrieve profiles of temperature and humidity with any practical microwave radiometer (with finite bandwidths).

The power absorption at a given frequency, $\alpha(\nu)$, due to a single species can be expressed as the sum of all absorption lines:

$$\alpha(\nu) = n \sum_{f,i} S_{fi}(T) F(\nu, \nu_{fi}) \quad (2.9)$$

where n is the number of molecules per unit volume, f and i denote energy levels of an isolated molecule, $S_{fi}(T)$ is the intensity for a single line for that molecule and $F(\nu, \nu_{fi})$ is a *line-shape function*.

In practice, it is common to limit the influence of each line to a finite bandwidth (e.g. 750 GHz) to speed up *line-by-line* radiative transfer calculations. To compensate for the small amount of absorption omitted in this process, the continuum contributions can be inflated accordingly. Different forms have been proposed for this line-shape function based on different models of the collisions between particles, from the simple *Lorentzian*, which is only valid in the sparse upper atmosphere to the more complex *Kinetic* model. However, the models considered here use the form derived by van Vleck and Weisskopf [1945]:

$$F(\nu, \nu_{fi}) = \frac{1}{\pi} \left(\frac{\nu}{\nu_{fi}} \right) \left[\frac{\gamma}{(\nu_{fi} - \nu)^2 + \gamma^2} + \frac{\gamma}{(\nu_{fi} + \nu)^2 + \gamma^2} \right] \quad (2.10)$$

where the *line-width parameter*, γ , is defined as half the frequency width at half peak intensity.

It is by exploiting the pressure broadening of spectral lines that we are able to get information on the species' vertical distribution needed to retrieve profiles of temperature and humidity with any practical microwave radiometer (with finite bandwidths).

As mentioned in Chapter 1, in the troposphere microwave absorption is dominated by oxygen and water vapour. These will be discussed in more details in the following sections. However, several minor species are also radiatively active.

2.2.2 Water Vapour

Rotational transitions of the electric dipole of water vapour produce resonant absorption lines extending from the microwave to sub-millimetre wavelengths, including 22.235 GHz and 183.31 GHz. There are numerous other water vapour lines at frequencies just above this region that also contribute to the microwave absorption.

In addition to these absorption lines, researchers have found it necessary to include a non-resonant contribution by a *water vapour continuum* for the models to match observations, especially in *window* regions between lines (see Figure 1-1). Although many theories have been developed to explain the process responsible for the water vapour continuum, it is usually represented as an empirical term used to fit observations to the theoretically based resonant terms. It is generally divided into two contributions, which are *self-* and *foreign-* broadened, due to collisions of water vapour molecules with molecules of the same species or other gases, respectively [Rosenkranz, 1998].

2.2.3 Oxygen

Oxygen produces strong resonant absorption due to transitions in the magnetic dipole spin-rotation band between 51.5-67.9 GHz. The oxygen spin-rotation band also includes non-resonant transitions leading to absorption that is relatively weak at 60 GHz, but that also contributes to line coupling in the resonant region. Although Zeeman splitting of the fine structure in this band can occur in the magnetic field, this is only noticeable in the upper atmosphere and its effect masked by emission from lower levels in the ground-based viewing geometry. Pressure broadening causes the multitude of lines in the 60 GHz oxygen complex to appear as a continuous feature in the ground-based viewing geometry.

2.2.4 Other species

There are numerous ozone lines at millimetre wavelengths. However, these have a negligible influence at the lower frequencies considered in this thesis, so it is not included in the RTM. As for nitrogen, otherwise inactive, there is collision induced absorption, which becomes significant above 100 GHz in very dry conditions. Although the strength of this term varies between the models, its impact is small. A simple model of the nitrogen continuum is included in the RTM used in this thesis [Rosenkranz, 1998]. Its partial pressure is easy to calculate, being a well-mixed gas.

2.2.5 Example Brightness Temperatures

The radiative transfer equation (2.7) has been evaluated to predict the down-welling brightness temperature (T_b) at zenith for the 1976 US standard atmosphere in clear conditions in Figure 2-2, using different absorption models. Channels in the 30-40 GHz window region show the lowest T_b , while those nearer the centre of the 22.235 GHz water vapour line, T_b

increases, but the atmosphere remains optically thin. It becomes optically thick for channels near the centre of the 60 GHz oxygen complex, where T_b approaches the average thermodynamic temperature of the lowest few hundred metres of the atmosphere. Although the model differences appear to be quite small on this scale, the figure serves to illustrate the range of T_b observed at different frequencies. The differences are discussed in the following section and their significance analysed in §3.4.2.

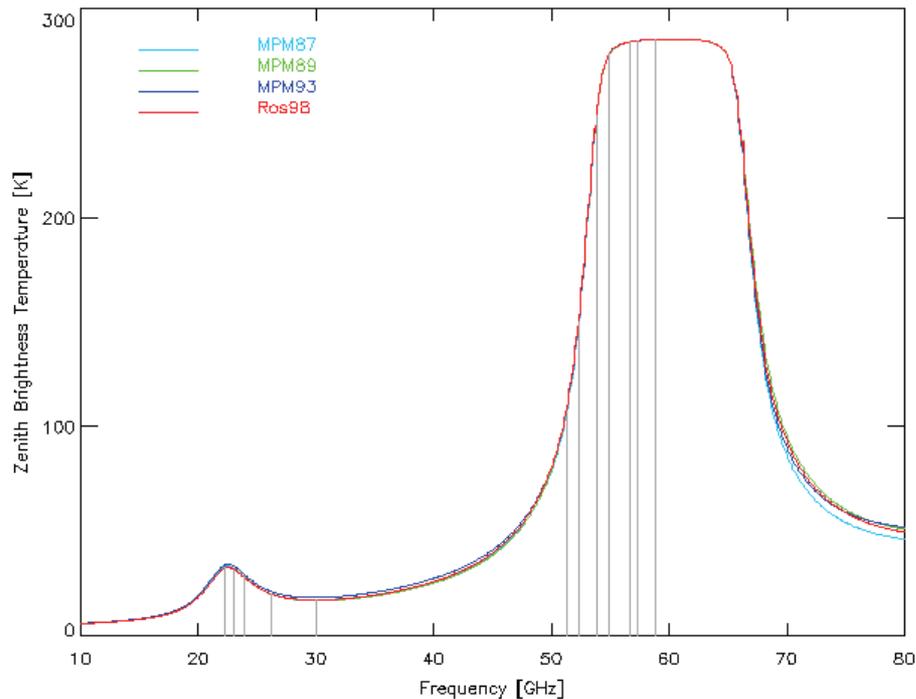


Figure 2-2 - Spectra of zenith brightness temperatures for clear standard atmosphere modelled with MPM87, MPM89, MPM93 and Ros98 absorption models. Vertical grey lines indicate centre frequencies of Radiometrics TP/WVP-3000 channels.

2.2.6 Absorption models available

The differences between the models pertinent to this study are discussed in the following subsections. A more general discussion is given in *Rayer [2001]*. Figure 2-2 shows the differences between the models are small outside the window regions, but these depend on the temperature, pressure, humidity and frequency. These models are compared to observations in §3.4.2, which led to the selection of Ros98 for use in the 1D-VAR retrievals developed in this study. However, other studies use the other models – for example, the neural network in Chapter 5 uses MPM87.

2.2.6.1 MPM87 [Liebe and Layton, 1987]

The clear air absorption part of the Millimetre-wave Propagation Model, MPM87 includes 30 water vapour lines and 44 oxygen lines all in the range 20-1000 GHz, based on theoretical

values and a Van-Vleck Weisskopf shape function. These are supplemented by an empirically derived water vapour continuum, fitted to laboratory observations at 138 GHz. However, these observations were limited to 282-316 K, and must be extrapolated for typical atmospheric conditions. Additional terms represent the non-resonant absorption due to the Debye spectrum of oxygen below 10 GHz and the pressure-induced nitrogen absorption above 100 GHz, which can become a significant contribution to the overall absorption in low humidity.

2.2.6.2 MPM89 [Liebe, 1989]

The 1989 revision of MPM modified the parameters describing the 183 GHz water vapour line, fitting the pressure broadened line width with four parameters, instead of one. Other components are the same as MPM87 for the purposes of this study.

2.2.6.3 MPM93 [Liebe *et al.*, 1993]

This version of MPM, has 34 water vapour lines between 20-1000 GHz, defined in a slightly different manner from MPM89. The 183 GHz line is 8.5% wider and 5% stronger than in MPM89. The water vapour continuum absorption is formulated as a *pseudo-line* near 2 THz, and has a different temperature dependence, based on newer observations. Like its predecessors, MPM93 includes 44 oxygen lines with the same line strengths, but 5% greater widths and 15% stronger mixing than MPM89. The non-resonant nitrogen absorption is essentially the same as MPM89 at the frequencies in this study.

2.2.6.4 Ros98 [Rosenkranz, 1998]

Ros98 uses 15 water vapour line parameters, which are very similar to the strongest lines used in MPM89. The other half of the lines have been omitted as they were judged to have negligible impact. Rosenkranz's investigations suggested a range of observations could be best modelled by using a water vapour continuum with a combination of MPM87's foreign-broadened component, and MPM93's self-broadened component. However, the water vapour lines used were truncated at ± 750 GHz, so the foreign- and self- broadened parts of the water vapour continuum were increased 15% and 3%, respectively to compensate. This model uses the same oxygen line parameters as MPM93, except at sub-millimetre frequencies, where values from the HITRAN database were used. It also uses a different form of non-resonant absorption due to pressure broadening by nitrogen.

2.2.7 Forward Model Spectroscopic Error

Random errors in the spectroscopy will be included in the observation errors in the retrievals and error analysis in Chapter 4. It would be difficult to perform a full error propagation through the radiative transfer model, even if estimates of the uncertainty on each spectroscopic coefficient were available, which they are not. Instead, forward model errors due to

spectroscopic uncertainties are estimated here as the difference in zenith brightness temperature calculated by 2 different absorption models, after removing their average difference.

The spectroscopic error has been evaluated using 1200 rocketsonde profiles spanning global conditions, and calculations of the zenith brightness temperature at 12 centre frequencies of Radiometrics TP/WVP-3000 channels for 2 different absorption models (MPM93 and Ros98). This was only done for clear skies.

The difference between the brightness temperatures calculated by the 2 models was then analysed. For the low frequency channels, it was found to depend on the Integrated Water Vapour (IWV), so the statistics have been divided into latitudinal bands: arctic, mid-latitude and tropical, as well as the global average. The covariance of the differences for mid-latitudes was then calculated (after removing average difference) to estimate the spectroscopic component of the forward model error in §4.4.1.2. The standard deviations of the brightness temperature differences are given in Table 2-1. As expected, the largest uncertainties in the model are found in the channels near the centre of the 22.235 GHz water vapour line and in the wings of the oxygen band at ~52 GHz. There is strong correlation between adjacent channels, which may further degrade the resolution of the retrievals. However, this empirical method may overestimate the spectroscopic contribution to the overall observation error – these terms are larger than are assumed in the operational assimilation of satellite data, although their impact would be reduced in the downward viewing geometry.

Table 2-1 - Estimated Forward Model Spectroscopic Error

Centre Frequency (GHz)	Arctic (K)	Mid-Lat (K)	Tropical (K)	Global Average (K)
22.235	0.54	0.79	0.65	1.01
23.035	0.54	0.79	0.65	1.01
23.835	0.50	0.73	0.62	0.94
26.235	0.41	0.58	0.52	0.74
30.000	0.40	0.54	0.50	0.69
51.250	0.71	0.95	0.80	1.20
52.280	0.52	0.70	0.65	0.88
53.850	0.14	0.18	0.16	0.23
54.940	0.02	0.03	0.02	0.03
56.660	0.01	0.01	0.01	0.01
57.290	0.01	0.01	0.01	0.01
58.800	0.01	0.01	0.00	0.01

2.3 Cloud extinction models

Although microwaves can penetrate thick cloud layers, liquid water has strong absorption due to its dielectric properties and this must be accounted for in the forward model.

2.3.1 Liquid water cloud

As we go from gases to liquids, the interaction between particles increases. Liquids may be regarded as an enormous molecule with a correspondingly increased number of degrees of freedom, leading to a large number of closely spaced spectral lines, effectively becoming continuous.

Electromagnetic radiation interacts with the dipole moment of cloud ice and liquid particles depending on their dielectric properties and geometry. For non-precipitating liquid water cloud, the droplets are much smaller than the wavelength of microwaves, so scattering can be neglected and the absorption coefficient is given by the Rayleigh approximation. This is accurate within several percent for particles with diameter less than 2% of the wavelength. For the frequencies below 60 GHz considered in this study, the Rayleigh approximation is valid for all drops less than 100 μm diameter. This includes almost all non-precipitating liquid water clouds.

In the Rayleigh limit scattering is not significant compared to absorption, so extinction is not sensitive to the particle's size and can be estimated as a function of the cloud's liquid water content and temperature. The frequency dependence of liquid water's dielectric permittivity is typically modelled following a Debye [1929] relaxation process, describing the response of an ideal population of non-interacting dipoles to an alternating external electric field. This thesis adopts the approach of Liebe [1989] who fitted observations to a double-Debye process with temperature-dependent relaxation frequencies of ~ 10 GHz and ~ 400 GHz. In this formulation, the cloud absorption coefficient increases as approximately the square of the frequency below the relaxation frequency, but linearly at higher frequencies. However, it is difficult to conduct such tuning with observational data of atmospheric clouds because of the absence of in situ measurements of liquid water content (by radiosondes) and the very large spatial and temporal variability of cloud. In MPM, this is implemented using values originally derived for sea water. However, there is very little sensitivity to salinity above 10 GHz [Ellison *et al.*, 2003].

Haze formed from condensation of water onto hygroscopic nuclei has very small particle sizes ($\ll 1 \mu\text{m}$) and very low liquid water content. It is, therefore, completely transparent to microwaves.

2.3.2 Supercooled Liquid Water Cloud

Liquid water can exist at temperatures below its nominal freezing point where there is a deficit of heterogeneous ice condensation nuclei, but is unstable to glaciation in the presence of ice crystals. Thin layers of pure supercooled liquid water cloud regularly occur in mid-latitudes [Hogan *et al.*, 2004], but tend to have low integrated liquid water (ILW) content, so microwave radiometers have low sensitivity to them. Supercooled liquid water is also prevalent in mixed phase cloud associated with dynamic frontal systems. These can have substantial ILW and present a serious icing hazard to aviation.

However, there is considerable uncertainty in the dielectric properties of liquid water below 0°C. Meissner and Wentz [2004] presented a model of the dielectric constant of sea water based a double Debye relaxation fitted to observations including supercooled water. The model of Liebe [1989] provides a good fit to their model at $T > -20^\circ\text{C}$ and $\nu < 100$ GHz. The differences increase at lower temperatures and higher frequencies.

2.3.3 Ice cloud

Unlike liquid water, the dielectric constant of ice is nearly independent of frequency and provides very little loss in the microwave region. Ice crystals found in cirrus clouds and dry falling snow are virtually transparent at microwave frequencies because of their very low dielectric loss. It is only when ice particles become large – in convective storms, for example – that scattering becomes significant below 60 GHz. Wet snow results in absorption and scattering similar to that of large rain drops due to its size and liquid water content.

2.4 Precipitation

As discussed in §2.3, scattering becomes significant for particles larger than a small fraction of the wavelength and becomes increasingly efficient for larger particles. At the frequencies considered here (20-60 GHz), scattering is significant for rain drops, wet snow, wet hail and graupel. The sensitivity of brightness temperatures to precipitation and its implications for retrievals is discussed in this section.

2.4.1 Extinction dependence on drop size and rain rate

Like cloud, precipitation droplets follow a continuous size distribution, spanning more than an order of magnitude. The modal size, number density and vertical extent of precipitation typically increase with the intensity of the precipitation, causing microwave radiometers to be much more sensitive to heavier precipitation.

Some early authors on the subject combined the effects of drop size, concentration and column height by measuring the extinction over a range of realistic conditions, and derived relationships between the extinction, κ_{er} and rain rate, R_r . e.g. Crane [1971] presented a

logarithmic model of the form, $\kappa_{er} = \kappa_1 R_r^b$ where κ_1 and b are coefficients that depend on frequency and polarisation. Extracting the relevant data from [Crane, 1971] for the average of V and H polarisations yields the results in Table 2-2:

Table 2-2 Simple model of extinction from precipitation as a function of Rain Rate for the average of V and H polarisations

Frequency (GHz)	Coefficients from Crane (1971)		Extinction (dB/km) at Rain Rate (mm/hr)		
	k_1	b	0.1	1	10
24	0.105	1.045	0.009	0.105	1.16
30	0.180	0.990	0.018	0.180	1.76
40	0.345	0.920	0.041	0.345	2.87
50	0.508	0.868	0.069	0.508	3.74
60	0.670	0.815	0.103	0.670	4.38

These data suggest a maximum rate of accumulation of drizzle of 0.1 mm/hr can be accommodated without compromising the accuracy of retrievals. A 300 m column of light drizzle at 0.1 mm/hr would introduce ~0.02 dB of loss at 50 GHz, which would contribute ~1 K to the zenith brightness temperature observed. This is approximately the limit of what could be resolved at this frequency because of the typical fluctuations in signal due to cloud variability. At higher frequencies, atmospheric absorption from oxygen is sufficiently strong for the extinction by light rain to be negligible.

However, it is difficult to validate such models due to serious problems measuring extinction (radiometers and radars are affected by liquid water), rain rate (tipping bucket type rain gauges are slow to respond to drizzle), vertical distribution of cloud and rain liquid water content and their drop size distribution and its vertical distribution [Jones, 1995].

In heavier rain, as drop sizes exceed ~1 mm, they become more oblate. This introduces a polarisation dependence in their scattering cross-section, which has been exploited to discriminate liquid water from clouds and rain [Czekala *et al.*, 2001].

2.4.2 Drizzle

Drizzle drops are usually defined as ranging in diameter from 0.2-0.5 mm [AMS Glossary], although radar meteorologists consider droplets greater than 50 μm to be drizzle, due to the strong sensitivity of radar reflectivity to very small numbers of larger drops. The latter is typical of drizzle from maritime stratus commonly found in the British Isles. Such drops are much smaller than the wavelengths used by the radiometer, so scattering will be negligible, as seen in Figure 1-2.

2.4.3 Practical considerations

In addition to extinction by airborne hydrometeors, liquid water on the radiometer's window (*radome*) also contributes to the observed brightness temperature by both emission and scattering, depending on the drop sizes. If this is allowed to accumulate, it will soon dominate the atmospheric signal. For example, a typical cloud may have an Integrated Liquid Water of 100 g/m^2 . This is equivalent to a $100 \text{ }\mu\text{m}$ thick layer of water on the radome.

To mitigate the effect of rain and condensation from dew or fog, microwave radiometers are often fitted with blowers and hydrophobic radomes. But even with these, there may still be residual water on the radome during and after precipitation, which influences the signal similar to rain droplets in the air. Ware *et al.* [2004] showed a new radome coating allows unbiased retrieval of Integrated Water Vapour (IWV) from zenith T_b observations at rain rates up to $\sim 1 \text{ mm/hr}$. This is consistent with the findings of this study comparing IWV retrieved by the radiometer and GPS with a co-located tipping bucket rain gauge (§5.5.2). It was also found that the difference between these IWV retrievals returned to nominal levels within a few minutes of the cessation of precipitation.

2.4.4 Maximum Integrated Liquid Water (ILW)

The Integrated Liquid Water path can be used as a simple check on the likelihood of the presence of precipitation. ILW can be retrieved from microwave radiometer observations as described in §1.3.1. In this study, the retrievals of temperature and humidity were compared to radiosondes at different levels and found to be independent of ILW for $\text{ILW} < 1000 \text{ g/m}^2$ (§6.5). However, the retrievals became noticeably biased for higher values of ILW due to scattering by precipitation forming by coalescence within the cloud. Hence this is taken as empirical threshold for ILW as a quality control throughout this study to indicate the presence of rain which may not reach the ground or be detected by the radiometer's rain sensor.

2.4.5 Data availability due to precipitation

The weather codes reported at the time of launch of radiosondes from Camborne were analysed over 1 year [Hewison and Gaffard, 2003]. The results showed that precipitation was reported 18.4% of the time, but more than half of this was slight rain or drizzle (less than 0.5 mm/hr). This suggests that even in Camborne's maritime climate, more than 90% of data should be uncontaminated by precipitation.

2.4.6 Rain Water Content

As drizzle particles only contribute to microwave radiances by emission, their effect is indistinguishable from that of liquid water cloud. So, the ideal choice of control variable for the retrieval would be the total liquid water – including clouds and rain. In practice, drizzle is unlikely to contribute a significant fraction of this total, so it may be neglected. However,

active instruments are more sensitive to the drop size of the scatterers, so it may be preferable to separate these terms in the control variable when combining observations from microwave radiometers and wind profilers or cloud radars.

2.4.7 Implications for retrievals

Accurate modelling of scattering by larger hydrometeors needs accurate data on their microphysical parameters which are not typically available in NWP models. Hence, it is likely that any observations in moderate to heavy precipitation will need to be rejected from the assimilation cycle, at least initially. This introduces a need for quality control to detect the presence of precipitation heavy enough to introduce significant scattering at the frequencies being used. This may be implemented by (the last two are used in this study):

- Looking for a polarisation difference – if available (not at zenith).
- Monitoring the residuals of the observations compared to the forward-modelled background as part of the assimilation process. A T_b excess, correlated with frequency is a tell-tale sign of scattering from precipitation. This is used in the operational processing of data from space-borne microwave radiometers.
- Checking the data from a co-located rain sensor (although this will not detect any precipitation unless it reaches the surface).
- First performing a simple retrieval of Integrated Liquid Water (ILW) and rejecting any observations where this is greater than a climate-dependent threshold (§4.8.3).

2.5 Speeding up Radiative Transfer Calculations

The *line-by-line* models described in §2.2.5 rely on summing the influence of every line in their database to calculate the absorption at a single frequency. This is a computationally expensive process, taking time t_{xl} . This becomes problematic when repeatedly evaluating it to calculate the matrix of the forward model's partial derivatives, known as the *Jacobian*, by *brute force* (see §4.5.1). To speed up the calculation a number of approximations and parametrizations can be made, in addition to reducing the number of lines in the spectroscopic database and truncating each line's influence mentioned in §2.2. These are described in this section.

2.5.1 Is a fast model needed?

For a line-by-line model to calculate the brightness temperatures at $n=12$ frequencies with a profile of temperature and humidity defined at $m=38$ levels takes $n.m.t_{xl} \approx 38$ ms using un-optimised PV-Wave code on a Linux workstation with a single 2.4 GHz Pentium-IV processor. To evaluate the Jacobian, \mathbf{H} , using *brute force* by perturbing each of $2m$ levels of the profile would take $2.m^2.n.t_{xl} \approx 2.9$ s. Using a non brute-force method to calculate Jacobians requires the differentiation of the forward model, but only takes twice as long as running the forward model itself. In principle, exact minimisation of a cost function, J , requires $2(m+1)$ evaluations

of J and \mathbf{H} , which would take $4(m+1).m^2.n.t_{\text{int}} \approx 225$ s by brute force. This is longer than the possible sampling interval of a microwave radiometer and is regarded as impractical. This could be reduced by:

- Differentiating the forward model to calculate \mathbf{H} – a factor of ~ 100
- Using some form of fast forward model – maybe by a factor of ~ 100
- Improving the code of the forward model – maybe by a factor of ~ 10
- Using approximations in minimisation – maybe by a factor of ~ 10
- Using faster computers – maybe by a factor of ~ 2

2.5.2 Representation of channels' finite bandwidth

Radiometers' channels are not truly monochromatic – they have finite bandwidths to reduce their radiometric noise. Heterodyne systems usually have a spectral response with one side bands on either side of a central frequency with central stop band to reject interference from the local oscillator (§3.1.3). However, the atmospheric absorption changes across the bandwidth of each channel. If it changes nonlinearly with frequency, a model of T_b evaluated at the channel's centre frequency can be significantly biased. To overcome this, the channels' passbands have been sampled at high spectral resolution and radiative transfer calculated for each frequency and the resulting T_b s averaged – the *band-averaged* T_b . However, this is computational expensive for a relatively small correction (~ 0.1 K), so it is desirable to use an *effective monochromatic frequency* (EMF) for each channel. The EMF was selected to correspond to the monochromatic frequency that minimises the difference with the band-averaged T_b for a representative data set of atmospheric profiles. As we will show in this section, the EMF does not always correspond to the nominal central frequency.

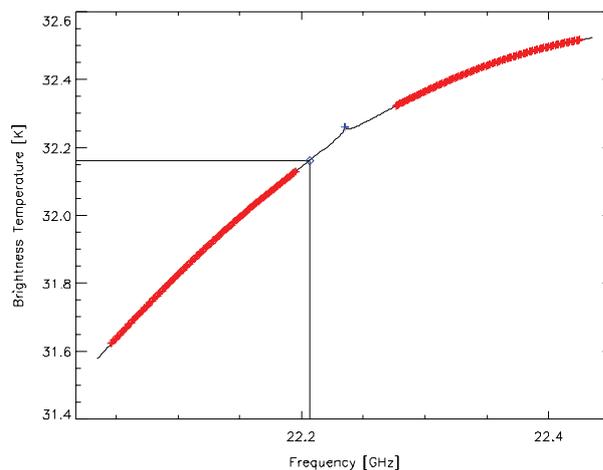


Figure 2-3 – T_b spectrum near 22.235 GHz channel of Radiometrics TP/WVP-3000.

Calculated using Ros98 for 1976 US Standard Atmosphere in zenith view.

Red crosses indicate the passband of the 22.235 GHz channel. Blue cross shows T_b evaluated at the channel's centre frequency. Blue diamond shows the band-averaged T_b and the Equivalent Monochromatic Frequency (EMF) at which this occurs.

From the TUC set of radiosonde profiles, T_b s were calculated for each channel at 40 frequencies separated by 5 MHz steps and located in a 200 MHz window centred at the nominal central frequency. Those were convolved with the channels' passband response (assumed rectangular) to calculate the band-averaged T_b s. Then, for each profile and for each channel, the frequency corresponding to the minimum difference with respect to the band-averaged T_b was identified as the EMF. This process is illustrated in Figure 2-3. The resulting EMFs shown in Table 2-3 are broadly consistent with a similar analysis done using data from Camborne [Hewison and Gaffard, 2003]. Table 2-3 also shows the statistics of the T_b differences between calculations at the nominal centre frequency, band-averaged and EMF. The mean T_b difference between band-averaged T_b and that evaluated at the centre frequency is less than 0.36 K, the largest values being found at 51 to 54 GHz channels, where the absorption changes most rapidly with frequency. Thus, the effect of a finite bandwidth is small, but not negligible. Because the SD is very small (less than 0.1 K) for all but the 22 GHz channel, most of the channels could take this effect into account as a mere bias. However, using the EMF reduces the uncertainty to negligible levels, so this approach has been adopted for the remainder of this study.

2.5.3 Fast Absorption Predictor (FAP)

To speed up the calculation, a *Fast Absorption Predictor* model is used to calculate the absorption as a third order polynomial function of pressure, temperature and specific humidity, following Löhnert *et al.* [2004]. The polynomial coefficients are derived *a priori* using a set of profiles by regression of the brightness temperatures modelled by a line-by-line model for each channel against the pressure, temperature and humidity at each level. This is only applied to pressures greater than 100 hPa as it was found to be inaccurate above this. At lower pressures, \mathbf{H} is not re-calculated, so this does not slow down the iterative calculations.

The data set of model profiles was divided in two – half was used to train the FAP model, the other half to validate it against T_b predicted by the line-by-line model used for the training. The statistics of this comparison, given in Table 2-3, show that FAP results in satisfactory levels of random error (i.e. smaller than the error of representativeness as will be shown in §4.4.1).

Table 2-3 - Centre frequency and equivalent monochromatic frequency (EMF) of Radiometrics TP/WVP-3000 calculated from TUC radiosondes, showing mean (bias) and standard deviation (SD) of difference in brightness temperatures, T_b , calculated at these frequencies and using Fast Absorption Predictor (FAP) compared to that averaged over the channels' passbands (band average).

Centre Freq.	EMF	Tb(centre freq) - Tb(band average)		Tb(EMF) - Tb(band-average)		Tb(FAP) - Tb(EMF)	
		Bias	SD	Bias	SD	Bias	SD
(GHz)	(GHz)	(K)	(K)	(K)	(K)	(K)	(K)
22.235	22.205	-0.144	0.289	-0.004	0.004	0.006	0.111
23.035	23.040	-0.021	0.009	-0.001	0.006	-0.030	0.205
23.835	23.835	0.000	0.001	0.000	0.001	-0.121	0.334
26.235	26.230	0.004	0.001	0.002	0.002	-0.249	0.331
30.000	30.005	0.001	0.000	0.000	0.000	-0.190	0.265
51.250	51.255	0.139	0.005	-0.005	0.003	-0.010	0.475
52.280	52.285	0.360	0.013	-0.099	0.014	0.031	0.433
53.850	53.855	0.348	0.047	0.056	0.052	0.007	0.166
54.940	54.935	-0.032	0.008	0.006	0.016	0.006	0.025
56.660	56.655	-0.003	0.002	0.001	0.001	-0.002	0.005
57.290	57.285	-0.001	0.001	0.000	0.001	-0.001	0.004
58.800	58.800	0.001	0.001	0.000	0.000	-0.002	0.005

2.5.4 RTTOV (Radiative Transfer for TOVS)

RTTOV [Saunders *et al.*, 1999] is a fast radiative transfer model for nadir-viewing passive infrared and microwave satellite radiometers, spectrometers and interferometers. Given an atmospheric profile of temperature, water vapour and, optionally, ozone and carbon dioxide, together with satellite zenith angle, surface temperature, pressure and, optionally, surface emissivity, RTTOV will compute the top of atmosphere radiances in each of the channels of the sensor being simulated. In addition to evaluating T_b using the forward model, RTTOV also returns its Jacobian matrix and its *tangent linear* values, either of which is needed for variational retrieval methods (see Chapter 4).

The coefficients used in RTTOV need to be specially trained for each instrument's channels and viewing geometry. However, the basic geometry assumed corresponds to nadir-viewing satellite instruments, and would need significant modification to adapt it to ground-based applications. Despite the benefits of providing a fast method of evaluating the forward model and its Jacobian, this effort was not justified for the purposes of this thesis. This may be reviewed in any future operational application where computational speed is of greater importance.

2.6 Profile Representation

To calculate T_b for any microwave radiometer sensitive to atmospheric emission, it is necessary to represent the atmospheric conditions along its optically path in the RTM. In this case, the vertical profile of pressure, temperature and specific humidity and liquid water (p, T, q, l) needs to be discretised to a manageable number of levels to allow fast calculation of T_b . The RTE (2.7) is then evaluated by integrating the emissions over the average (p, T, q, l) between the profile levels. Therefore, the choice of levels for discretisation represents a compromise between efficiency and accuracy of calculation. If the level spacing is too large to represent features of the profile that generate significant microwave emission, errors will be introduced in the modelled T_b .

In this study atmospheric profiles are derived from radiosondes and NWP models, which are discussed below.

2.6.1 Radiosonde profiles

Balloon-borne radiosondes can provide useful profiles of (p, T, q) for validation of radiometer observations and retrievals if they are sufficiently close in space and time. Modern radiosondes are capable of providing accurate profiles with high vertical resolution thanks to improved sensors, transmitters and receivers. For example, Vaisala claim a total uncertainty ($1-\sigma$) for their RS92 package of 0.75 hPa for pressure, 0.25 K for temperature and 2.5% for relative humidity with a response time of <0.5 s at low levels (the RH sensor has a much slower response in low temperatures). These specifications have been confirmed by comparison with other radiosondes, including devices using a chilled mirror hygrometer [Smout *et al.*, 2005]. The reproducibility between sensors in dual flights was found to be excellent with a standard deviation of 0.5 hPa, 0.2 K and 2%RH.

It is possible to quantify the likely magnitude of radiosonde errors and their impact on forward modelled T_b s by comparing the results from different sensors flown on the same balloon. This was done during the TUC experiment [Hewison *et al.*, 2006]. The SD was found to be largest (~ 0.5 K) for the channels near 22 GHz and in the wings of the oxygen complex near 52 GHz most sensitive to water vapour, reducing to ~ 0.2 K for the channels nearer the centre of the 60 GHz oxygen complex.

In operational use, radiosondes are required to ascend at 5.0-5.5 m/s through the troposphere. Thus it takes ~ 2000 s to reach a typical tropopause height of ~ 10 km, during which time the balloon will have drifted with the wind from its launch location, typically ~ 10 -100 km. In the UK it is operational practise to sample the profiles at 2 s intervals, corresponding to a vertical resolution of ~ 10 m. However, their data are currently transmitted (and assimilated) in a format with degraded resolution. Their profiles are first interpolated to a

set of 21 *standard levels* (only 6-8 of which are in the troposphere). *Significant levels* are first added at the surface and tropopause. Further significant levels are then added iteratively by a procedure which identifies where T or RH in the interpolated profile deviate by more than 1 K or 10% from the original. This is referred to as the *TEMP* code message and reduces data volume by a factor of ~ 100 compared to the 2 s data. These thresholds have been experimented with here to establish how much error is introduced by this discretisation process (§2.6.3).

2.6.2 Mesoscale model levels

Profiles from the Mesoscale version of the Met Office Unified Model (§1.1.1) in this study used a staggered Charney-Phillips vertical grid (G3), where T and q are defined at the surface (1.5 m) and on 38 levels and p , ρ and winds are defined at intermediate layers. These levels are more concentrated in the boundary layer (see Figure 2-4), which is convenient for our retrievals.

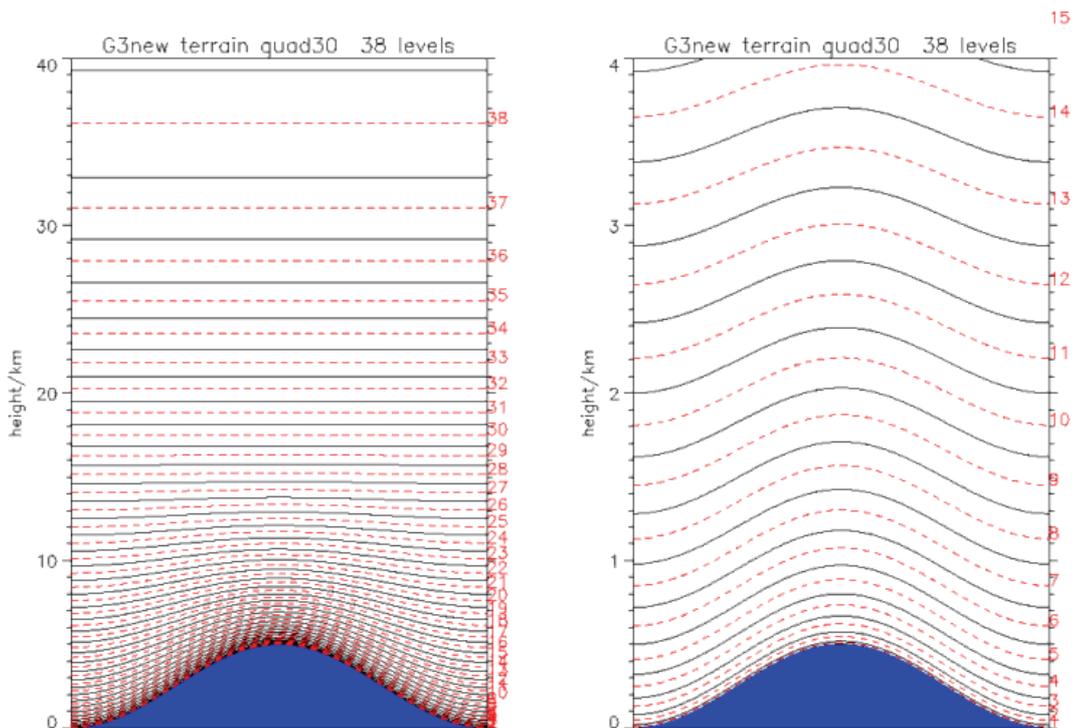


Figure 2-4 – G3 levels of Met Office Unified Model. Black lines show levels at which temperature and humidity are defined. Red dotted lines show intermediate levels for pressure, density and winds. (Right-hand panel is a zoom of the left hand panel.)

The UM has since been revised to a 50 level model by the addition of extra levels in the stratosphere and above. This change would not affect these retrievals. However, possible future improvements in the vertical grid will reduce the errors due to discretisation.

2.6.3 Impact of profile discretisation

It is important that the discretisation of the profile does not introduce excess errors in T_b calculated by the RTM. To test this, a set of high-resolution radiosondes was used as input to the RTM and the T_b calculated for the 12 channels of the Radiometrics TP/WVP-3000. These were compared with T_b calculated using the same profiles reduced by different methods of discretisation. *WMO Standard* and *UMG3* refer to profiles which have been interpolated at standard and mesoscale model levels only, respectively. *TEMP* refers to profiles that have been interpolated to WMO standard levels then extra levels added iteratively, defined as above. *Significant* is an improved version of this method, with tighter tolerances on departures from the original (0.5 K and 5 %RH), and subsequently more levels, but smaller errors. *Neural Network* refers to the fixed vertical levels used in the neural network retrievals: every 100 m from the surface to 1 km, then every 250 m up to 10 km (§5.2.1).

Table 2-4 Bias and r.m.s. error introduced to zenith brightness temperatures modelled for 57 clear radiosonde profiles using different discretisation methods

Channel Frequency (GHz)	Significant		TEMP		UMG3		WMO Standard		Neural Network	
	Bias (K)	rms (K)	Bias (K)	rms (K)	Bias (K)	rms (K)	Bias (K)	rms (K)	Bias (K)	rms (K)
22.235	0.15	0.21	0.38	0.55	-0.02	0.22	0.96	1.73	0.01	0.11
23.035	0.12	0.19	0.34	0.50	-0.04	0.21	0.88	1.63	0.01	0.11
23.835	0.09	0.14	0.25	0.38	-0.05	0.18	0.67	1.30	0.00	0.09
26.235	0.04	0.07	0.12	0.18	-0.05	0.11	0.31	0.65	-0.02	0.05
30.000	0.03	0.05	0.08	0.12	-0.06	0.09	0.20	0.42	-0.03	0.04
51.250	0.07	0.09	0.15	0.20	-0.39	0.47	0.48	0.60	-0.34	0.26
52.280	0.06	0.07	0.13	0.16	-0.46	0.56	0.41	0.49	-0.39	0.31
53.850	-0.02	0.03	-0.03	0.07	-0.36	0.42	-0.18	0.31	-0.28	0.24
54.940	-0.07	0.09	-0.14	0.18	-0.08	0.09	-0.63	0.74	-0.06	0.05
56.660	-0.06	0.11	-0.10	0.19	-0.04	0.06	-0.45	0.71	-0.06	0.10
57.290	-0.06	0.12	-0.09	0.20	-0.03	0.06	-0.39	0.71	-0.07	0.12
58.800	-0.06	0.13	-0.09	0.23	-0.03	0.06	-0.32	0.74	-0.08	0.14

The results are shown in Table 2-4 show that interpolating profiles at *Standard* levels only introduces significant errors in all channels. The *TEMP code* version of the profile adds substantial uncertainty in the channels near the water vapour line, but its errors are approximately halved by halving the tolerances for temperature and humidity departures in the *Significant* scheme. The levels used in the *Neural Network* retrievals, however, introduces similar, low error levels to all channels. Interpolating to UMG3 levels improves this for all except those channels on the edge of the oxygen band, which are dominated by other errors (§4.4.1.4). This shows that the UMG3 levels are suitable for use in retrievals without introducing excess noise.

The covariance of these discretisation errors is a significant component of the forward model error (§4.4.1.2), and includes substantial correlation between the channels close to the centre of the 22 GHz water vapour line and between channels on the edge of the oxygen band.

2.6.4 Atmospheric propagation path

In the case of a zenith viewing instrument, the RTE integration is straightforward – vertically along the propagation path. However, as a radiometer scans towards the horizon, corrections are needed due to the bending of the path by refractive index gradients and to account for the Earth’s curvature. Han and Westwater [2000] showed these effects were small (<0.2 K) for zenith angles less than 75°, so these can be neglected here, as there are other practical problems in exploiting radiometer observations close to the horizon, as will be discussed in §3.2.1.

At high zenith angles, it can also become important to account for the finite beamwidth of the microwave radiometer. This can be addressed by use of an effective zenith angle, increased slightly from the true angle by an amount depending on the atmospheric opacity and beamwidth – see §3.2.1.

2.6.5 Topping-up profiles

Radiosonde balloons usually reach 20 – 30 km before bursting. This can leave much of the stratosphere and all the mesosphere unrepresented in the profile used in the RTM. Although this part of the atmosphere has a very small contribution to T_b measured at the surface, it is possible to correct for this by appending a reference profile [McClatchey *et al.*, 1972] above the upper limit of the radiosonde measurements. Given the poor performance of some radiosonde hygrometers at the low temperature and humidity in the stratosphere, it is advisable to also substitute their measurements here with a reference profile. It was found that the choice of reference profile used for this *top-up* produced negligible differences to the modelled T_b s.

This is less problematic when using profiles from the mesoscale model, as the uppermost level of these is at 39 km currently. Although the temperature profile may not be very accurate near the model top, ground-based radiometers will be quite insensitive to errors here. Also the assumption of local thermodynamic equilibrium breaks down at higher levels (>70 km) and the atmospheric constituents become increasingly monatomic and require separate treatment in radiative transfer calculations, although their contribution to T_b at the surface is negligible.

2.6.6 Representing cloud in profiles

Radiosondes do not measure liquid water, so cannot be used to validate the retrievals’ cloud parameters. The mesoscale model includes prognostic cloud liquid water content at each

level of the profile. This can be used as first guess for retrievals and its covariance estimated by the NMC method [Parish and Derber, 1992], as explained in §4.3.2.

2.7 Validation of Radiative Transfer Models

Radiative Transfer Models (RTMs) and their component models of absorption and scattering by different media are often developed from theoretical bases which usually include assumptions and approximations or are based on laboratory measurements. When these models are applied to real world situations their validity should be checked against observations of the model's predictand. Validation is part of an iterative cycle of model development, allowing refinement of certain model parameters. Validation can be used to investigate the model's systematic biases, estimate its accuracy and also check it is capable of representing the full range of observed conditions.

Validation needs to include measurements of all model inputs, together with estimates of their uncertainty. Care is needed to ensure the observations of model inputs accurately represent the range of conditions influencing the model's predictand – this is a problem of sampling the observed scales at sufficient resolution and coverage. Often the uncertainty introduced by representativeness dominates the error budget of the validation experiment. The uncertainty in each component of the validation should be analysed – both in terms of random uncertainties which may be reduced by averaging and systematic biases (which may not).

A range of techniques can be applied to validate different aspects of RTMs, including:

2.7.1 Laboratory measurements for model validation

Although it is not always possible to reproduce real world conditions, laboratory measurements allow model predictors to be varied in a controlled way. This makes it easier to isolate the influence of each predictor, which is difficult in nature. It is also much easier to ensure observations of model inputs and outputs are representative in a lab environment, as the sampling can be controlled, though it may not be possible to explore the full range of parameter space.

The measurements of resonance spectrometer attenuation by Liebe and Layton [1987] illustrate this technique. They measured the absorption at 138 GHz over the range of 281-316 K, 0-3 bar in pure water vapour and a water vapour/nitrogen mix. These measurements were used empirically to fit the coefficients of self- and foreign- broadened water vapour continuum, including their temperature coefficients. The temperature and partial pressures of water vapour and nitrogen in the resonance cell were measured and controlled. If the chamber is homogeneous and in steady-state, the measurements are assumed to accurately represent the conditions influencing millimetre-wave absorption. However, this does not necessarily represent the full range of conditions found in nature.

2.7.2 Validation by ground-based radiometers and radiosondes

Ground-based radiometer or spectrometer measurements can be used for validation of RTMs given *in situ* observations of the model's input variables – typically from co-located radiosondes. If long-term data sets are available, it is possible to validate the full range of a model's inputs for a single site. As both the radiometer and radiosonde sample similar, small spatial volumes, representativeness error can be minimised if the validation is performed in homogeneous conditions. It is extremely difficult to validate cloud absorption models because of the large spatial and temporal variability of cloud properties and the shortage of good quality *in situ* measurements of microphysical parameters needed, so validation by this technique is limited to atmospheric absorption models in clear air. This technique is applied to radiometer observations from Camborne in §3.4.2.

Cimini *et al.* [2003a] compared observations of ground-based microwave radiometers at 20-30 GHz with coincident radiosonde profiles in clear air modelled using Ros98. They found a dry bias of ~5% using Vaisala RS80 humidity sensors, which was much reduced with RS90 sensors. Westwater *et al.* [2003] used similar radiometers in the tropics, but found variability in radiosonde humidity sensors dominated differences between absorption models at 23.8 GHz. They showed that MPM87 fitted observations better than Ros98 at 31.4 GHz, which is dominated by the water vapour continuum. Earlier work [Westwater *et al.*, 2001] had shown that Ros98 best fitted observations at 31.4 GHz in cold, dry arctic conditions.

Cimini *et al.* [2003b] conducted a similar analysis on data from 4 radiometers with a total of 19 channels between 20-59 GHz. They found that of these models Ros98 gave the best results at 20.6-20.7 GHz channels, while MPM93 was preferable close to 22.2 GHz. These 2 models stayed within 0.3 K at 23.8 GHz, but Ros98 gave the best results in the atmospheric window (~30 GHz), where emission is dominated by the water vapour continuum. Three models (MPM87, MPM93, and Ros98) showed a sharp change, from +2 to -2 K, when comparing 51.2 and 52.2 GHz channels with observations. At higher frequencies (55-60 GHz), they found the models were almost equivalent, but showed a negative bias (~1 K) with respect to radiometric observations.

Liljegren *et al.* [2005] analysed 5 channels between 22-30 GHz and showed that using a 5% smaller width of the 22.235 GHz line in Ros98 resulted in smaller bias in comparison with observed T_b and in retrieved profiles. Even temperature profiles above 3 km were improved by this modification, which is strange, as this information comes from the 51-52 GHz channels, which should not be sensitive to the width of this line. However it may be possible that the consistency between 20-30 and 51-53 GHz channels might improve the retrievals. Their results also suggested that replacing the Ros98 water vapour continuum with MT_CKD [Mlawer *et al.*, 2003] improved the fit with T_b trends. Mattioli *et al.* [2005] compared T_b s from 3 co-located Radiometrics instruments, each with channels at 23.8 and 31.4 GHz using 2

calibration algorithms, with Vaisala RS90 radiosonde measurements forward modelled with various absorption models. Their results supported the modification of the continuum proposed by Liljegren *et al.* [2005], but not the change in width of the 22 GHz line.

A similar analysis was performed as part of this study [Hewison *et al.*, 2006] using observations from TUC (§1.5) in a winter continental climate. The results show a systematic bias in the brightness temperatures measured by the Radiometrics profiler at 55-59 GHz, which has since been corrected in the control software. Observations at lower frequencies (22-30 GHz) in these cold conditions do not support Liljegren's proposed changes to the width of the 22.235 GHz water vapour line, although this is subject to the assumption of no residual bias in the radiosonde humidity. In these conditions, MPM89 and Ros98 provided the best fit to the radiometer observations. At intermediate frequencies (51-54 GHz), the absorption models produce large differences, which may be due to differences in oxygen line coupling [Boukabara *et al.*, 2003] and highlight the need for further laboratory measurements at low temperatures.

2.7.3 Airborne radiometer/spectrometer and in situ sampling

Installing a radiometer on a research aircraft offers the attractive potential to simultaneously sample the radiative and *in situ* properties of the atmosphere. With careful experiment design, this can ensure errors of representativeness are minimised. For example, the aircraft can provide a profile of pressure, temperature and humidity for input to a radiative transfer model. This can be compared with radiances measured in clear skies during level runs to validate the absorption models. This is akin to comparing a ground-based radiometer with forward modelled radiosondes, although the analysis can be repeated at different levels in the atmosphere. The errors of representativeness can be further reduced by comparing radiance measurements made *during* an aircraft profile with *in situ* measurements. The absorption can then be calculated over atmospheric layers for which simultaneous measurements are available. However, this requires careful control of the radiometer's calibration, which can change with temperature during aircraft ascents/descents. Hewison [2006] used this technique to validate absorption models at 89, 157 and 183 GHz, and found all models to underestimate the observed absorption in very dry conditions in the window region, while MPM93 tended to overestimate the strength of the water vapour continuum in warm, humid conditions.

English [1995] compared models and measurements of liquid water cloud absorption from the same airborne radiometer using *in situ* measurements of liquid water content. He found this to be very variable on small spatial scales.

2.7.4 Satellite measurements with *in situ* data

Satellite observations of model predictands are able to represent their global range of conditions. However, in practice, careful data selection is needed to isolate the variable of interest from the other factors which may also influence the scene radiance. The need for *in situ* measurements of model input variables also imposes stringent selection criteria. This may reduce a huge global data set to a small number of satellite observations co-located with *in situ* observations – e.g. radiosondes, buoys. Co-location is needed both in space and time. With polar-orbiting satellites, this can reduce the available data set for comparison to zero.

The fundamental problem in using satellite data to validate RTMs with *in situ* data is that these observations sample different spatial scales. The co-location criteria must be chosen by consideration of the spatial and temporal variability of the parameters in question. To minimise errors due to representativeness, it is often necessary to impose criteria on the variability of observations – both spatially for satellite data and temporally for *in situ* data.

An example of this technique is given by Buehler *et al.* [2004]. They compared radiances from the 183 GHz channels of the Advanced Microwave Sounding Unit (AMSU-B) with forward modelled radiances from radiosonde profiles. Comparisons were only made for satellite data within 30 minutes of the nominal time for radiosondes launched from Lindenberg, Germany and where the sonde drifted by less than 50 km before reaching the tropopause. They also filtered the data to reject dense ice cloud. Although this technique includes contributions from errors in the radiometers and *in situ* observations, radiative transfer model as well as uncertainty in the representativeness of the two observations, Buehler *et al.* were able to show their results were consistent with a dry bias in radiosonde humidity sensors in very dry conditions. Subsequently, this has been suggested as a method of evaluating the accuracy of different radiosonde observations.

2.7.5 Validation with Numerical Weather Prediction analysis

Another technique to exploit satellite data for model validation is by comparing it with Numerical Weather Prediction (NWP) analysis fields in lieu of *in situ* data. They have the advantage of global availability and, often, lower errors of representativeness, as they are sampled on scales similar to the resolution of space-borne microwave radiometers (~10 km). The underlying assumption is that the NWP field represents an unbiased state of the atmosphere as it is constructed to be consistent with a diverse set of current and past observations. A common method is to monitor the *Observation-minus-Background* (O-B) residuals as part of the variational assimilation of observations into NWP. This process calculates differences between the observations and the NWP model's short-range forecast. These *residuals* often show systematic biases, which can be further analysed in multivariate space – e.g. as a function of temperature, humidity, pressure, wind speed, scan angle, etc.

It is inevitable that this technique includes the observations' and the NWP model's bias as well as the RTM. However, these are equivalent from a data analysis point of view and their differences can be corrected empirically as part of the assimilation process as if they were all due to observation biases.

Ellison *et al.* [2003] applied this technique to validate models of ocean emissivity. They compared observations from various space-borne microwave radiometers with the models' predictions based on various NWP models. To isolate the component of interest (the permittivity of sea water), they filtered the available data sets to remove the influence of liquid water cloud and high wind speeds, which induce surface spray/foam. The results were then analysed over different temperature regimes. A similar technique was also used by Meissner and Wentz [2004] to validate a new fit to laboratory measurements of the permittivity of pure and sea water. They showed this model provides the best fit to SSM/I observations over a range of temperatures and frequencies when applied to sea surface emissivity. Their new model produced similar results to other models when applied to cloud liquid absorption (including super-cooled conditions).

It is also possible to apply this technique to observations from ground-based measurements, although this introduces the problems of representativeness due to different sampling scales.

2.7.6 Choice of clear air absorption model

Based on the validation work described above no single absorption model clearly performed best over a wide range of conditions. MPM93 was found to over-estimate the water vapour continuum in warm, humid conditions, while MPM89 and Ros98 generally had low bias compared to observations. Ros98 was selected as the default absorption model for this study on the basis that it has fewer spectral lines in its database, so should provide faster calculations.

2.8 Infrared Radiative Transfer Model

The Radiometrics TP/WVP-3000 system incorporates a zenith-viewing infrared radiometer – the Heimann KT19.85 (§3.1.6). This measures the infrared radiance in the 9.6-11.5 μm window to give information on cloud base temperature. If its observations are to be included in the retrieval a forward model is needed.

A simple forward model defines the infrared brightness temperature, T_{ir} , as the temperature of the lowest level with any cloud. However, this would give sharp transitions in T_{ir} when cloud starts to form if the humidity in a model layer increases. A slightly more sophisticated radiative transfer model is developed here, which accounts for absorption by atmospheric water vapour

and finite extinction in liquid water cloud. This model gives more Gaussian error characteristics and smoother Jacobians, due to less abrupt transitions at cloud boundaries.

Although it is common to express the signal from infrared radiometers as brightness temperatures, the Planck function does not allow the Rayleigh-Jeans approximation to be used, as the radiance is not linear with T_b . Hence, its radiative transfer calculation must be conducted as radiances, starting from zero at the top of the atmosphere. Once the radiance has been propagated down to the surface, it is converted to the infrared brightness temperature, T_{ir} , and a threshold of 223.2 K is applied as this is the lower limit of the instrument's range.

If the measurements of cloud emissivity reported by Chylek and Ramaswamy [1982] are linearly averaged across the 9.6-11.5 μm band, an extinction coefficient of 0.072 m^2/g is suggested. However, when this was applied to forward model T_{ir} from radiosondes or mesoscale model data, it was found to underestimate T_{ir} compared with its observed distribution over the course of several months. Bloemink *et al.* [1999] observed that typical clouds become optically thick at $\lambda=9\text{-}12$ μm when their Integrated Liquid Water (ILW) content exceeds 20 to 30 g/m^2 . If this is interpreted to mean a liquid water cloud has unity opacity when $\text{ILW}=0.03$ kg/m^2 , it implies an extinction coefficient of 33.3 $\text{Np}/\text{km} \cdot (\text{kg}/\text{m}^3)^{-1}$. This extinction coefficient was found to improve the distribution of forward modelled T_{ir} compared to the observations.

Ice cloud is treated in the same way as liquid in the infrared forward model, as unlike microwave frequencies it has a strong influence on T_{ir} .

Although the nominal minimum value of T_{ir} measured with the radiometer is 223.2 K, higher values are often observed in clear skies in warm, humid conditions. This is because it is also sensitive to emission from atmospheric water vapour as a result of its weak emission lines and continuum in the 9-12 μm band. Although it is possible to model these accurately using a line-by-line model, this would be too slow for operational applications, so it has been parametrized empirically here. An extinction coefficient due to water vapour of 0.02 $\text{Np}/\text{km} \cdot (\text{kg}/\text{kg})^{-1}$ has been estimated based on comparisons of time-series of observed T_{ir} with those calculated from the forward model, based on mesoscale model profiles.

2.9 Summary of Radiative Transfer Modelling

This chapter described the radiative transfer model, which forms the core of the forward model used to transform state space to observation space. This is needed for all *physical* retrieval methods, along with estimates of its uncertainties.

The radiative transfer equation and brightness temperatures were first defined. Then the emission mechanisms for atmospheric oxygen, water vapour and cloud and precipitation were discussed to provide a theoretical basis of radiative transfer modelling for both microwave and the thermal infrared. Different models were compared to estimate the uncertainty in the forward model due to spectroscopic errors. Other contributions to the uncertainty were also analysed, including those introduced by approximations made in the fast absorption model, channel averaging and profile discretisation.

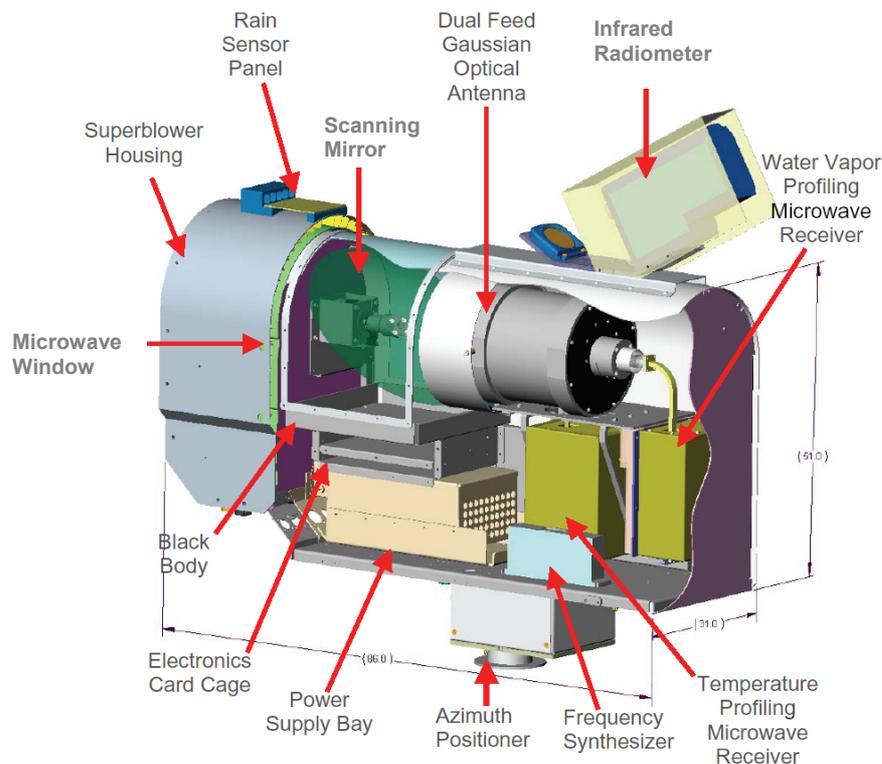
These models have been compared with observations to validate their predictions and assess their relative biases, which will bias the profiles retrieved from them. The absorption model of Rosenkranz [1998] was selected for use in this study.

Chapter 3 Microwave Radiometer Hardware

This chapter describes the microwave radiometer used in this study – the Radiometrics TP/WVP-3000. Its principles of operation, components and calibration are introduced and compared with those of other radiometers. It is important to understand the instrument's calibration accuracy and radiometric noise as these affect the accuracy of the products retrieved from its observations. These aspects of the system's performance are analysed in this chapter, together with practical aspects of its operation. This understanding has helped improve the instrument further during the period of this work. Finally, likely future technological developments are presented and their implications discussed.

3.1 Radiometer Design

The Radiometrics TP/WVP-3000 system includes the microwave radiometer, sensors for atmospheric pressure, temperature and humidity, a rain sensor and an infrared radiometer. When the radiometer views the sky, incoming radiation is transmitted through a microwave window onto a planar scanning mirror and into a Gaussian optics antenna, comprising a dielectric lens, polarising grid (to couple two receiver bands onto a common axis) and two feedhorns. These feed two independent receiver chains, which are mounted in thermally insulated enclosures. The cut-away diagram in Figure 3-1 illustrates the components of the radiometer system, which are described in the following sections.



**Figure 3-1 Cut-away Diagram of Microwave Radiometer TP/WVP-3000 System
[Courtesy of Radiometrics]**

3.1.1 Microwave Window

The microwave radiometer views the atmosphere through a dielectric window. This has now been treated with a hydrophobic coating, which was developed during the course of this study. Together with a stronger blower motor this was found to be effective at preventing the accumulation of liquid water, which would otherwise contribute to the radiometer's signal (§5.5.2) in fog and light rain. The window has a dielectric loss of approximately 0.01 dB (0.2%) at 22-30 GHz [Mike Exner, Radiometrics, personal communication], so its thermal emission will contribute ~0.5 K to the receive signal in these channels. But since its loss is the same in each view, its thermal emission is effectively included in the calibration coefficients.

3.1.2 Antenna System

The antenna comprises all the quasi-optic components in the radiometer's front end: the microwave window, scanning mirror, dielectric lens, polarising grid and corrugated feedhorns. Together they define the antenna response function, which is approximately Gaussian, with the 3dB beamwidth ranging from 6.2° at 22.235 GHz to 2.4° at 58.8 GHz [Radiometrics, 2001]. The instrument's beamwidth is chosen as a trade-off between antenna size (and cost) and resolution. The beamwidth becomes important when viewing low elevation angles (§3.2.1).

3.1.3 Receiver

The incoming radiation from the feedhorns is then mixed with a reference *Local Oscillator* signal from a frequency synthesiser to produce an Intermediate Frequency (IF) signal, ranging from 0-300 MHz. The frequency synthesiser centre determines the centre frequency of the radiation detected instantaneously by the radiometer, which is switched sequentially to allow the radiometer to sample different frequency bands, known as *channels*. Although any number of channels can be selected within the receiver's 22-30 GHz and 51-59 GHz bands, the radiometer has been operated on a fixed set of 12 channels throughout this study, shown in Table 3-1. These values were derived as an optimal set by selecting those frequencies which produced eigenvalues with the maximum information content [Solheim *et al.*, 1996]. There is a delay of 180 ms between sampling each channel, so it takes several seconds to sample a set of 12 channels, which can introduce random noise on the observations (§3.5). The stability of the frequency synthesiser is quoted as 20 ppm. This is equivalent to <100 kHz at 60 GHz and is not expected to contribute a significant uncertainty in the radiative transfer.

Table 3-1 Channel characteristics of Radiometrics TP/WVP-3000

Frequency (GHz)	22.235	23.035	23.835	26.235	30.000	51.250	52.280	53.850	54.940	56.660	57.290	58.800
Beamwidth (FWHM)	6.3°	6.2°	6.0°	5.5°	4.9°	2.5°	2.5°	2.5°	2.5°	2.4°	2.4°	2.4°

Massive amplification is needed to boost the incoming thermal signals to detectable levels. This is provided by low noise IF amplifiers with gain ~ 70 dB, which are highly sensitive and their output is prone to drift with small variations in temperature and power supply fluctuations.

The signal is then passed through bandpass filters, which define the bandwidth of radiation that will be detected by the radiometer. These are common to all channels, which have dual sidebands with 3dB bandwidths of 150 MHz, centred 115 MHz from the channel centre. So the radiometer detects radiation in bands $\pm(40-190)$ MHz each side of the centre frequency. The response is assumed to be uniform over the bandwidth.

The IF signal is converted to an analogue voltage by square law diode detectors, for which output is proportional to their input power. This is then further amplified before being digitised and integrated over a user-defined period – typically ~ 200 ms/channel.

3.1.4 Alternative receiver designs

Other microwave radiometers currently available, such as HATPRO [Rose *et al.*, 2005] sample a fixed set of channels simultaneously by splitting the incoming signal between a discrete set of bandpass filters, referred to as a *filterbank*. This allows faster observing sequences at the expense of a greater component count. In this way they can also exploit recent technological advances allowing *direct detection* of incoming radiation (after amplification), without the need to mix with a reference local oscillator.

3.1.5 Surface sensors

The Radiometrics TP/WVP-3000 incorporates ambient pressure, temperature and relative humidity sensors, mounted within the blower unit to ensure good ventilation and to minimise the effect of solar radiation. It also includes a rain sensor to detect precipitation and allow it to be flagged in subsequent retrievals.

3.1.6 Infrared radiometer

An optional infrared radiometer (Heimann KT19.85) is mounted on top of the instrument, which views the zenith in a small mirror. This uses a lens to produce a 5° beamwidth, which approximately matches that of the microwave radiometer's channels in the water vapour band. This instrument detects radiation emitted in the $9.6-11.5 \mu\text{m}$ band, which is used to provide information on the cloud-base temperature.

3.2 Observations

The radiometer's scan mirror can be rotated to allow the receivers to detect incoming radiation at any arbitrary elevation angle (with a 0.45° resolution). It can also view an internal reference target (§3.3.1). A series of views can be scheduled sequentially under user control.

3.2.1 Elevation scan

The radiometer nominally views in the zenith direction, but can also measure a series of elevation angles, usually symmetric about zenith. The elevation scans can be used for tip curve calibration of the water vapour channels (§3.3.4) or to provide extra observations to improve the retrieval of temperature profiles (§5.4). However, observations away from zenith are subject to several additional sources of bias, described below, which should be corrected.

At high zenith angles, the brightness temperature changes rapidly, and becomes sensitive to small errors in the beams' alignment or the beamwidth over which radiation is integrated. This is illustrated in the schematic in Figure 3-2. The finite beamwidth can be accounted for by artificially increasing the zenith angle, following Appendix 1 of Han and Westwater [2000]. The vertical gradient of refractive index and Earth's curvature cause small deflections of the radiometer's beam from its nominal path, which can become significant as the atmosphere becomes optically thick and the radiometer is sensitive to emission over extended paths [Han and Westwater, 2000].

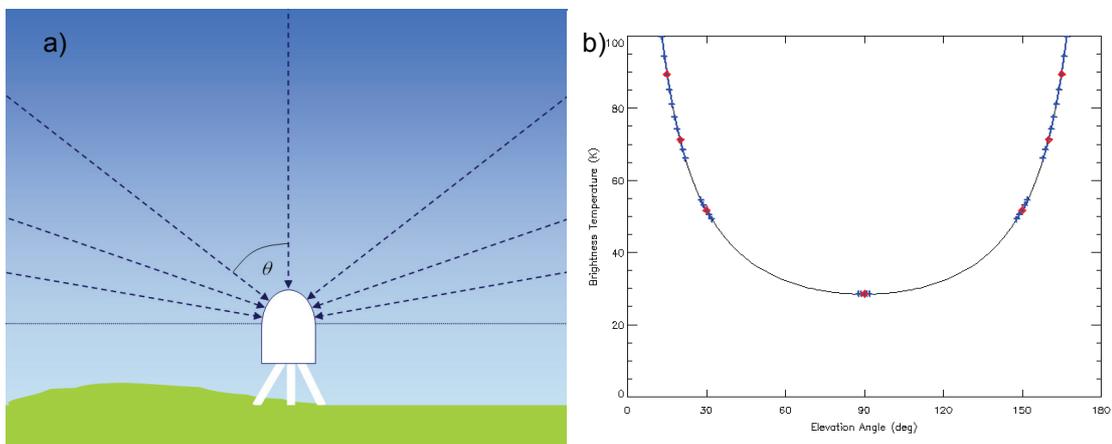


Figure 3-2 a) Schematic diagram of observations at 7 elevation angles, ϕ , and b) example brightness temperatures at different ϕ for 30.000 GHz channel, highlighting beamwidth of channel (blue crosses) and nominal angle (red diamonds).

Optically thin channels and those near the edge of the oxygen absorption band (51-52 GHz) are very sensitive to changes in elevation angle. If the radiometer beam is systematically misaligned in elevation, biases will be introduced in these observations at high zenith angles.

Examination of the brightness temperatures measured in nominally symmetric elevation angles on opposite sides of zenith show systematic differences as much as 6 K. This is most likely to be caused by misalignment of the instrument by $\sim 0.5^\circ$ – comparable to the accuracy with which it can be levelled manually on installation. To reduce the influence of this, brightness temperatures from angles on opposite sides of zenith are co-averaged. This reduces the uncertainty caused by instrument misalignment to < 0.1 K. However, it relies on the assumption of horizontal homogeneity, which will be investigated in §5.4.1.

Closer examination of the brightness temperatures measured in the 51-52 GHz channels shows they are not consistently asymmetric – they jump between 2 quantised levels several kelvin apart. This was identified as being due to different alignment triggering used to control the stepper motor, which has 0.45° resolution. This suggests that finer control is needed to ensure elevation angles are repeatable and accurate.

The radiometer's antennae are not 100% efficient and are also sensitive to emission from outside the main Gaussian beam in the form of antenna *sidelobes*. These can introduce bias at low elevation angles due to emission from the surface, which can have ~ 200 K contrast to the sky T_b at low frequencies. Analysis by Hewison and Gaffard [2003] suggested that as much as 2% of the power received by the 22.235 GHz channel at $\varphi=20^\circ$ may have been emitted from the surface, which would bias the T_b measured here by ~ 4 K. The antenna's beam efficiency improves at higher frequencies, so this bias would reduce to ~ 0.6 K at 30 GHz.

No mirror is perfect. As the scan mirror rotates, the polarisation angle incident on it also rotates. Because the mirror's reflectivity varies with polarisation angle, this can bias the brightness temperatures measured in all views [McGrath and Hewison, 2001], including those of the reference target. Hewison and Gaffard [2003] showed this does not bias the zenith view, as the polarisation angle incident on the mirror is the same as in the opposite view of the reference target. However, this effect can bias T_b s measured in other views by ~ 0.2 K, which can be significant in the tip curve calibration and retrievals from elevation scans.

3.2.2 Observing Cycle

The total integration time of each channel is user-configurable. This allows a trade-off between low radiometer noise for longer integrations and low atmospheric variability, which will be discussed in §7.1.5. Higher sampling rates may allow more realistic representation of small-scale structure in the atmosphere. Typically, the radiometer may be configured to integrate each of 12 channels for 170 ms, sequentially over a period of 7 s. In this way, views at 5 elevation angles and a view of the internal black body calibration target can be obtained in a 1 minute observing cycle.

The instrument is usually orientated to scan in the North-South plane. In this way views at low elevation angles can be screened to prevent contamination by the emission from the Sun using a simple time of day filter: avoiding the use of these views within 1 hr of local noon.

3.3 Calibration of Microwave Radiometer

To convert the radiometer's output voltages, V_b , when viewing the sky into brightness temperatures, T_b , they must be compared to those measured viewing a reference scenes of known radiances to determine the system's gain, G , and offset, a . In the first instance the linear relationship is assumed:

$$V_b = GT_b + a \quad (3.1)$$

This calibration against references must be performed sufficiently frequently that gain fluctuations do not contribute significantly to the overall error budget. References are provided by injecting signals from noise diodes in the receivers' input and viewing an internal reference target.

In a recent revision of the operating software, Radiometrics included an additional term in (3.1) to account for the nonlinear response of the radiometer's detectors. This has a maximum magnitude $\sim 1\text{K}$ for $T_b \sim 200\text{ K}$, but this correction is not expected to change with time, so it will only contribute to the bias, not the random error. This change was applied to data gathered during the later part of the data set analysed here.

3.3.1 Noise Diodes

Noise diodes generate stable white noise with Gaussian characteristics and can be switched on and off rapidly, requiring a settling time of only $\sim 1\text{ ms}$. One noise diode is coupled into the input of each receiver to inject a reference signal, which can be used to control the radiometer's drift by alternating views of the scene with the noise diode active and inactive at high frequency ($\sim 100\text{ Hz}$). The noise diode's contribution to the signal is expressed in terms of an equivalent brightness temperature, T_{ND} , but does not provide an absolute reference and must be calibrated against an external reference (§3.3.4 and §3.3.5). In fact, T_{ND} includes a number of terms related to the loss of the receiver front end. This loss is dependent on the temperature of these components, so operationally a correction is applied based on the black body temperature and coefficients derived by Radiometrics during testing in an environmental chamber. The noise diodes are set up to produce typical values of $T_{ND} \sim 200\text{-}300\text{ K}$.

3.3.2 Internal black body reference target

The other reference scene that can be viewed regularly is the internal black body target. This comprises a microwave absorber, for which the emissivity is assumed to be one – so its brightness temperature is the thermodynamic temperature measured with thermometers, T_{BB} .

The target's emissivity is not expected to introduce significant bias because, if the emissivity is 0.99 for example, 1% of the power received by the radiometer will be that emitted by the antenna reflected by this surface and its temperature is likely to be within 3 K of the black body. Otherwise, the accuracy of the brightness temperature of the black body target is limited by that of the temperature sensors: 0.2 °C [Radiometrics, 2001], which has been taken as the uncertainty in T_{BB} . This has a direct and linear impact on the final brightness temperatures.

The target's temperature is monitored by two sensors, for which the mean has a random error of 0.05 K. The target is not thermally controlled, and is typically ~10.5 K warmer than ambient with an r.m.s. rate of change over one year at Camborne of 1.6 K/hr. If its temperature is sampled within 1 min of the radiometer viewing it, temperature gradients should introduce negligible errors.

3.3.3 System Equations

Together, the noise diodes and black body target provide reference brightness temperatures against which the radiometer's gain and offset can be calculated to calibrate its output voltage by the *system equations*:

$$T_b = T_{BB} - \frac{(V_{BB} - V_b)}{G} \quad (3.2)$$

where

$$G = \frac{(V_{BB+ND} - V_{BB})}{T_{ND}} \quad (3.3)$$

where T_b is the scene brightness (antenna) temperature,
 T_{BB} is the temperature of the black body target,
 V_{BB} , V_{BB+ND} and V_b are the voltages measured by the radiometer when viewing the black body, the black body plus the noise diode and the scene, respectively,
 G is the gain, and
 T_{ND} is the brightness temperature of the noise diode.

Because T_{ND} is added to the reference scenes, this system relies on extrapolating the calibration coefficients, which are derived for scenes with $T_b \gg T_{BB}$, while for typical scenes $T_b \sim 20-300$ K. So it is important to estimate them as accurately as possible as this extrapolation can tend to amplify calibration errors.

3.3.4 Tip curve calibrations

The *tip curve* [Han and Westwater, 2000] is one method to derive an accurate reference scene temperature against which to calibrate a radiometer. This technique uses radiance

measurements over a range of elevation angles calculated using an initial estimate of the calibration coefficients to derive an accurate estimate of the true radiance at zenith. This is then used to update the calibration coefficients. The underlying assumptions are that the atmosphere is horizontally stratified and optically thin, such that its opacity is a linear function of the slant path. In these conditions, the radiative transfer equation (2.7) can be simplified to the following expression for the down-welling brightness temperature, T_b , at zenith angle, θ :

$$T_b(\theta) = T_{MR} (1 - e^{-\tau \sec \theta}) + T_{CMB} e^{-\tau \sec \theta} \quad (3.4)$$

where T_{MR} is the mean radiative temperature of the atmosphere,
 T_{CMB} is the effective brightness temperature of the cosmic microwave background,
 τ is the opacity at zenith,
 θ is the zenith angle.

The opacity over a range of zenith angles is calculated from measured T_b s, using equation (3.4), based on an initial calibration. For an optically thin, horizontally stratified atmosphere, the opacity is expected to increase linearly with slant path, $\sec \theta$.

Every few minutes the radiometer can take measurements at 5 or 6 angles, symmetric around zenith. The measurements are then fitted to the theoretical function to derive an estimate of the true zenith brightness temperature. If the fit is deemed satisfactory, equations (3.2) and (3.3) are then inverted to adjust T_{ND} . The values of T_{ND} for successful tips are logged and a weighted average of these is used to provide the 'current' value used in equations (3.2) and (3.3). In practice tips are usually only successful in cloud-free conditions.

Systematic biases in the tip curve calibration were analysed in detail by Hewison and Gaffard [2003], following the methods described by Han and Westwater [2000]. The magnitudes of each term and uncertainties are summarised here.

The current tip curve calibration uses a simple method to estimate T_{MR} as constant for each channel, which can bias the calibrations at high or low temperatures. Although this had little impact on the water vapour channels, T_{MR} was found to dominate the error budget when attempting to apply tip curve calibrations to higher frequency channels because the atmosphere becomes optically thick as we approach the centre of the oxygen band. The vertical gradient of refractive index and the Earth's curvature have negligible impact on T_b s below 50 GHz for $\theta < 70^\circ$. The correction for the radiometer's finite beamwidth is small (~ 0.1 K) and has low uncertainty for $\theta < 60^\circ$. Similarly, averaging T_b s measured in opposite views reduces the impact of misalignment of the radiometer's elevation angle on the tip curve to negligible levels. However, there are two important factors mentioned in §3.2.1 that were

not included in the analysis of Han and Westwater [2000]. One is due to changes in scan mirror's reflectivity as it rotates, which can introduce a small bias in the tip curve ~ 0.2 K. The other is the systematic bias introduced by exposure of antenna sidelobes to emission from the surface at high zenith angles. This can introduce a bias of ~ 0.5 K for $\theta < 60^\circ$, with a large uncertainty, which dominates the error budget of the tip curve calibration. If higher zenith angles are included in the tip curve, the biases in the lowest frequency channels (and their uncertainties) increase exponentially. It is recommended, therefore, that tip curves are applied only to zenith angles $\theta \leq 60^\circ$.

3.3.5 Liquid Nitrogen Calibrations

The tip curve technique can only be applied to channels where the atmosphere is optically thin. So, another external reference is needed to calibrate T_{ND} for the oxygen-band channels: a black body at cryogenic temperatures. For practical reasons these calibrations cannot be performed as frequently as tip curves, so allowing greater T_{ND} drifts. However, the typical T_b of these channels is much higher (shown in Table 3-2), so their calibration is less sensitive to errors in T_{ND} . Radiometrics recommend calibration every "several months" or after transporting the radiometer.

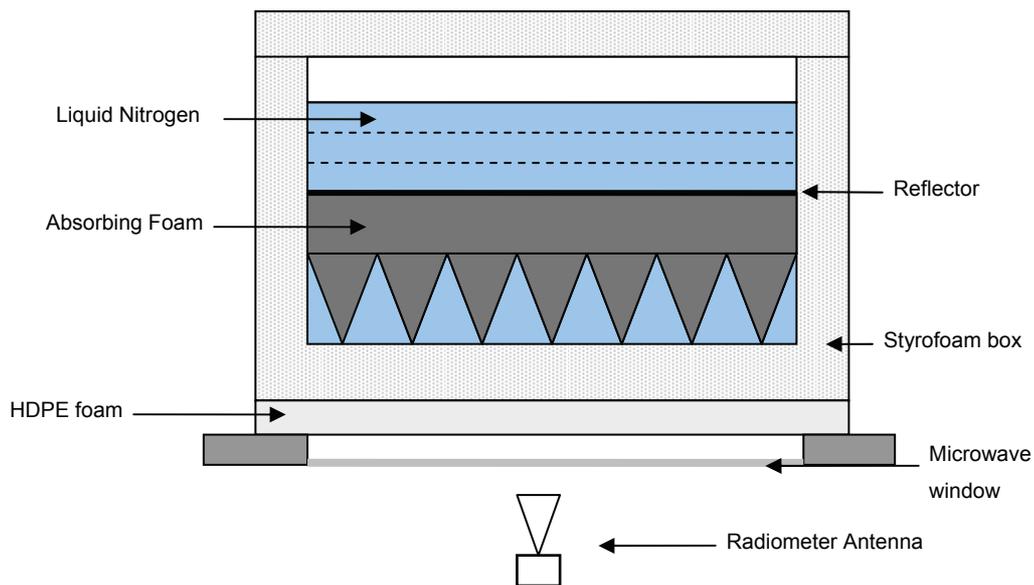


Figure 3-3 Cryogenic Target used to Calibrate T_{ND} for 51-59 GHz channels.

The manufacturer supplied a calibration target, illustrated in Figure 3-3, which comprises a box of expanded polystyrene foam (*Styrofoam*), containing a permeable microwave absorber. This is filled with liquid nitrogen and placed on top of the radiometer, so it views emissions from the absorber through the base of the Styrofoam box, which has a low loss at microwave frequencies. The cryogen is assumed to be isothermal at the boiling point of nitrogen at atmospheric pressure, incremented a little to account for the hydrostatic pressure due to a

20 cm column of liquid nitrogen. This provides a black body at a known low temperature: 77.55 K at 1013 hPa, but varying by ± 0.2 K over the typical range of pressures experienced at Camborne. The target has been found to be sufficiently stable when used in a relatively dry atmosphere to delay the formation of condensation on the outside of the target and allow it to be used for periods of ~ 30 min. T_{ND} is adjusted after each liquid nitrogen calibration so as to force the brightness temperatures measured by the radiometer to match the expected values over this period.

Each component of the liquid nitrogen calibration target was also analysed in detail by Hewison and Gaffard [2003] to calculate its contribution to the uncertainty in its brightness temperature by emission, absorption and scattering. These are summarised below.

The dissimilar dielectric properties of liquid nitrogen and Styrofoam cause some of the radiation incident at the interface to be reflected. However, it is believed that liquid nitrogen gradually penetrates between the beads of polystyrene and breaks down the interface. This effect reduces the magnitude of this term by an unknown amount. The brightness temperature of the target has been observed to reduce by ~ 1 K during the first few minutes after filling [Hewison and McGrath, 2001]. It has been assumed that this term has a negligible contribution for low-density Styrofoam targets (1.75 lb/ft^3). However, there remains a large uncertainty on this ($\sim \pm 0.8$ K), which dominates that of the liquid nitrogen calibrations. For this reason, the same procedure should always be followed to ensure repeatable calibrations.

Emission from the base of the Styrofoam box was found to increase with frequency from 0.4 K at 22.235 GHz to 1.2 K at 58.8 GHz and has a large uncertainty of $\pm 50\%$. The beads of polystyrene are estimated to be 2 mm diameter – i.e. $D \sim \lambda$ – the regime in which volume scattering is most efficient. However, the low dielectric constant of polystyrene foam results in scattering producing a small but significant contribution to the total extinction. A 6 mm layer of high-density polyethylene foam (HDPE) is attached to the underside of the Styrofoam target to act as a thermal barrier to prevent the formation of condensation. This has very similar dielectric properties to Styrofoam and air, so no significant reflections are expected at these interfaces. The low loss tangent of HDPE foam means its emission is also negligible.

3.4 Calibration Bias

Each term in the radiometer system equations (3.2) and (3.3) is potentially subject to systematic errors, which can bias the observed brightness temperature and hence the profiles retrieved from them. In this section the known sources of bias are modelled and compared with the bias observed in a validation exercise, comparing radiometer measurements with forward-modelled radiosondes.

3.4.1 Modelled calibration bias

A full review of all known mechanisms that can introduce bias in the calibration was conducted by Hewison and Gaffard [2003], including estimating the magnitude of each term and its uncertainty. Data from liquid nitrogen and tip curve calibrations were used together with theoretical modelling of the processes that can introduce biases. Table 3-2 shows the total bias introduced in observations at typical scene brightness temperatures for the current configuration of liquid nitrogen and tip curve calibrations. It also shows the uncertainty with which each bias is known. These uncertainties contribute directly to the overall error budget of the calibrated brightness temperatures.

Table 3-2 Total bias introduced to typical scene brightness temperatures by all known mechanisms in Liquid Nitrogen and Tip Curve Calibrations and uncertainty.

The mean bias found in comparison (§3.4.2) of observations with forward-modelled radiosondes is also shown for three absorption models: MPM89, MPM93 and Ros98.

Standard deviation of the differences between observations and model is also shown.

Frequency (GHz)	Typical T_b (K)	Total Bias on Typical T_b from LN2 calibration (K)			Total Bias on Typical T_b from Tip Curve calibration (K)			Mean Bias	Mean Bias	Mean Bias	S.D.
		Bias	±	Uncertainty	Bias	±	Uncertainty	(Obs- MPM89) ◇ (K)	(Obs- MPM93) × (K)	(Obs- Ros98) + (K)	(Obs- MPM89) (K)
22.235	27.5	1.12	±	0.80	0.76	±	0.52	1.81	0.72	1.92	1.30
23.035	27.0	1.15	±	0.82	0.62	±	0.40	1.35	0.23	1.43	1.12
23.835	24.0	1.22	±	0.84	0.60	±	0.40	1.44	0.36	1.48	0.94
26.235	17.1	1.40	±	0.92	0.44	±	0.29	1.00	0.01	0.93	0.51
30.000	15.0	1.54	±	1.01	0.31	±	0.21	0.78	-0.35	0.57	0.46
51.250	105.5	0.80	±	1.06	0.26	±	0.35	0.57	0.91	0.18	0.66
52.280	148.9	0.55	±	0.89	0.35	±	0.78	-1.08	-0.08	-1.39	0.42
53.850	248.6	0.22	±	0.38	0.52	±	1.52	0.90	0.16	-0.08	0.33
54.940	278.7	0.11	±	0.24	N / A			1.21	0.89	0.94	0.33
56.660	283.4	0.09	±	0.22	N / A			1.36	1.29	1.31	0.64
57.290	283.8	0.09	±	0.22	N / A			1.76	1.72	1.74	0.44
58.800	284.1	0.10	±	0.22	N / A			1.18	1.22	1.24	0.37

Table 3-2 shows that the tip curve introduces lower uncertainties than the liquid nitrogen calibration for the lowest frequencies. At the higher frequencies where the atmosphere is optically thick, liquid nitrogen is better, as the tip curve becomes unfeasible. Table 3-2 also includes the mean bias found between the observations and three different absorption models in the following section.

3.4.2 Validation of observations/model

As discussed in §2.7, radiometer observations can be used to validate radiative transfer models (RTMs) or vice-versa. This is illustrated here by comparing the brightness temperatures (T_b s) measured by the radiometer with those forward modelled from co-located

radiosondes in clear skies. This can highlight biases in the radiometer, radiosondes or RTM, which can affect retrievals and their validation.

Hewison *et al.* [2003] analysed the first year of radiometer observations from Camborne. The zenith T_b s measured within 10 minutes of the launch of each radiosonde were compared to those modelled from the high-resolution radiosonde profiles using different absorption models. The following absorption models were used in this study: MPM87, MPM89, MPM93 and Ros98, which are described in §2.2.5. This comparison was only performed in clear sky conditions, as radiosondes do not measure liquid water content, which would be necessary to model cloudy cases.

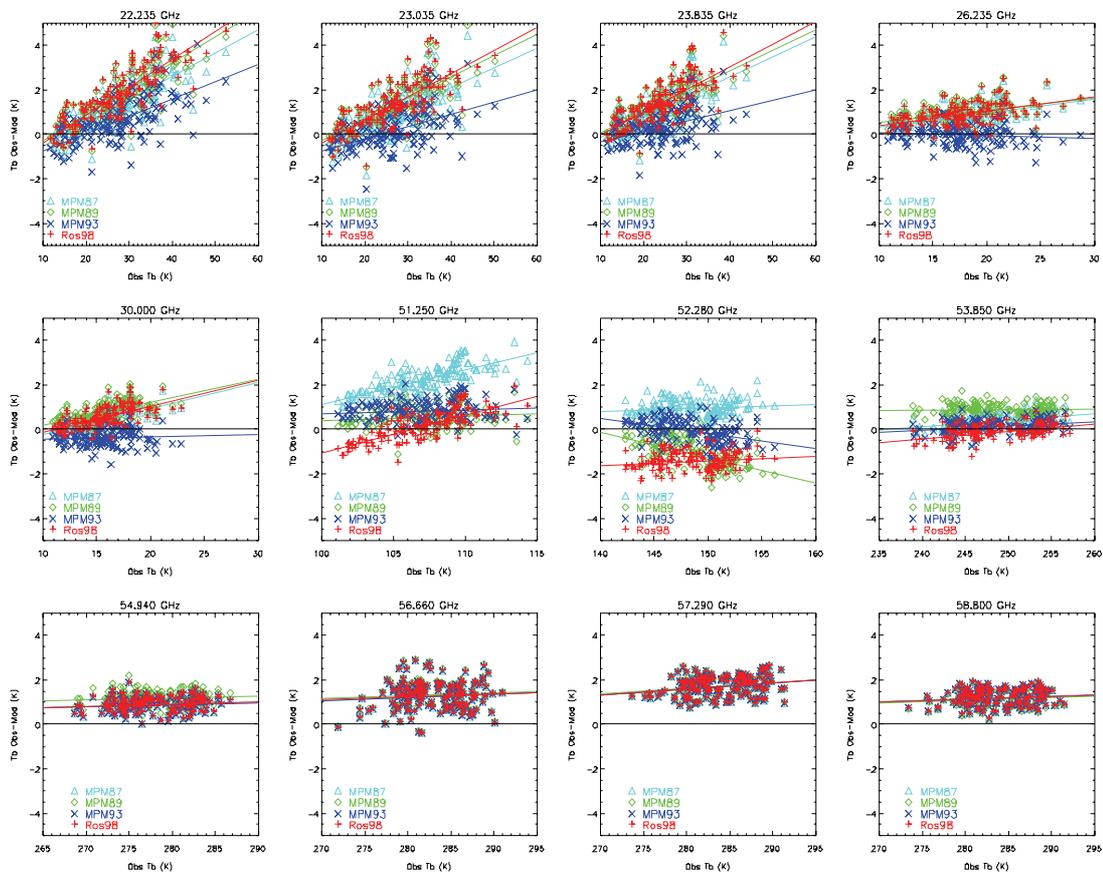


Figure 3-4 Observed-Modelled zenith brightness temperatures. Each point is one radiosonde, modelled using MPM87(Δ), MPM89(\diamond), MPM93(x), Ros98(+). Solid lines show the linear regression through all point for each model (same colour). Based on 145 RS80H radiosondes with $\leq 1/8$ low cloud, 22/2/02-3/2/03, Camborne, UK.

Figure 3-4 shows the difference between observed and modelled brightness temperatures, plotted against the observed value. The comparison of the channels dominated by water vapour emission proved problematic, due to the higher variability of atmospheric humidity, which introduces a large scatter on the data. Despite this, the lowest frequency channels

show a strong slope, with all models showing a significant bias in humid conditions. The RS80H radiosondes used then at Camborne are known to have a dry bias due to sensor contamination and solar heating [Nash *et al.*, 1995]. The mean maximum humidity measured in low cloud during this period was 97.0 %RH. When the vapour pressure in the profiles input to the RTMs was increasing by a factor of 1.03, the difference between observations and model decreased, but a positive bias remained for all models. MPM93 then remained the closest model to the observations, although it overestimates the strength of the water vapour continuum at 26-30 GHz.

There was a consistent, positive bias of 1-2 K at the highest frequencies. This has been identified as an instrument bias, and has subsequently been reduced by modification of the control software. The channels at intermediate frequencies (30-52 GHz) show significant differences between the models, where MPM93 provides the best fit to the observations, although there is a large systematic uncertainty (~ 1 K) in the observations at 51.250 GHz mostly due to the calibration. There remains uncertainty in the accuracy of the absorption models at the edge of the oxygen absorption band near 51-53 GHz due to their sensitivity to the line coupling parameters [Boukabara *et al.*, 2004].

3.5 Radiometric Noise

Random noise on oxygen band channels is the single most important parameter of the radiometer in determining the accuracy of the temperature profile retrievals (§4.4.1). Atmospheric fluctuations also introduce random “noise”, which will be considered in §4.4.1.3, so although noise can be improved as a trade-off with integration time, the optimum result is not necessarily the same for each channel (§8.1). It is also useful to minimise the random noise on the observations in order to better evaluate biases. There is a contribution to the radiometer’s noise from every component of the receiver and antenna system. These were analysed by Hewison and Gaffard [2003] to allow optimisation of radiometer noise and are summarised here.

3.5.1 Modelled radiometric noise

The radiometric resolution, ΔT_{\min} , is defined as the minimum difference in scene brightness temperature that is resolvable at the receiver’s output. This is equivalent to the standard deviation of the brightness temperatures measured by the radiometer while viewing a thermally stable scene, excluding contributions from noise on the calibration coefficients. The radiometric resolution depends on the system noise temperature, T_{SYS} , the pre-detection bandwidth, B , and the integration time, τ , and gain stability, according to equation (3.5) [Ulaby *et al.*, 1981]:

$$\Delta T_{\min} = T_{\text{sys}} \left[\frac{1}{B\tau} + \left(\frac{\Delta G}{G} \right)^2 \right]^{1/2} \quad (3.5)$$

where ΔG is the change in gain between observations not accounted for by the calibrations.

However, ΔT_{\min} does not include noise on the calibration coefficients that affect the final T_b s. This can be evaluated by differentiation of the system equations (3.2) and (3.3) to yield:

$$\varepsilon T_b^2 = \varepsilon T_{BB}^2 + 2 \cdot \Delta T_{\min}^2 + \left[\frac{(T_{BB} - T_b)}{T_{ND}} \right]^2 \cdot (2 \cdot \Delta T_{\min}^2 + \varepsilon T_{ND}^2) \quad (3.6)$$

where εT_{BB} is the uncertainty in the temperature of the black body, T_{BB} ,

εT_{ND} is the uncertainty in brightness temperature attributed to the noise diode, T_{ND} ,

εT_b is the uncertainty on the scene brightness temperature, T_b .

Note that εT_b derived from equation (3.6) contains ΔT_{\min} terms multiplied by $\sqrt{2}$. This is due to the implementation of the system equations (3.2) and (3.3), which are applied independently to each observation cycle, differencing the views of the sky and the black body to account for the radiometer's offset. The overall noise on brightness temperatures, εT_b , could be reduced substantially (by as much as a factor of $\sqrt{2}$) if the gain and offset could be averaged over a number of observation cycles. This is only possible if the radiometer is calibrated more frequently than the typical timescale for their variation.

The most difficult term to quantify is εT_{ND} because of the large noise on estimates of T_{ND} provided by tip curve calibrations and the irregularity and variability due to changes in the design of the liquid nitrogen target. Hewison and Gaffard [2003] analysed a series of tip curve calibrations in clear skies to fit a non-stationary model of the evolution of T_{ND} to estimate εT_{ND} as a function of the typical time between calibrations. Table 3-3 shows this term dominates the noise on the brightness temperatures when applying liquid nitrogen calibrations to the water vapour channels, which typically have low T_b s and need to extrapolate the calibration points furthest. However, there is a large uncertainty on this term.

Table 3-3 shows the contributions to the total uncertainty on typical zenith brightness temperatures due to each term affecting calibrations using liquid nitrogen and tip curves. These should be compared to the radiometric noise, shown in Table 3-4. It shows that to minimise the random noise on T_b , the tip curve should be used for the low frequency channels that see lower T_b s, while liquid nitrogen calibrations should be used for channels above ~ 50 GHz. The optimum interval between these calibrations can also be determined from an extension of this analysis. It is clear that the average interval of 16 hr between tip curves does not cause noise diode drifts to dominate the calibration noise, while it does for average interval of 3 months since the most recent liquid nitrogen calibration.

**Table 3-3 Analysis of Contributions to Noise on Typical Zenith Brightness
Temperatures from Liquid Nitrogen and Tip Curve Calibrations**

Freq. (GHz)	Contribution to uncertainty on Tz using LN2 Calibration					Contribution to uncertainty on Tz using Tip Curve Calibrations						
	Black Body Noise (K)	LN2 Noise (K)	Tnd Drift (K)	Radi- meter Noise (K)	Total (K)	Black Body Noise (K)	Atmos. Noise (K)	Noise Diode Noise (K)	Tnd Drift (K)	Radi- meter Noise (K)	Tmr Noise (K)	Total (K)
22.235	0.06	0.09	1.07	0.25	1.10	0.06	0.24	0.16	0.13	0.09	0.26	0.47
23.035	0.06	0.09	0.63	0.06	0.64	0.06	0.24	0.13	0.11	0.06	0.25	0.45
23.835	0.06	0.09	0.67	0.05	0.67	0.06	0.22	0.12	0.10	0.05	0.22	0.40
26.235	0.06	0.17	1.10	0.07	1.11	0.06	0.19	0.08	0.12	0.07	0.15	0.31
30.000	0.06	0.16	0.61	0.11	0.64	0.06	0.19	0.09	0.13	0.11	0.12	0.31
51.250	0.06	0.08	0.56	0.08	0.57	0.06	0.17	0.12	0.08	0.08	1.09	1.55
52.280	0.06	0.07	0.44	0.09	0.45	0.06	0.16	0.11	0.06	0.09	1.55	2.20
53.850	0.06	0.02	0.14	0.06	0.16	0.06	0.10	0.10	0.02	0.06	2.62	3.70
54.940	0.06	0.01	0.05	0.14	0.15	0.06	0.10	0.12	0.01	0.14	2.81	3.98
56.660	0.06	0.01	0.04	0.02	0.05	0.06	0.19	0.09	0.01	0.02	2.88	4.08
57.290	0.06	0.01	0.04	0.01	0.04	0.06	0.21	0.08	0.01	0.01	2.90	4.10
58.800	0.06	0.01	0.04	0.00	0.04	0.06	0.12	0.10	0.01	0.00	2.90	4.10

3.5.2 Measured radiometric noise

It is also possible to evaluate the radiometric resolution of the whole system, including any noise introduced by the calibration process, εT_b . This can be estimated as the standard deviation of the brightness temperature measured while viewing a stable reference scene, as shown in Table 3-4 for views of the ambient black body reference target and a liquid nitrogen target over the same 30 min period. All the channels have noise levels between 0.1-0.2 K, except for the 57.29 GHz channel, which is consistently noisier. There was no statistically significant difference between the variance measured with the two scenes, which shows that noise diode fluctuations do not contribute significantly to the radiometer noise on short timescales (<~1 hr).

**Table 3-4 Mean and Uncertainty of Standard Deviations of Brightness Temperatures
measured viewing Ambient (Amb) and Liquid Nitrogen (LN2) reference scenes**

Frequency (GHz)	22.235	23.035	23.835	26.235	30.000	51.250	52.280	53.850	54.940	56.660	57.290	58.800
LN2 S.D.	0.17	0.12	0.11	0.13	0.21	0.18	0.15	0.17	0.17	0.18	0.55	0.17
±	0.03	0.02	0.02	0.03	0.04	0.04	0.03	0.03	0.03	0.04	0.11	0.03
Amb S.D.	0.18	0.18	0.10	0.13	0.20	0.22	0.17	0.13	0.21	0.18	0.55	0.16
±	0.04	0.03	0.02	0.03	0.04	0.04	0.03	0.02	0.04	0.03	0.11	0.03
Ave S.D.	0.18	0.15	0.11	0.13	0.20	0.20	0.16	0.15	0.19	0.18	0.55	0.17
±	0.02	0.02	0.01	0.02	0.03	0.03	0.02	0.02	0.03	0.02	0.08	0.02

3.5.3 Radiometric noise covariance

The radiometric noise is an important contribution to the *Observation Error Covariance Matrix* (§4.4.1), which is needed in variational assimilation of its data. It can be estimated from the same time-series of calibrated brightness temperatures. The variance-covariance matrix, \mathbf{E} is a symmetric ($n \times n$) matrix, for which the diagonal elements are the sample variances of n variables, and the other elements are the covariances among the variables [Wilks, 1995], defined by:

$$\mathbf{E} = (k - 1)^{-1} \mathbf{y}'^T \mathbf{y}' \quad (3.7)$$

where \mathbf{y}' is the ($k \times n$) matrix of brightness temperature anomalies from the mean of k observations in n channels, and T is the matrix transpose operator.

The radiometric noise covariance matrix, \mathbf{E} , was estimated while viewing a liquid nitrogen target over 30 min. It is plotted as an image in Figure 3-5, with each element represented as a square, shaded according to the colour scale. This includes 15 channels – the 12 channels of the microwave radiometer, followed by its ambient temperature and humidity sensors (converted to $\ln q$ – see §4.2.2) and the infrared brightness temperature. Figure 3-5 shows \mathbf{E} is an approximately diagonal matrix, for which the diagonal terms are given in Table 3-4.

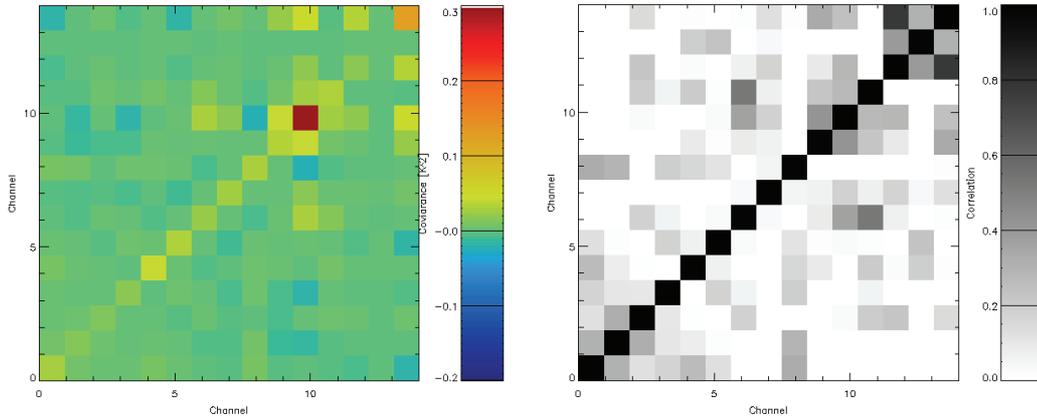


Figure 3-5 Covariance (left) matrix of radiometric noise estimated viewing LN2 target. Correlation (right) Matrix calculated from same data set, showing uncorrelated noise. First 12 channels are Radiometrics TP/WVP-3000 brightness temperatures [K^2]. Last 3 channels are ambient temperature [K^2], total water [$\ln q^2$] and IR T_b [K^2].

3.5.4 Effect of calibration drift

It is possible to demonstrate the effect of calibration drift by comparing periods of data collected by the radiometer with and without regular calibration against views of the ambient reference target. This can be achieved by calculating the *Time Structure Function* (TSF) [Kitchen, 1989], defined here as the r.m.s. difference (divided by $\sqrt{2}$) between brightness temperatures re-sampled over different intervals, τ , from an extended time-series, $\mathbf{y}(t)$:

$$TSF(\tau) = \frac{1}{\sqrt{2}} \left\langle [\mathbf{y}(t+\tau) - \mathbf{y}(t)]^2 \right\rangle^{1/2} \quad (3.8)$$

This is closely related to the *Allan Variance* [Allan, 1966]. It allows the variability of stochastic processes to be quantified over specific timescales and can be used to analyse the stability of microwave radiometers [Vasić *et al.*, 2005].

TSFs have been evaluated for two data sets of 58.8 GHz observations in Figure 3-6, each covering 1106-1306 UTC over 11 day periods in Summer 2005, during CSIP. The TSF starts to increase after 100 s without viewing the ambient black body target, due to calibration drift. Similar results were obtained using the 56.66 and 54.94 GHz channels. Therefore, it is recommended that the observation cycle includes calibrations against the internal black body at least every 60 s.

It is also recommended that the calibration coefficients are averaged over several minutes to reduce their impact on the noise of the calibrated brightness temperatures. In the current configuration, this accounts for half the variance observed in T_b . Exploiting the calibration immediately following the observation would also improve its representativeness, effectively halving the time available for the radiometer's calibration to drift.

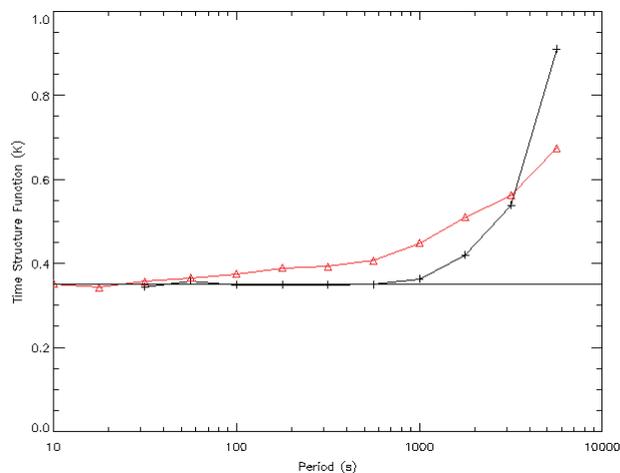


Figure 3-6 Time structure functions of 58.8 GHz T_b over 1106-1306 UTC during CSIP: 21/7 – 31/7/05 with calibration against ambient black body every 30 s (black), 26/6 – 7/7/05 without calibration, but sampled every 12s (red) [multiplied by $\sqrt{2}$]. Horizontal line indicates white noise level on calibrated brightness temperatures.

On periods longer than 1000 s, the TSF starts to increase as atmospheric variability starts to dominate over radiometric noise. It is recommended that the brightness temperatures are not averaged for longer than this period before retrieving the profile. This break point occurs at much shorter time scales for lower frequency channels, which are sensitive to humidity and

cloud. Even by 54.94 GHz, this break point occurs ~300 s. At lower frequencies, the observations are not noise limited even on the shortest timescales, so the above technique cannot be used to measure the calibration drift. This is discussed in more detail in §4.4.1.3.

3.6 Radio Frequency Interference (RFI)

Interference poses a serious risk to the viability of future operations of microwave radiometers. Passive microwave radiometers are designed to detect thermal signal levels and can be extremely sensitive to transmissions within their channels' passbands (or even from adjacent bands if inadequately filtered). Interference can also be picked up in other parts of the radiometer, if inadequately screened – particularly in the intermediate frequencies (IF). Although no RFI was experienced during the trials of the Radiometrics TP/WVP-3000, and its electromagnetic compatibility with other systems required for future integrated profiling systems has been demonstrated, it is likely to become more problematic in future.

As the number of microwave radiometers in a network increases so does the risk of compromising the quality of their observations due to interference, as it will be more difficult to site them at “quiet” locations. It will also become more difficult to monitor applications from other users of the electromagnetic spectrum to establish nearby transmitters in, or adjacent to, the passbands of the radiometers. There is also increased commercial pressure on the electromagnetic spectrum from mobile communications, fixed data links, automotive collision avoidance radars, and low power, spread-spectrum applications, such as WiFi.

The conflicting demands on the spectrum are managed by national and international regulation. Several of the channels used by microwave radiometers are in narrow bands designated as exclusively for passive applications (23.60-24.00 GHz, 52.6-55.78 GHz) [Ofcom, 2004]. However, others are in shared allocation (22.21-23.0 GHz, 55.78-59.3 GHz), or completely unprotected bands (e.g. 23.0-31.3 GHz, 50.4-52.6 GHz). Although the centre of the 60 GHz oxygen band is shared, strong atmospheric absorption here provides a good level of protection. It is recommended that, in future, the channels currently operating at 30.0 GHz and 51.25 GHz are moved to adjacent protected passive bands at 31.3-31.5 GHz and 50.2-50.4 GHz, respectively. Neither of these channels has great sensitivity to the vertical distribution of water vapour or oxygen, so this change should not significantly alter the performance of the radiometer or retrievals derived from its data. However, it would still be necessary to account for this change in the radiative transfer modelling.

3.6.1 Interference Mitigation

There are a range of techniques that have been applied in passive radiometry to mitigate against interference. Many of these rely on identifying anomalous differences between signals in channels or views with similar characteristics to detect the typical narrow range of the interfering signal's spectrum, directivity or polarisation. For example, the passband of each

channel can be divided into discrete sub-bands, each of which is detected independently. If the radiances measured in all sub-bands pass a test of statistical consistency, they can be co-averaged and processed as normal; if not, that measurement is flagged as suspect interference. A similar technique can be applied to radiances measured at close separation of elevation or azimuth angles, which has the advantage of not requiring any additional hardware. The fact that thermal emissions are not expected to be polarised can also be exploited to detect anthropogenic emissions, which are typically linearly, or circularly polarised. However, again, this requires increased hardware complexity by having either a quasi-optic, a waveguide polarisation switch or multiple antennae.

The recent development of agile digital detectors for microwave radiometers allows very rapid sampling of the signal, which in turn allows the calculation of higher order moments of the signal's probability density function than have traditionally been used. The *kurtosis* of the signal is the ratio of its 4th:2nd moments. Higher kurtosis means more of the variance is due to infrequent extreme deviations, as opposed to frequent modestly-sized deviations, and is a very sensitive indicator of RFI compared to natural, thermal emissions. Misra *et al.* [2006] described a 7.4 GHz radiometer system that combines this with digital sub-band filtering to calculate the first 4 moments. Although this technique requires complex hardware and computation, and has only been applied at relative low microwave frequencies, it offers much promise for the future.

It is recommended that in addition to implementing one or more of these techniques to mitigate against RFI in the radiometer's design, mechanisms are put in place to monitor licence applications and feedback any observations of instances of RFI to the regulatory authority to allow them to enforce the limited protection provided to passive applications. If the radiometer is capable of scanning frequency and elevation and/or azimuth angles, full scans should be occasionally scheduled in the observing sequence, specifically to look for new RFI sources.

3.7 Future hardware developments

During the past decade the price of a microwave profiler has reduced by a factor of 2. At the same time, technological developments, such as direct detection or better switching, have improved their accuracy and stability. Yet, at ~£100k they are still too expensive to deploy in a dense observing network.

Recent developments in Monolithic Microwave Integrated Circuit (MMIC) technology have allowed the integration of key components of a water vapour profiling radiometer into a single package at unit component costs of 32 k€ [Iturbide-Sanchez *et al.*, 2004]. This includes a local oscillator, mixer, power splitter, amplifiers, channel definition filters, detectors and pin diode switches to reference the scene to noise diodes and chip terminations, but not the

antenna. The advantages of MMIC technology is that it allows mass production of identical, compact and now low-noise systems. The component costs of additional radiometers would be ~9 k€ each. This MMIC technology is available now at frequencies up to 30 GHz, and is expected to be extended to 60 GHz by 2009 [S. Reising, personal communication, 2006]. This opens up the potential for mass production of radiometer receivers suitable for temperature as well as humidity profiling.

If these receivers could be integrated into simplified radiometers to keep their production costs low, they could be deployed in dense national networks of dozens, or hundreds across Europe. This requires a substantial investment by the manufacturers, but also in the assimilation of their data. It is, therefore, recommended that these developments are closely monitoring by continuing liaison with the manufacturers, and that instruments based on this new technology are tested for operational suitability.

3.7.1 Simplifying the radiometer for weather station integration

If dense networks of microwave radiometers are to be considered, they may be integrated as standard instruments in weather stations. In this case, the radiometer surface pressure, temperature and humidity sensors would obviously be redundant and could be omitted. A number of other measures are suggested below to simplify the radiometer package further and improve affordability, which warrant further investigation.

The scan mechanism could be removed. This would reduce the size, weight and power requirements considerably. It would also remove the need to couple multiple frequency bands onto the same optical axis. However, it would not allow elevation scanning, tip curve calibration or routine calibration against an ambient black body reference. In this case, the radiometer would need to be sufficiently stable for its drift to be reliably calibrated by monitoring the observations-minus-background as part of the cycle of assimilating its data into an NWP model. For this technique alone to provide sufficiently accurate calibration may require the radiometer to be stable enough so it does not drift significantly over the course of ~1 month. This may be feasible, given current performance levels with two internal references, if the radiometer can be thermally isolated from the environment – for example in a thick insulating enclosure. However, it may also be possible to use data from a co-located GPS sensor in clear conditions as an absolute calibration reference for the channels in the water vapour band, as both are sensitive to the integrated water vapour. (§4.9.6, §7.1.4)

The feedhorn antennae could be replaced by a more compact design, such as a waveguide array antenna [Iturbide-Sanchez *et al.*, 2004].

The blower mechanism could be removed to save costs and power requirements and improve reliability by reducing the number of moving parts. This would be more attractive for

radiometers with a fixed view away from zenith, so they could be built into an enclosure, which acts as a rain shelter. In this case, the microwave window could also be inclined to encourage rain to run-off, whilst still including a hydrophobic coating. A simple rain sensor adjacent to the microwave window could provide a useful quality flag to indicate when the radiometer's calibration is likely to be compromised by liquid water.

3.8 Summary of Radiometer Hardware

The performance of the Radiometrics TP/WVP-3000 has been analysed in terms of calibration bias and radiometric noise. Although similar microwave radiometers are available commercially, data collected with this instrument during a series of trials and experiments conducted during this study was used in this analysis. This radiometer is capable to scanning a range of elevation angles. However, it was shown that observations at low elevation angles are susceptible to errors due to misalignment, which may be reduced by averaging symmetric views and improved by better mirror control. Low frequency channels were also found to be sensitive to emission from the surface, which limits their application at low elevation angles.

The Radiometrics TP/WVP-3000 includes reference noise diodes and an internal ambient black body calibration target, which are viewed frequently to calibrate the receiver's gain and drift. Referencing the radiometer's signal to these adds noise to its calibrated brightness temperatures, which could be reduced by averaging them over a period of a few minutes. However, the noise diodes do not provide an absolute reference, and their brightness temperatures must be calibrated against an external reference scene occasionally to provide an absolute reference and account for their long-term drift. The low frequency channels, where the atmosphere is optically thin, can be calibrated by the *tip curve* method, which relies on the assumption that the opacity increases linearly with path length through the atmosphere for different elevation angles. At higher frequencies, the atmosphere becomes optically thick and the absolute calibration is provided by viewing a cryogenic reference target for a short period. Fortunately, the noise diodes were found to be very stable and only need calibration every few months. The radiometer's noise levels was assessed during calibrations against a cryogenic reference target. Its calibration accuracy was assessed by comparison with forward modelled radiosondes in clear sky, which identified biases in the instrument's calibrations, radiosondes' humidity and absorption models. These results and analysis in later chapters led to a recommended observing sequence for this instrument (§7.1.6).

It is recommended that future developments in MMIC technology are monitored as this may allow the mass production of microwave radiometers operating at frequencies up to 60 GHz, which may be suitable for temperature profiling. Also, techniques to mitigate Radio Frequency Interference (RFI) should be implemented to reduce the risk of this jeopardising the future viability of microwave radiometry at these frequencies.

Chapter 4 Variational Retrievals (1D-VAR)

The fundamental retrieval problem can be stated mathematically as a transformation from an observation vector, \mathbf{y} , which includes all the channels of an observing system, to a state vector, \mathbf{x} , which describes the vertical profile of temperature, humidity and cloud. A *forward model* is a function, H , which transforms state space into observation space, $\mathbf{y} = H(\mathbf{x})$, based on the radiative transfer model described in Chapter 2. The *inverse* problem involves mapping observations into state space, which requires a statistical approach as described in this chapter.

However, in this case, the dimension of state space is much larger than that of the observations, so the retrievals are ambiguous – there are an infinite number of profiles that match a given set of observations. This is known as an *ill posed* problem, which can only be solved by the application of *a priori* constraints. This problem is further complicated by noise on the observations. This chapter describes the development of a variational retrieval method, which takes its *a priori* background data from NWP and combines it optimally with the observations in a way that is statistically consistent with the errors of each. This process is illustrated in Figure 4-1 and explained in detail in this chapter.

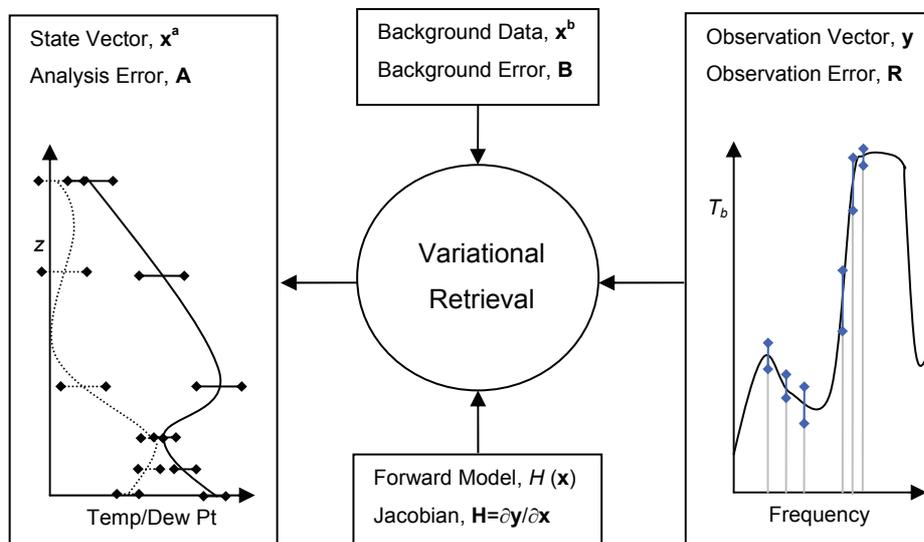


Figure 4-1 Schematic of inverse problem of retrieving profiles from observations. Right hand panel indicates set of brightness temperature observations with error bars corresponding to observation vector, \mathbf{y} , and its error covariance, \mathbf{R} . Left hand panel shows profile of retrieved temperature and dew point with error bars, corresponding to state vector, \mathbf{x}^a and its error covariance, \mathbf{A} .

4.1 Theory

Bayes' Theorem allows us to know how the probability distribution of a measurement, \mathbf{y} , maps into state space, \mathbf{x} and combines with prior knowledge:

$$P(\mathbf{x} | \mathbf{y}) = \frac{P(\mathbf{y} | \mathbf{x})P(\mathbf{x})}{P(\mathbf{y})} \quad (4.1)$$

where $P(\mathbf{x})$ and $P(\mathbf{y})$ are the *probability density functions* (pdf) in state and observation space respectively, $P(\mathbf{y}|\mathbf{x})$ is the pdf of \mathbf{y} given \mathbf{x} and $P(\mathbf{x}|\mathbf{y})$ is the pdf of \mathbf{x} given \mathbf{y} .

We seek the best estimate of \mathbf{x} , knowing *a priori* its *background* value, \mathbf{x}^b (e.g. given by an NWP model), and the coincident observation vector, \mathbf{y}^o . When the errors in the observations and the errors in the background have Gaussian distributions and are uncorrelated, then the logarithm of the conditional *pdf* of \mathbf{x} given \mathbf{y}^o is given the cost function [Rodgers, 2000]:

$$-2\ln P(\mathbf{x} | \mathbf{y}^o) = [\mathbf{x} - \mathbf{x}^b]^T \mathbf{B}^{-1} [\mathbf{x} - \mathbf{x}^b] + [\mathbf{y}^o - H(\mathbf{x})]^T \mathbf{R}^{-1} [\mathbf{y}^o - H(\mathbf{x})] + c_3 \quad (4.2)$$

where \mathbf{B} and \mathbf{R} are the error covariance matrices of the background, \mathbf{x}^b , and observations, \mathbf{y}^o , respectively (described in sections 4.4.1 and 4.3.2), $H(\mathbf{x})$ is the forward model operator, used to translate state space into observation space, and c_3 is a constant. This uses the standard notation of Ide *et al.* [1997].

Bayesian methods do not provide an exact answer to the inversion of observations into state space, but a posterior *pdf* from which the expected value or most probable state can be derived, together with an estimate of its uncertainty. For example, the maximum *a posteriori* probability can be found by maximising the conditional *pdf*, which is equivalent to minimising the cost function given by (4.2). Because $H(\mathbf{x})$ is nonlinear, this minimisation needs to be performed iteratively, as described in §4.7.

Because this minimisation uses the derivative of $H(\mathbf{x})$, it is referred to as a *variational* retrieval, which is a special case of *optimal estimation* [Eyre, 1990]. This refers to the fact that the observations and background are combined in an *optimal* way, accounting for the error characteristics of both. This is an example of a *physical* statistical retrieval method, as it relies on a physical, rather than empirical, model to relate the state and observation space.

This chapter reviews the requirements for One-Dimensional Variational (1D-VAR) retrieval of temperature and humidity profiles from radiometer data. This follows a similar approach to the recently *Integrated Profiling Technique* (IPT) recently published by Löhnert *et al.* [2004].

However, while the IPT concentrated on the retrieval of cloud properties and takes its prior data from nearby radiosonde profiles, these 1D-VAR retrievals use short-range forecasts from NWP models as backgrounds and concentrate on temperature and humidity profiling.

Different configurations are discussed and experiments conducted to evaluate their benefits.

4.2 Defining State Space

When defining the retrieval problem we have a choice of the definition of state space – both the control variables for temperature and humidity and the vertical grid.

4.2.1 Choice of state space – level heights and range

It is convenient to use levels of the mesoscale model (§1.1.1) as this provides the background for the retrievals and its levels are concentrated near surface, where the radiometer has the highest resolution (§4.6). In §2.6.3 it was shown that discretising the true profile at these levels produces a small, but significant quantisation error.

4.2.2 Choice of control variables

As the radiative transfer calculations are performed using thermodynamic temperature, this seems a natural choice for the control variable. However, this should be converted to a conservative quantity, like potential temperature, when calculating layer averages to ensure they are not biased.

In this study, the humidity components of the state vector are defined as the natural logarithm of total water, $\ln q_t$, where q_t is the total of the specific humidity and liquid water content. This control variable was first proposed by Deblonde and English [2003] for assimilation of cloud information from microwave satellite data. They used a simple linear partition function, where water in excess of $RH_1=0.95$ is portioned by a fraction, $C_{split}=0.50$ between vapour and liquid, according to (4.3). Water in excess of $RH_2=1.05$ all appears in condensed form.

$$\begin{aligned}
 RH_{qt} < RH_1 \dots q &= q_t \\
 RH_1 \leq RH_{qt} \leq RH_2 \dots q &= RH_1 q_{sat} + C_{split} (q_t - RH_1 q_{sat}) \\
 RH_2 \leq RH_{qt} \dots q &= q_{sat} [RH_1 + C_{split} (RH_2 - RH_1)]
 \end{aligned} \tag{4.3}$$

where
$$RH_{qt} = \frac{q_t}{q_{sat}} \tag{4.4}$$

and
$$q_L = q_t - q \tag{4.5}$$

This partition function gives a discontinuity in the derivative at RH_1 , which may produce unstable iterations in the retrieval when the retrieved humidity is close to this threshold. A modified version of this partition function has been developed to produce a smooth transfer function between water vapour for $q_t/q_{sat} < 90\%$ and liquid water for $q_t/q_{sat} > 110\%$ (where q_{sat} is the specific humidity at saturation), following (4.6):

$$\begin{aligned}
 RH_{qt} \leq RH_1 \dots \frac{dq_L}{dq_t} &= 0 \\
 RH_1 < RH_{qt} < RH_2 \dots \frac{dq_L}{dq_t} &= \cos^2 \left[\frac{(RH_{qt} - RH_2) \pi}{(RH_2 - RH_1) 2} \right] \\
 RH_2 \leq RH_{qt} \dots \frac{dq_L}{dq_t} &= 1
 \end{aligned}
 \tag{4.6}$$

where $RH_1=90\%$ and $RH_2=110\%$ were selected to approximate the partition function from Deblonde and English [2003].

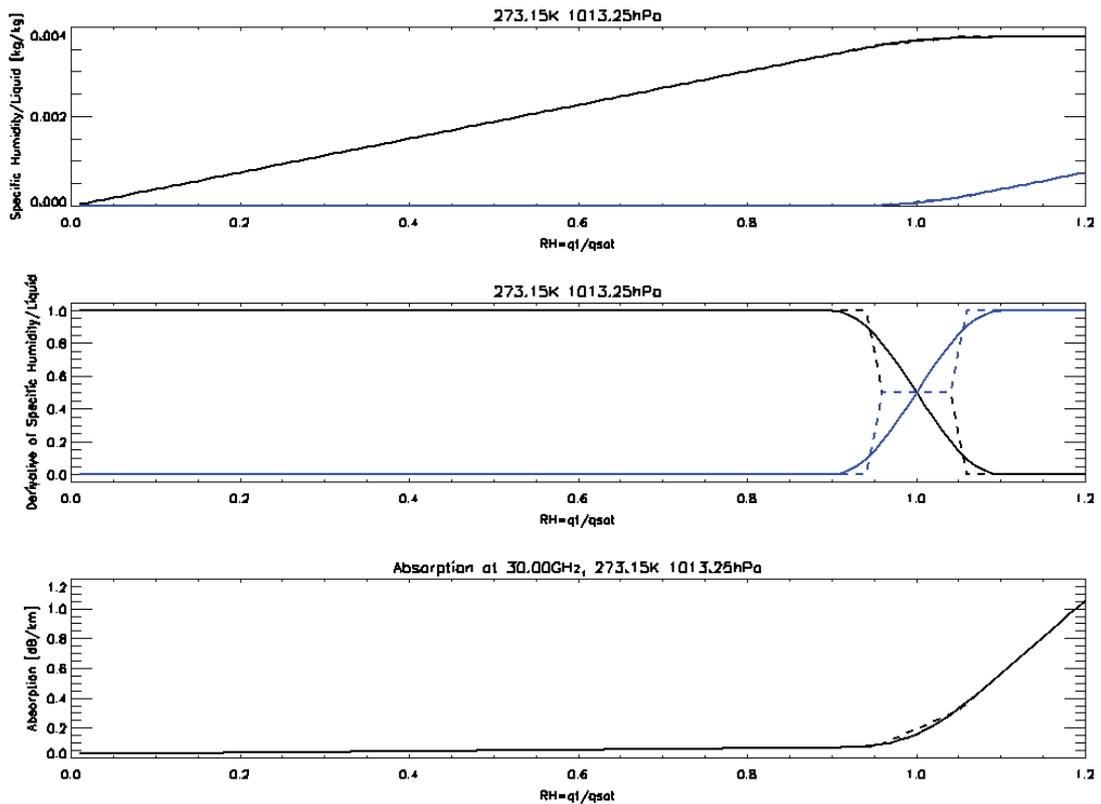


Figure 4-2 Total Water partition functions by Deblonde and English (dashed) and this study (solid) for a range of $q_t = (0 - 1.2)q_{sat}$.

Upper panel shows specific humidity, q (black) and liquid water mixing ratio, q_L (blue).

Middle panel shows derivatives dq/dq_t (black) and dq_L/dq_t (blue).

Lower panel shows 30 GHz absorption modelled by Rosenkranz'98 and Liebe (cloud).

Figure 4-2 illustrates how total water, q_t , is partitioned between vapour and condensate according to RH_{qt} . The upper panel shows how specific humidity (black line) increases linearly with RH_{qt} until $RH_{qt}=90\%$. At higher values of RH_{qt} increments in q_t result in the addition of liquid water (blue line). The middle panel of Figure 4-2 shows that the derivatives dq/dq_t (black) and dq_L/dq_t (blue) for the proposed function (solid lines) are continuous, in contrast to those in Deblonde and English [2003] (dashed lines). The lower panel shows the absorption increases steadily with q_t in clear air, but much more rapidly as the air becomes saturated.

Although sinusoids are used here in this partition function, it would also be possible to use a hyperbolic tangent, which also follows a smooth transition and has the desirable property of being continuously differentiable.

The choice of total water has the advantage of reducing the dimension of the state vector, enforcing correlation between humidity and liquid water and an implicit super-saturation constraint because of the stronger absorption of liquid water than vapour. The use of the logarithm creates error characteristics that are more closely Gaussian and prevents unphysical retrieval of negative humidity.

Initially the microwave forward model further partitioned condensed water between liquid and solid fractions using the function proposed by Jones [1995]. However, this was intended for application to mixed phase clouds and allows very little liquid (~20%) to form, even at temperatures just below 0°C, which does not apply to more general situations. It has been replaced by the simpler parametrization in (4.7), based on a fraction of liquid water increasing linearly with temperature from 0% at -40°C to 100% at 0°C, similar to that proposed by Wu and Weinman [1984]:

$$\begin{aligned} T < -40^{\circ}\text{C} \dots f_l &= 0 \\ -40^{\circ}\text{C} < T < 0^{\circ}\text{C} \dots f_l &= (T + 40)/40 \\ T > 0^{\circ}\text{C} \dots f_l &= 1 \end{aligned} \tag{4.7}$$

where f_l is the fraction of condensed water forming as liquid and T is its temperature (°C).

The remainder forms as ice and is effectively lost from the retrieval, as it is assumed to have negligible extinction at the frequencies of this study. This partitioning scheme inhibits the formation of liquid water at very low temperatures, although this was not found to have a strong impact on the retrievals. However, this partitioning is not applied in the infrared forward model, which treats ice and liquid in the same way, as ice cloud has a strong influence on T_{ir} , unlike at microwave frequencies.

4.3 Background

Additional information is needed to resolve the ambiguity introduced into the retrieval due to the under-constrained nature of the problem. For the variational retrievals, this *background* information is provided by *a priori* knowledge of the atmospheric state in the form of a short-range forecast from an NWP model. Although NWP model analyses are available for retrievals conducted retrospectively, they would not be available for real-time operational systems, (and may contain information from radiosondes, which are used in the validation of the retrievals). Our observations must improve on this background if they are to have any impact on the forecast model.

4.3.1 Mesoscale Model

The mesoscale version of the Met Office Unified Model (§1.1.1) is used to provide background data for the retrievals in the form of profiles of temperature, humidity and liquid water. The model grid points (12 km apart) are interpolated to the position of the observations. This model is initiated every six hours, including data from radiosonde stations. A short-range forecast (T+3 to T+9 hr) valid at the hour nearest the observations is used for the background, as would be available to operational assimilation schemes. This is independent of any radiosondes launched at observation time, which may be used to validate the retrievals.

The background was found to have a consistent bias with respect to co-located radiosondes used in this study, particularly during TUC in Payerne (Switzerland). This is believed to be due to the mesoscale model's limited representation of the orography. This bias was corrected empirically prior to using the background in the retrieval by subtracting the average difference between the model and radiosonde profiles.

4.3.2 Background Error Covariance Matrix

The background error covariance, \mathbf{B} , describes the expected variance at each level between the forecast and true state vector and the correlations between them. A good specification of background error covariances is an essential part of any state-of-the-art data assimilation system [Ingleby, 2001]. As will be shown, the specification of \mathbf{B} can be somewhat empirical, yet it has critical implications for the retrievals and their information content. For this reason it is important to validate the choice of \mathbf{B} , as will be shown in §4.6.3. Kadygrov and Pick [1998] estimated \mathbf{B} based on a diagonal matrix, with correlation between levels at different heights defined by a simple Gaussian function. In this study, \mathbf{B} was initially taken from that used to assimilate data from satellite instruments operationally at the Met Office, but it has also been evaluated by another method in the following subsections.

4.3.2.1 \mathbf{B} used for 1D-VAR of ATOVS data

Initially, the version of \mathbf{B} used for 1D-VAR assimilation of ATOVS data in the northern hemisphere into the global version of the Unified Model was investigated, referred to here as $\mathbf{B}_{\text{ATOVS}}$. Its temperature components are defined on 43 RTTOV (pressure) levels, while the humidity is defined in terms on $\ln q$ on the lowest 26 levels, but contains no correlation between temperature and humidity. The temperature and humidity sub-matrices were re-mapped onto the UMG3 vertical grid independently by interpolating their eigenvectors following the approach described in Collard [1998]. It is expected that the variances and correlations of \mathbf{B} would be a little smaller for the mesoscale version of the UM, because of its improved ability to represent small-scale features.

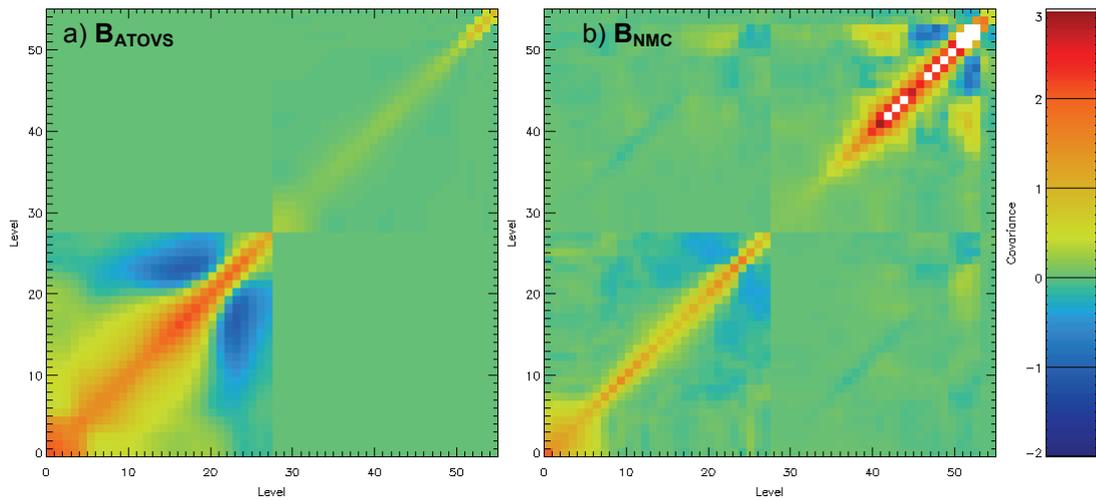


Figure 4-3 B matrices used in assimilation of ATOVS satellite data in the northern hemisphere by Met Office 1D-VAR, interpolated to lowest 28 UMG3 levels (left) and calculated by NMC method using T+6 – T+12 mesoscale model forecast profiles for Larkhill over 1 year (right). The state vector comprises temperature [K] at 28 levels, followed by total water ($\ln q_t$) [dimensionless] at the same 28 levels.

4.3.2.2 The NMC method

B has also been estimated using the *NMC method*, first described by Parrish and Derber [1992] (then at the National Meteorological Center, USA), which uses the covariance of the differences in state vector from sequential runs of an NWP model valid at the same forecast time to estimate **B**. This assumes that the statistical structure of forecast errors varies little over the period. Under this assumption, the spatial correlations of background error should be similar to the correlations of differences between two forecasts verifying at the same time. The advantage of the method is that it is straightforward to calculate the required statistics. The disadvantage is that the underlying assumption that the statistical structure of 48 hr - 24 hr forecast differences are similar to that of background error is difficult to justify.

The NMC method was applied to estimate **B** as the covariance of the difference between state vectors from one year of successive T+6 and T+12 hr mesoscale model forecasts for Larkhill. This site was chosen as representative of inland UK locations, which have larger diurnal variations than coastal areas. The resulting covariance matrix was scaled as described in §4.6.3 to give **B_{NMC}**, for which the magnitude was consistent with the observed *Observation minus Background (O-B)* statistics.

Figure 4-3 shows **B_{NMC}** is approximately diagonal, but with significant correlation between adjacent levels, especially near the surface. The terms of $\sqrt{\text{diag}(\mathbf{B}_{\text{NMC}})}$ are also shown in Figure 4-11, which are ~ 1.0 K for temperature components at all levels, similar in magnitude to **B_{ATOVS}**, but with a different correlation structure. The diagonal terms of the total water

components of $\sqrt{\text{diag}(\mathbf{B}_{\text{NMC}})}$ increase from 0.25 near the surface to 1.00 by 3.5 km. (An error of 0.1 in $\ln q$ is equivalent to a 10% error in specific humidity.) There is also significant correlation between temperature and humidity components, which is often neglected. Combined with the some channels' sensitivity to both temperature and humidity, this reduces the independence of these variables in the retrievals.

Other inland sites gave similar statistics. Histograms of the difference between T+6 and T+12 hr forecasts were examined and found to be approximately Gaussian, with few outliers. Approximately 1% of cases lie outside the 3-sigma limit for both temperature and humidity (although the surface temperature and stratospheric humidity have more outliers). This is, however, higher than would be expected from a purely Gaussian distribution (0.27%).

4.4 Observation Vector

In this study the observation vector, \mathbf{y} , initially comprises the zenith brightness temperatures in the 12 channels of the Radiometrics TP/WVP-3000 microwave radiometer, $T_{b1}, T_{b2}, \dots, T_{b12}$, the ambient air temperature, T_{AMB} , and humidity, $\ln q_{\text{AMB}}$, both measured at ~ 1.2 m. This vector is later suffixed with the brightness temperature from the infrared radiometer, T_{ir} :

$$\mathbf{y} = [T_{b1}, T_{b2}, \dots, T_{b12}, T_{\text{AMB}}, \ln q_{\text{AMB}}, T_{\text{ir}}] \quad (4.8)$$

4.4.1 Observation Error Covariance

The observation error covariance, \mathbf{R} , has contributions from the radiometric noise (\mathbf{E}), forward model (\mathbf{F}) and representativeness (\mathbf{M}) errors:

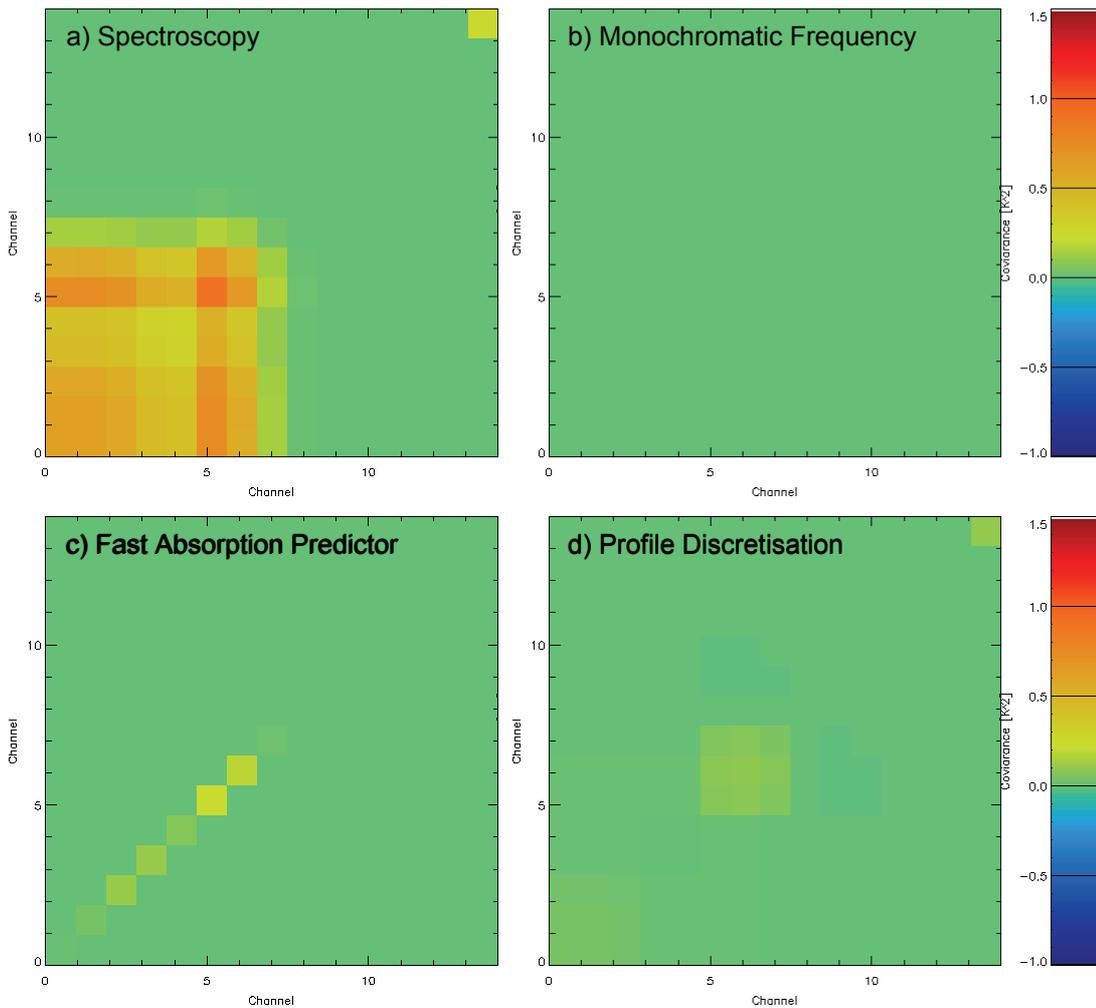
$$\mathbf{R} = \mathbf{E} + \mathbf{F} + \mathbf{M} \quad (4.9)$$

4.4.1.1 Radiometric Noise

The radiometric noise, \mathbf{E} , of the microwave radiometer was evaluated in detail in §3.5.3 as the covariance of the T_b s measured while viewing a stable scene over a short period. As shown in Figure 3-5, this term is approximately diagonal – with diagonal terms $\sim (0.1\text{-}0.2 \text{ K})^2$ – i.e. the channels are independent.

4.4.1.2 Forward Model Error

The forward model errors include random errors introduced in the spectroscopy, any fast model parametrizations and the profile representation in the radiative transfer model, which were evaluated in §2.2.7, §2.5.2, §2.5.3 and §2.6.3, respectively. The contributions of these terms to the overall forward model error covariance are shown in Figure 4-4, which clearly shows it is dominated by the uncertainties in the spectroscopy, which are the most difficult to estimate accurately. However, as noted in §2.2.7, this term may be over-estimated.



**Figure 4-4 Contributions to Forward Model error covariance due to
a) spectroscopic errors, b) monochromatic frequency approximation,
c) Fast Absorption Predictor and d) profile discretisation.**

**First 12 channels are Radiometrics TP/WVP-3000 brightness temperatures [K²].
Last 3 channels are ambient temperature [K²], total water [lnq_t²] and IR T_b [K²].**

4.4.1.3 Representativeness Error

“In objective analysis the desired resolution determines a scale, which may be described by saying a grid-point value represents an average in space and time. It is usual to define the analysis and background errors as deviations from the truth also truncated to this scale. Thus it is possible to conceive a ‘perfect’ analysis or background, even though they do not represent all atmospheric scales. To compensate for this it is usual to add errors of representativeness to the instrumental errors to get a larger observational error. It is important to realize that these errors are as much a function of background representativeness as the observational representativeness; they can be considered explicitly as the errors in the generalized interpolation from background to observations.” [Lorenç, 1986]

The representativeness error, \mathbf{M} , allows for the radiometer's sensitivity to fluctuations on smaller scales than can be represented by the NWP model. \mathbf{M} has been estimated here by studying the fluctuations in the radiometer's signal on typical time scales taken for atmospheric changes to advect across the horizontal resolution of the NWP model. While it would have been better to estimate \mathbf{M} from coincident time series of observations from identical systems separated by an appropriate distance, this was not possible.

The representativeness error has been estimated from a time-series of observations during a rare 6 day period without rain at Camborne (23-29 July 2004) when the radiometer was sampling the zenith view once a minute. The first day of this period was free of cloud; the next 2 days had variable amounts of low cumulus cloud. This was followed by 2 overcast days and 1 more day of variable cloud. Although the variability during this period may be smaller than average due to the anti-cyclonic conditions that pervaded, it is difficult to identify a longer rain-free period of data for analysis. The extinction by rain drops tends to produce highly erratic time-series of T_b , which dominate the time structure function and make retrievals unreliable, for the reasons discussed in §2.4. Figure 4-5 shows the time-series of microwave and infrared brightness temperatures, ambient temperature and humidity measured during this period.

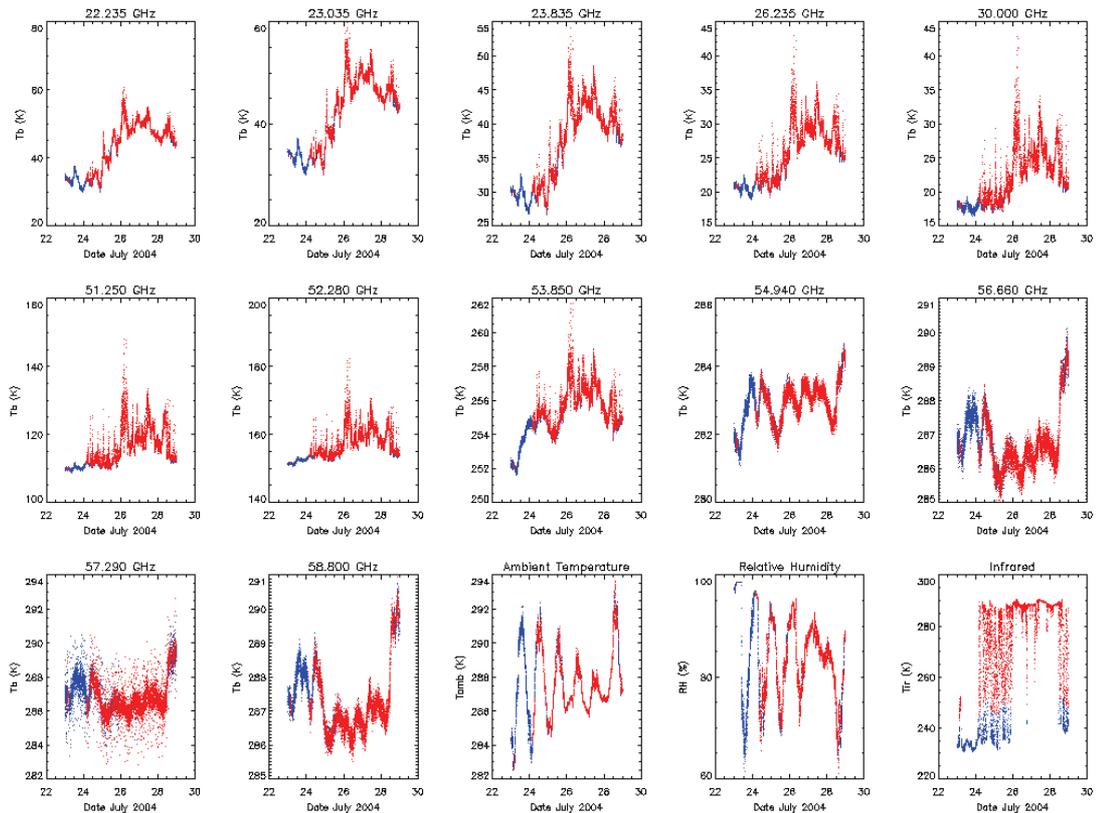


Figure 4-5 Time-series of observations used to calculate representativeness covariances. Blue points are classified as *Clear Sky*, Red points as *Cloudy*.

The Time Structure Functions (TSFs) were calculated over various sampling intervals as in §3.5.4 from this data set. These are plotted as scalars on logarithmic axes in Figure 4-6 for all the channels of the Radiometrics TP/WVP-3000, including ambient temperature, humidity (as $\ln q$) and T_{ir} . The low frequency channels, which are sensitive to humidity and cloud, show a steady increase from the shortest timescales (~ 1 min). The highest frequency channels, on the other hand, are sensitive only to the temperature in the lowest few hundred metres of the atmosphere. On short timescales ($<10^3$ s) their TSFs *bottom out* at the noise level of the radiometer, as it is unable to resolve atmospheric fluctuations below this level. On longer time scales the TSF often peak for periods ~ 12 h, corresponding to the diurnal cycle. This shows how the TSF combines the atmospheric variability and the instruments' noise.

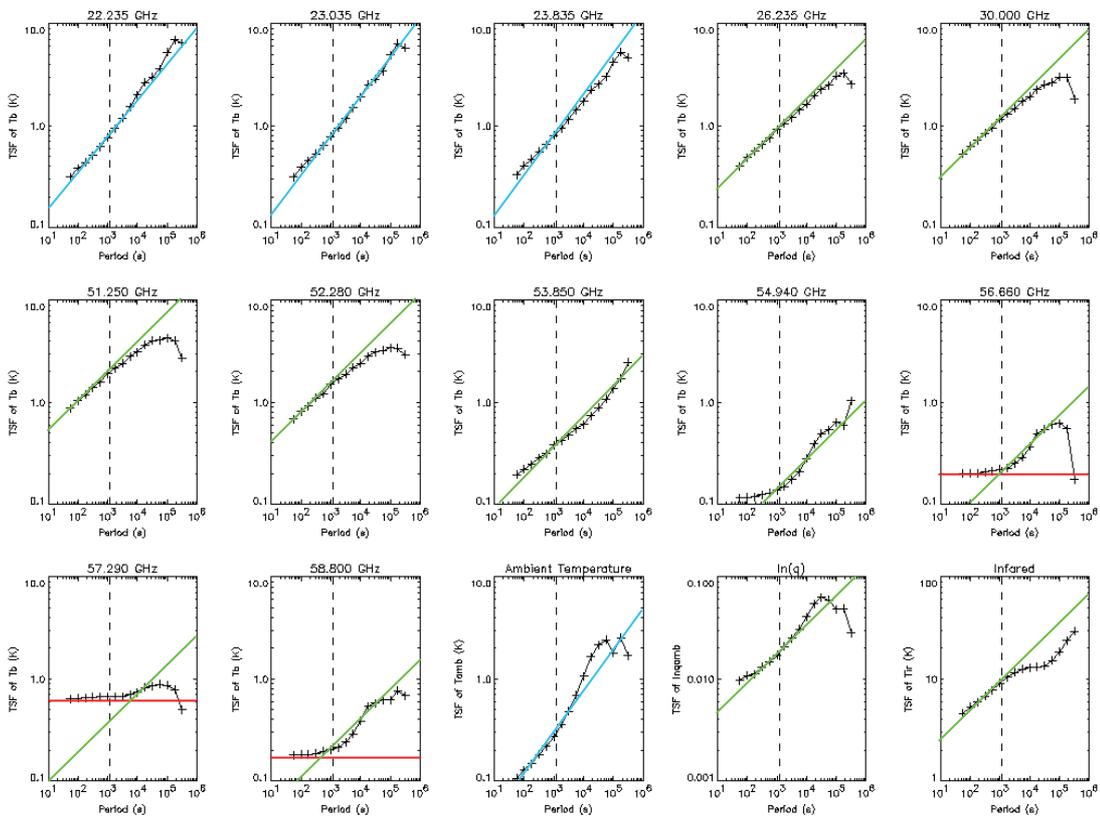


Figure 4-6 Time Structure Functions of Radiometrics TP/WVP-3000 observations measured at Camborne 23-29 July 2004 in all dry conditions. Coloured lines show reference slopes: Red=0, Blue=1/3, Green=1/4.

It is also possible to calculate a covariance matrix from these data. In the case of the mesoscale model with a 12 km grid, a time scale, $\Delta t = 1200$ s was chosen to represent a typical advection timescale and the covariance was calculated from this to estimate the typical representativeness error, \mathbf{M}_{typ} :

$$\mathbf{M} = \varepsilon(\mathbf{y}^o(t + \Delta t) - \mathbf{y}^o(t))(\mathbf{y}^o(t + \Delta t) - \mathbf{y}^o(t))^T \quad (4.10)$$

where $\mathbf{y}^o(t)$ and $\mathbf{y}^o(t+\Delta t)$ are the observations at times separated by interval, Δt , and ε is the expectation operator.

When TSFs are calculated over shorter data sets they show great variability, depending on the atmospheric conditions. In clear skies the TSF of the low frequency channels can be an order of magnitude smaller than in cloudy conditions. The data set shown in Figure 4-5 was divided between clear and cloudy periods, according to a simple threshold of infrared brightness temperature (see §4.8.1) and the TSF covariances were recalculated for each classification. The corresponding covariances, $\mathbf{M}_{\text{clear}}$ and $\mathbf{M}_{\text{cloudy}}$ are shown on different colour scales in Figure 4-7. $\mathbf{M}_{\text{clear}}$ shows strong correlation between the low frequency channels most sensitive to humidity, while the temperature sounding channels are almost diagonal, but most components are small ($\ll 0.05 \text{ K}^2$). (The peak in the 11th channel is an instrumental error on this version of the radiometer.) $\mathbf{M}_{\text{cloudy}}$ shows has much larger terms ($\gg 1 \text{ K}^2$) in those channels most sensitive to cloud. This term dominates the observations error budget of these channels in cloudy conditions. \mathbf{M} would be smaller for higher resolution models in proportion to the ratio of their resolutions to the power of the slope of the TSF in Figure 4-6. e.g. In a model with 3 km resolution, $\sqrt{\text{diag}(\mathbf{M})}$ would be $\sqrt{2}$ smaller for the 26-54 GHz channels than in the 12 km model shown here.

However, this method is likely to underestimate the spatial variability for the surface sensors, which are strongly coupled to surface properties.

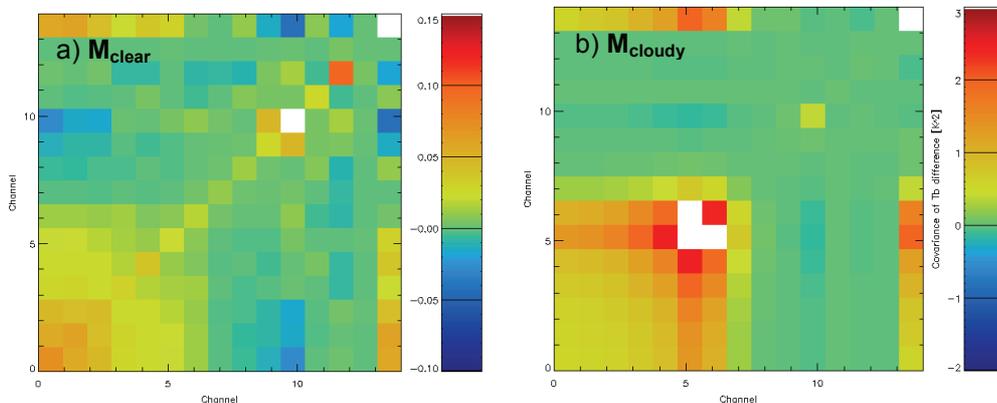


Figure 4-7 Covariance matrices of TSFs calculated from observations in clear and cloudy conditions over 1200 s periods 23-29/7/04 (n.b. different colour scales). First 12 channels are Radiometrics TP/WVP-3000 brightness temperatures [K^2]. Last 3 channels are ambient temperature [K^2], total water [$\text{In}q_t^2$] and IR T_b [K^2]. This estimates the representativeness error in applying point measurements to a 12 km grid box, with typical advection speeds of 10 m/s.

The representativeness error has also been evaluated dynamically, by recalculating **M** from the time-series of observations within a 1 hr window around each observation used for retrieval. It is hoped that this technique will allow the observation errors to be reduced in periods of atmospheric stability, when more confidence can be placed that the radiometer observations are representative of the model's state.

4.4.1.4 Total Observation Error Covariance

The magnitudes of the diagonal components of each term of **R** are shown in Table 4-1 for the 12 channels of the microwave radiometer, surface temperature and humidity sensors (as dimensionless $\ln q$) and infrared radiometer. In general the contributing error terms have similar magnitudes, which shows appropriate approximations have been made in the parametrizations. However, it is also noticeable that the channels near the 22 GHz line are dominated by spectroscopic uncertainties, whilst those channels most sensitive to cloud are dominated by their representativeness errors and the highest frequency channels, which are only sensitive to the temperature in the lowest few hundred metres, are dominated by radiometric noise. The total observation error covariances are shown in Figure 4-8.

While inaccurate estimates of **R** would cause the variational method to produce results that are not strictly optimal, it alone would not cause the observations to be rejected. However, given the large uncertainty associated with the estimation of the background error covariances, this is unlikely to be the dominant source of non-optimality.

Table 4-1 Diagonal components of Observation Error Covariance Matrix, $\sqrt{\text{diag}(\mathbf{R})}$ evaluated for all dry weather conditions. Dominant terms highlighted in bold.

Channel	Measurement Noise, E	Modelling Errors, F	Representativeness Error, M	Total Uncertainty, R	Units
22.235 GHz	0.17	0.83	0.65	1.07	K
23.035 GHz	0.12	0.84	0.67	1.08	K
23.835 GHz	0.11	0.82	0.69	1.08	K
26.235 GHz	0.13	0.67	0.78	1.04	K
30.000 GHz	0.21	0.61	1.00	1.19	K
51.250 GHz	0.18	1.10	1.70	2.04	K
52.280 GHz	0.15	0.88	1.35	1.62	K
53.850 GHz	0.17	0.35	0.32	0.50	K
54.940 GHz	0.18	0.06	0.10	0.14	K
56.660 GHz	0.19	0.05	0.10	0.22	K
57.290 GHz	0.54	0.05	0.40	0.67	K
58.800 GHz	0.18	0.06	0.11	0.22	K
T_{AMB}	0.17	0.00	0.22	0.28	K
$\ln q_{\text{AMB}}$	0.01	0.00	0.02	0.02	
T_{ir}	0.78	0.27	9.10	9.14	K

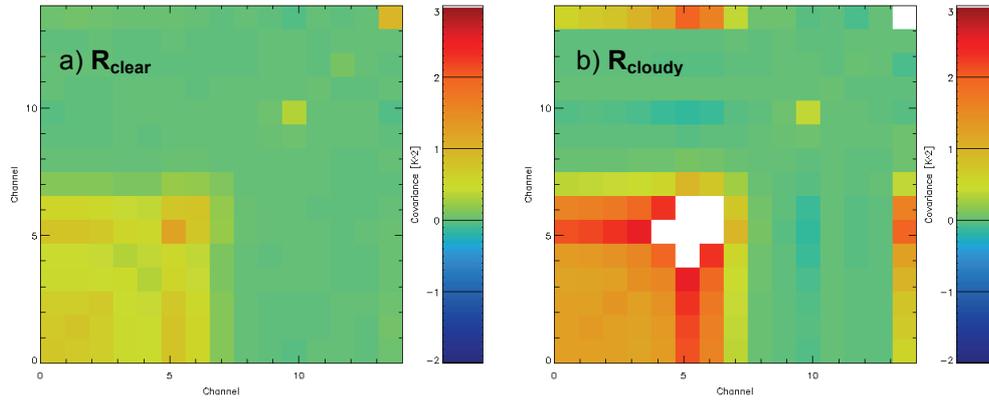


Figure 4-8 Total Observation Error Covariances, R_{clear} and R_{cloudy} as Figure 4-7.

4.5 Jacobians

All variational retrieval techniques rely on the derivative of the forward model. This is defined by the *Jacobian* matrix, $\mathbf{H}(\mathbf{x})$, which is a weighting function describing the sensitivity of the observation vector to changes in the state vector:

$$\mathbf{H}(\mathbf{x}) = \frac{\partial \mathbf{y}}{\partial \mathbf{x}} \quad (4.11)$$

The elements of \mathbf{H} contain the partial derivatives $\partial y_i / \partial x_j$, where the subscript i refers to the channel number and j to the position in state vector.

The Jacobian is needed in the minimisation of the cost function to determine the most probable atmospheric state as part of the retrieval or assimilation process. As will be shown in §4.6.3, this minimisation is an iterative procedure, which involves multiple evaluations of the Jacobian as it is nonlinear, in that it depends on the reference state, \mathbf{x} . The Jacobian is also needed to conduct an error analysis.

4.5.1 Calculation of Jacobians

\mathbf{H} , can be calculated by repeated evaluation of the forward model by perturbing each element of the reference state vector in turn by a small amount and examining the change in the observation vector. This is known as the *brute force* method and, although slow, is robust and simple to implement, so provides the basic technique used in these investigations. The perturbed profiles are then converted to p , T , q and q_L and linearly averaged into layers, before re-evaluating the forward model. (The layer thickness also needs to be re-calculated to account for changes in density.)

Perturbations of 1K in T and 0.001 in $\ln q_t$ were chosen to calculate Jacobians by brute force. (The latter is equivalent to a 0.1% change in absolute humidity.) These represent a compromise between perturbations that are large enough not to introduce significant rounding errors in single precision calculations yet small enough for the response to be linear. This

selection was not trivial for the case of total water, especially near the threshold at which cloud starts to condense. These choices were tested for linearity and computational stability by checking the differences between the forward model results applying positive and negative perturbations are much less than the expected observation errors – i.e. $H(\mathbf{x}+\delta\mathbf{x})-H(\mathbf{x}-\delta\mathbf{x})\ll\text{diag}(\mathbf{R})$.

Figure 4-9 shows the Jacobian's temperature components for the oxygen band channels and humidity components for the lower frequency channels of the Radiometrics TP/WVP-3000, scaled by the model layer thickness to show *weighting functions*. These show the highest frequency channels nearest the centre of the oxygen absorption band are most sensitive to temperature changes near the surface, whilst lower frequency channels in the oxygen band have more uniform sensitivity to temperature changes from 0-4 km. However, the Jacobian's humidity components have very similar shapes as the atmosphere is optically thin at these frequencies. This shows that there is less information available on the vertical distribution of humidity than temperature.

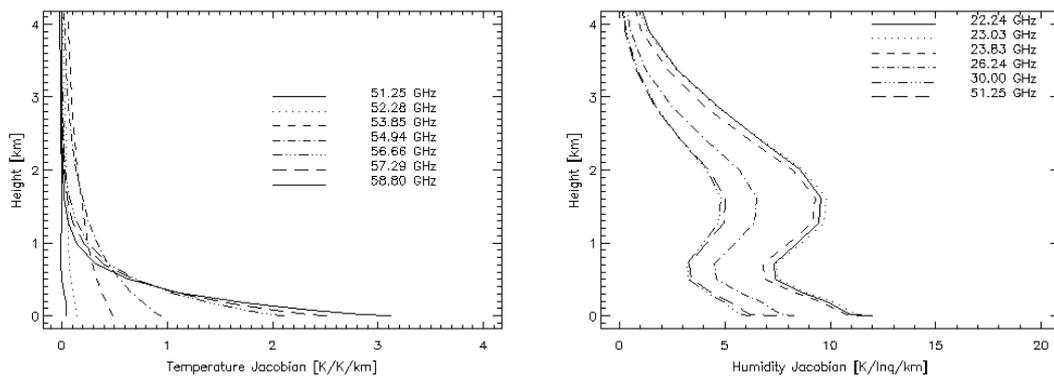


Figure 4-9 Jacobian's temperature component of 51-59 GHz channels (left) and humidity component for 22-51 GHz channels of Radiometrics TP/WVP-3000 (right), scaled by model layer thickness, Δz : $H/\Delta z = (\partial y/\partial x)/\Delta z$. Calculated for clear US standard atmosphere.

It would be much more efficient to calculate Jacobians analytically, but this requires the differentiation of the forward model. This is implemented in the fast model, RTTOV [Saunders *et al.*, 1999], described in §2.5.4. However, RTTOV has been designed for nadir-viewing satellite geometry and would require extensive modification to be applied to ground-based instruments. Other utilities are available to provide automatic differentiation of forward models – for example, TAMC [Giering, 1999]. However, this has not been pursued given the complexity of the code used in the radiative transfer model here.

4.5.2 Truncating Jacobian calculations at higher levels

Methods to speed up the forward model calculations were discussed in §2.5. There are, however, some additional measures that can be introduced to reduce the time taken to calculate Jacobians by brute force. The Jacobians are only calculated for levels below 8 km, and are assumed to be negligible at higher levels in the ground-based viewing geometry based on Figure 4-9, so higher levels are fixed to $\mathbf{H}=0$. This means that the retrievals will not modify the *a priori* profile above this level.

4.6 Error Analysis

Part of the beauty of Bayesian methods is that they allow quantification of the performance of assimilation and retrieval methods by propagating statistical errors from the observations and background. They can also provide an assessment of the information content and vertical resolution of the observations, accounting for the assumed error characteristics. These are useful tools to quantify the benefits of different combinations of channel frequencies, view angles, noise levels, sampling sequences and when optimising instrument configurations as well as assessing the accuracy of their products and their impact on NWP.

4.6.1 Analysis Error Covariance

Following the assumption that the observations and background *pdfs* are Gaussian and uncorrelated, it is possible to estimate the expected error in the analysis, \mathbf{x}^a . Rodgers [2000] shows in the linear case, the covariance of the analysed state vector, \mathbf{A} , is a simple function of the background error covariance, \mathbf{B} , and that of the observations, \mathbf{R} , mapped into state space by \mathbf{H} :

$$\mathbf{A} = (\mathbf{H}^T \mathbf{R}^{-1} \mathbf{H} + \mathbf{B}^{-1})^{-1} \quad (4.12)$$

Rodgers [2000] later goes on to show that equation (4.12) can also be applied for nonlinear problems that can be linearised about the analysis solution, \mathbf{x}^a , within the error bars of the retrieval. In this case, \mathbf{A} will be state-dependent, as it depends on the Jacobians evaluated about the analysis state.

\mathbf{A} has been evaluated for different instrument combinations in a clear US standard atmosphere. The full observation vector, \mathbf{y} , was used to evaluate the influence of the radiometer and surface sensors together. This was later supplemented with T_{ir} . The influence of only the surface sensors was investigated by removing the other channels from \mathbf{R} and \mathbf{H} and repeating the calculation of \mathbf{A}_{rs} . This analysis was repeated using the errors currently assumed in operational assimilation of radiosonde data (\mathbf{R}_{rs}) in the Met Office Unified Model, which are diagonal and dominated by representativeness, but are not necessarily accurate as they were derived empirically for an older version of the UM. In this case $\mathbf{H}=\mathbf{I}$, the identity matrix.

The diagonal terms of \mathbf{A} are plotted in Figure 4-10. This shows the error in the temperature profile retrieved by the radiometer (including surface sensors) is expected to approach 0.1 K near the surface, but increases linearly with height, to exceed 1 K above 5 km. For the humidity profile, \mathbf{A} varies greatly with the \mathbf{x} due to the nonlinearity of \mathbf{H} . In this example the $\ln q$ error in the radiometer's retrieval increases from 0.05 near the surface to 0.40 by 3 km. These errors correspond to uncertainties of 5% and 40% in humidity, respectively. The analysis error covariance, \mathbf{A} , for the radiometer's retrievals represent a substantial improvement on the background, $\mathbf{B}_{\text{ATOVs}}$, in the lowest 3 km (and >5 km for the temperature) and compared to the surface sensors only, which only influence the lowest 500 m. (The surface sensors' influence is determined by the correlation between the lowest levels in \mathbf{B} .) Figure 4-10 shows the retrievals from the radiometer are of comparable accuracy to the radiosonde profiles below ~1 km, but poorer above this.

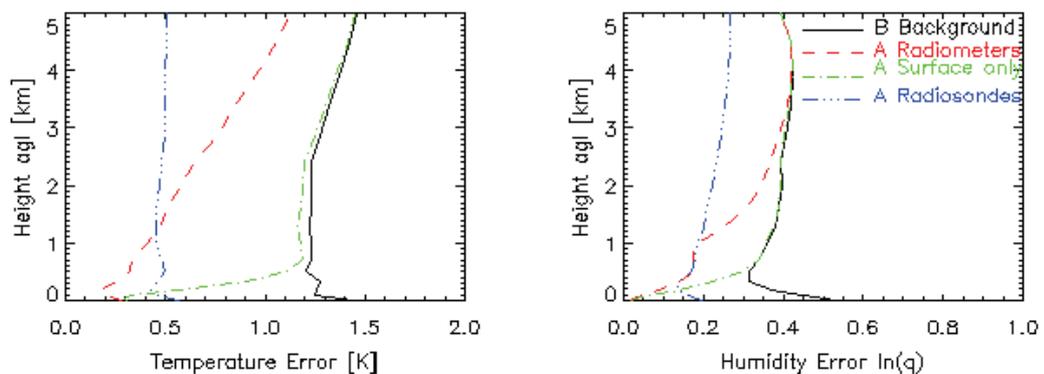


Figure 4-10 Error Analysis for 1D-VAR retrievals. Background error covariance matrix from mesoscale model, $\sqrt{\text{diag}(\mathbf{B}_{\text{ATOVs}})}$ (black) and analysis error covariances matrices, $\sqrt{\text{diag}(\mathbf{A})}$, with surface sensors only (green), radiometers and surface sensors (red), and radiosonde only (blue). Diagonal components plotted for lowest 5 km of temperature [K] and humidity ($\ln q$).

When \mathbf{B}_{NMC} is used instead of $\mathbf{B}_{\text{ATOVs}}$, \mathbf{A} gives more limited improvements in temperature, up to ≈ 4 km, while the analysed humidity profile improves the background slightly at all levels up to ≈ 5 km, although the humidity errors were much larger at high levels than with $\mathbf{B}_{\text{ATOVs}}$. With \mathbf{B}_{NMC} retrievals from the radiometer system are expected to have larger errors than the radiosondes above ~ 500 m.

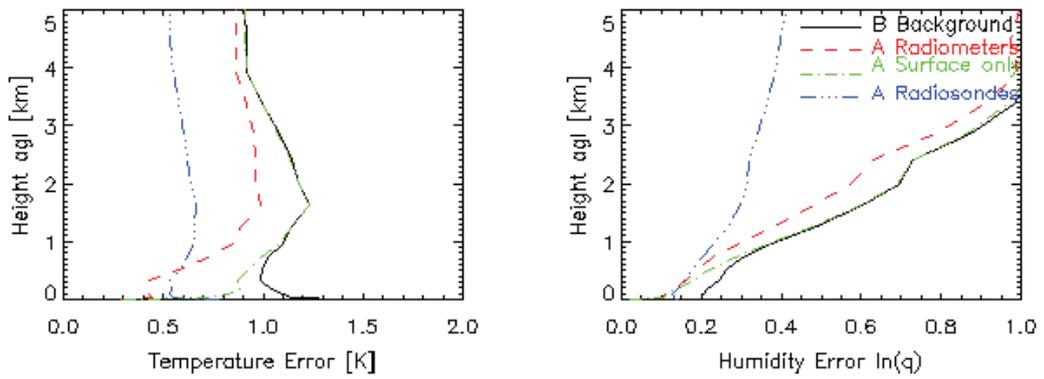


Figure 4-11 As Figure 4-10, but calculated with B_{NMC} , instead of B_{ATOVS}

4.6.2 Vertical Resolution

However, A only tells part of the story. The other important aspect of the retrieval's performance is the vertical resolution – its ability to resolve a perturbation in state space. One simple, robust definition of the vertical resolution, Δz , is the reciprocal of the diagonal of the *averaging kernel matrix* [Rodgers, 2000]:

$$\Delta z = \delta z / \text{diag}(\mathbf{A}\mathbf{H}^T\mathbf{R}^{-1}\mathbf{H}) \quad (4.13)$$

where δz is the level spacing.

This follows the concept of effective data density [Purser and Huang, 1993], whereby the trace of A represents the number of degrees of freedom at each level and its reciprocal is the number of levels per degree of freedom, hence vertical resolution.

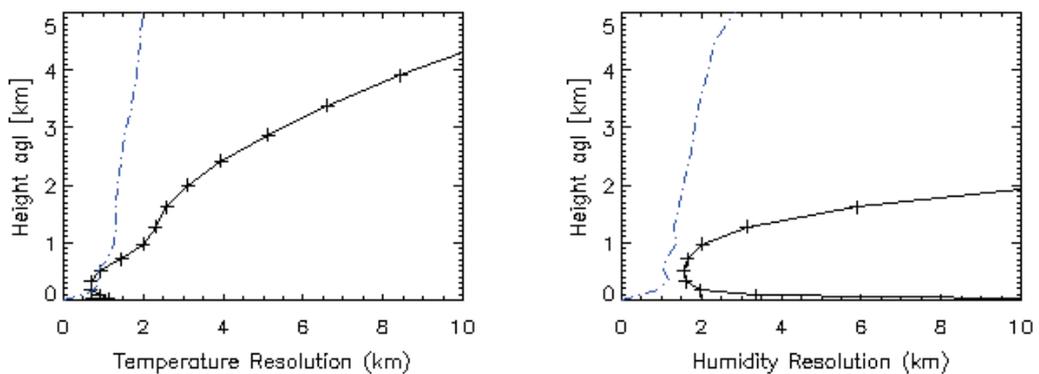


Figure 4-12 Vertical Resolution, Δz , of temperature and humidity ($\ln q$) from microwave radiometer and surface sensors (black solid line) and radiosondes (blue, dashed line) in clear standard atmosphere using B_{ATOVS}

Δz is evaluated in Figure 4-12, which shows that the vertical resolution of temperature profiles degrades with height, from ~700 m near the surface, approximately linearly as twice the height from 0.5-4 km. For lnq , it degrades very rapidly above 1.6 km, from ~1.6 km near the surface, but is critically dependent on the reference state, \mathbf{x} , due to non-linearity in \mathbf{H} .

The vertical resolution is also sensitive to the choice of \mathbf{B} , which makes it difficult to compare these results with other definitions, which tend to produce more optimistic results. For example, Smith *et al.* [1999] defined the vertical resolution from the vertical covariance of the difference between the retrieved profiles and the truth (co-located radiosondes). Liljegren *et al.* [2005] used this definition to calculate the vertical resolution of retrievals from a similar microwave radiometer as 0.75 km at the surface, increasing to 2.75 km by 4 km for the retrieved temperature profile, and 0.75 km at the surface, increasing to 1 km by 4 km for the humidity profiles. However, all definitions of Δz are much larger than the vertical spacing of the model levels, shown as the crosses in Figure 4-12, which is important if the retrievals are to be well represented on this vertical grid. However, this definition does have the advantage that it should be less dependent on the choice of humidity control variable.

The apparent degradation of vertical resolution near the surface is also due to the assumed correlations in \mathbf{B} . If the correlations between the 6 lowest levels in \mathbf{B} are suppressed by a factor of 10 for both temperature and humidity, the resulting vertical resolution do not increase near the surface in this way. These correlations in \mathbf{B} represent the trade-off between vertical resolution and suppression of noise in the retrievals and can be adjusted to suit particular applications.

The vertical resolution was also calculated for radiosondes if they were assimilated using $\mathbf{B}_{\text{ATOVs}}$. Figure 4-12 shows the vertical resolution even available from the radiosondes is limited by the correlations in \mathbf{B} to $\Delta z \sim 1$ km near the surface, and $\Delta z \sim 2$ km by 4 km for both temperature and humidity.

4.6.3 Validation of \mathbf{B} matrices by O-B statistics

Although the derivation of \mathbf{B} matrices is often seen as something of a black art, it is possible to check their consistency with the \mathbf{R} matrix by analysing the O-B statistics (§2.7.5) of a series of real *Observations minus Backgrounds* residuals. O-B is the term given to the difference between the observation and the background and is usually calculated in observation space. The covariance of this difference can be calculated as:

$$\mathbf{S}_{\text{O-B}} = \varepsilon \left(\mathbf{y} - H(\mathbf{x}^b) \right) \left(\mathbf{y} - H(\mathbf{x}^b) \right)^T \quad (4.14)$$

This can be compared to the theoretical value, $\mathbf{A}_{\text{O-B}}$, given by:

$$\mathbf{A}_{\text{O-B}} = \mathbf{H}^T \mathbf{B} \mathbf{H} + \mathbf{R} \quad (4.15)$$

The diagonal terms of \mathbf{A}_{O-B} are compared with the standard deviation of the observations' residuals from the background in clear sky conditions over 1 year of data in Figure 4-13. This shows that $\text{diag}(\mathbf{S}_{O-B}) > \text{diag}(\mathbf{A}_{O-B})$ for the 22-24 GHz channels, which are most sensitive to humidity. This figure also shows that $\mathbf{H}^T \mathbf{B} \mathbf{H}$ dominates the $\mathbf{R}_{\text{clear}}$ terms for these channels, which suggests that the humidity error is somewhat under-estimated in $\mathbf{B}_{\text{ATOVs}}$.

This analysis has also been applied using the covariance calculated by the NMC method, \mathbf{B}_{NMC} . This initially produced an $\mathbf{H}^T \mathbf{B} \mathbf{H}$ term which underestimated the observed $O-B$ variance. A better match was obtained by multiplying the whole covariance matrix by a factor of 3 and scaling the humidity quadrant by a further factor of 3. This scaling was applied to estimate \mathbf{B}_{NMC} in §4.3.2.2 to give $O-B$ statistics that are consistent with expectations.

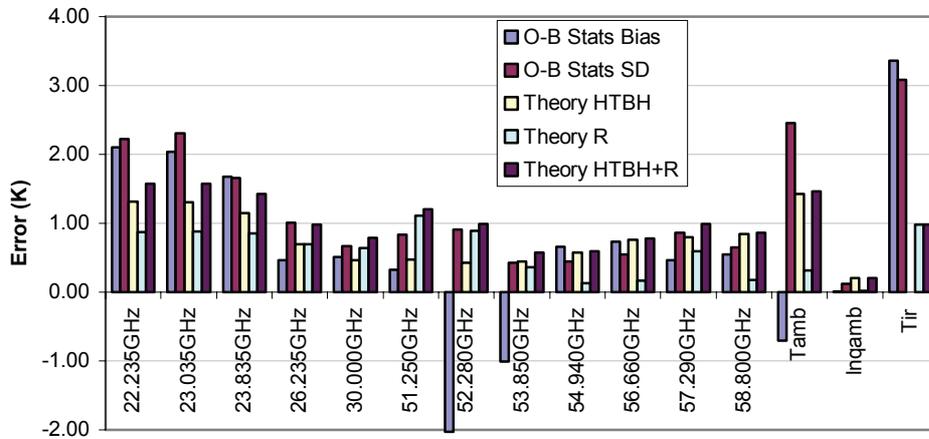


Figure 4-13 O-B statistics for clear cases compared to expected values, $\sqrt{\text{diag}(\mathbf{H}^T \mathbf{B} \mathbf{H} + \mathbf{R})}$

The $O-B$ statistics can also be used to correct the observations' bias empirically. It is important that they are consistent as this will encourage convergence to the true minimum in nonlinear problems, which may have multiple minima. Figure 4-13 also shows the mean bias of the radiometer observations over ~1 year with respect to the model background. These biases are broadly consistent with those found in the absorption model validation exercise (§2.7.2).

4.7 Cost function minimisation

The conditional probability density function given by equation (4.2) can also be thought of as a *cost function*, $J(\mathbf{x})$: its minimum value corresponds to the maximum *a posteriori* probability state, \mathbf{x}^a :

$$J(\mathbf{x}) = [\mathbf{x} - \mathbf{x}^b]^T \mathbf{B}^{-1} [\mathbf{x} - \mathbf{x}^b] + [\mathbf{y}^o - H(\mathbf{x})]^T \mathbf{R}^{-1} [\mathbf{y}^o - H(\mathbf{x})] \quad (4.16)$$

4.7.1 The Linear Case

It is possible to minimise equation (4.16) analytically in the limiting case where $H(\mathbf{x})$ is a linear function and \mathbf{x}^b and \mathbf{y} have Gaussian error characteristics, described by their covariances, \mathbf{B} and \mathbf{R} , respectively.

The application of microwave channels near the centre of the strong 60 GHz absorption band to retrieve temperature profiles in the lowest few hundred metres of the atmosphere is an example of a linear problem, where the cost function can be minimised analytically. An example of this is the retrieval of temperature profiles up to 600 m by Kadygrov and Pick [1998].

4.7.2 Gauss-Newton Minimisation

However, the humidity component of the retrievals is nonlinear. This also becomes important for lower frequency channels in the oxygen band used to profile temperature above ~ 1 km. Minimising a nonlinear cost function can be achieved by finding the value of \mathbf{x} for which its gradient is zero. Newton's method [Newton, 1669] can be applied iteratively to find the minimum of a nonlinear scalar function of a vector, $f(\mathbf{x})$:

$$\mathbf{x}_{i+1} = \mathbf{x}_i - [\nabla_{\mathbf{x}} f(\mathbf{x}_i)]^{-1} f(\mathbf{x}) \quad (4.17)$$

where \mathbf{x}_i and \mathbf{x}_{i+1} are the state vectors before and after iteration i .

The vector derivative of the cost function in the form of equation (4.16) is given by:

$$\nabla_{\mathbf{x}} J = -[\nabla_{\mathbf{x}} H(\mathbf{x})]^T \mathbf{R}^{-1} [\mathbf{y} - H(\mathbf{x})] + [\mathbf{x} - \mathbf{x}^b]^T \mathbf{B}^{-1} [\mathbf{x} - \mathbf{x}^b] \quad (4.18)$$

Further differentiation of $f = \nabla_{\mathbf{x}} J$, re-writing $\nabla_{\mathbf{x}} H(\mathbf{x})$ as \mathbf{H} , yields the *Hessian*:

$$\nabla_{\mathbf{x}} f = \mathbf{B}^{-1} + \mathbf{H}^T \mathbf{R}^{-1} \mathbf{H} - [\nabla_{\mathbf{x}} H(\mathbf{x})]^T \mathbf{R}^{-1} [\mathbf{y} - H(\mathbf{x})] \quad (4.19)$$

The third term on the right hand side of equation (4.19) is small in moderately nonlinear problems and can be ignored in *small residual* problems [Rodgers, 2000]. Substituting the remaining terms into (4.17) yields the Gauss-Newton method:

$$\mathbf{x}_{i+1} = \mathbf{x}_i + (\mathbf{B}^{-1} + \mathbf{H}_i^T \mathbf{R}^{-1} \mathbf{H}_i)^{-1} [\mathbf{H}_i^T \mathbf{R}^{-1} (\mathbf{y}^o - H(\mathbf{x}_i)) - \mathbf{B}^{-1} (\mathbf{x}_i - \mathbf{x}^b)] \quad (4.20)$$

It is usual to start the iterations at the background state, \mathbf{x}^b , and continue iterations until a convergence criterion is satisfied (see §4.7.4), at which point \mathbf{x}_{i+1} gives the analysis, \mathbf{x}^a , the retrieved best estimate of the truth. However, it is also possible to modify the first guess from \mathbf{x}^b to linearise the problem (see §4.8.2).

4.7.3 Levenberg-Marquardt Minimisation

A modification of the Gauss-Newton method was suggested by Levenberg and Marquardt [Levenberg, 1944] and [Marquardt, 1963] to improve the convergence in moderately nonlinear problems, where the first guess is too far from the truth for the increment to be reliable. This method adds a parameter, γ , to modify the direction of the step, which is adjusted after each iteration. In the form suggested by Rodgers [2000], the Levenberg-Marquardt method can be written as:

$$\mathbf{x}_{i+1} = \mathbf{x}_i + \left((1 + \gamma)\mathbf{B}^{-1} + \mathbf{H}_i^T \mathbf{R}^{-1} \mathbf{H}_i \right)^{-1} \left[\mathbf{H}_i^T \mathbf{R}^{-1} (\mathbf{y}^o - H(\mathbf{x}_i)) - \mathbf{B}^{-1} (\mathbf{x}_i - \mathbf{x}^b) \right] \quad (4.21)$$

where γ is a factor, which is adjusted after each iteration depending on how the cost function has changed:

- If $J(\mathbf{x})$ increases, reject this step, increase γ by a factor of 10 and repeat the iteration,
- If $J(\mathbf{x})$ decreases, accept this step, decrease γ by a factor of 2 for the next iteration.

These factors were adjusted empirically to minimise the number of steps needed to reach convergence, starting with an initial value of $\gamma = 2$.

It can be seen that as $\gamma \rightarrow 0$, the step tends towards that same as the Gauss-Newton (4.20). As $\gamma \rightarrow \infty$, the step tends toward the steepest descent of the cost function, allowing the search to jump out of a local minimum, towards a more global one. This is shown schematically in Figure 4-14. Such local minima are often found in nonlinear problems [Dharssi *et al.*, 1992].

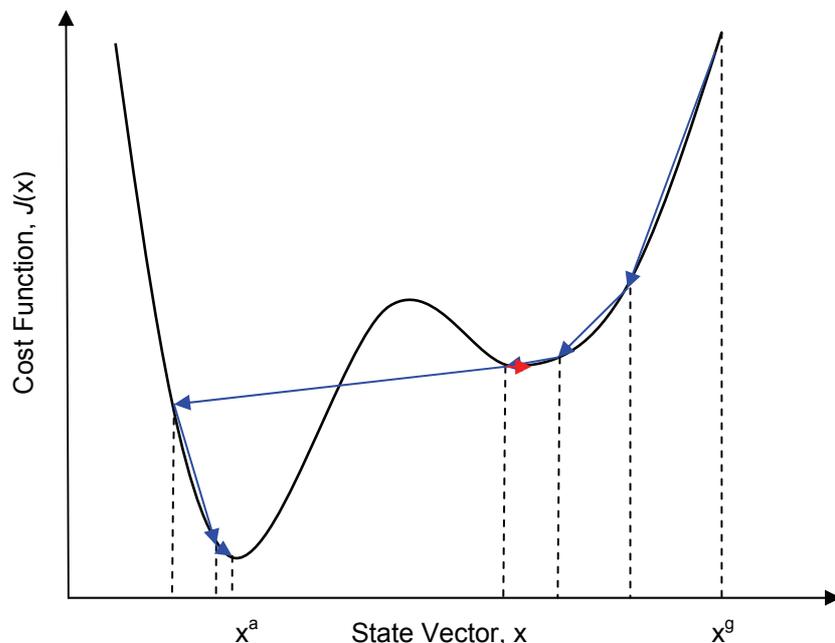


Figure 4-14 Schematic of minimisation of nonlinear cost function by Levenberg-Marquardt method, showing state vector, x , at each iteration (blue arrows) from first guess, x^g , to analysis, x^a . Red arrows are iterations that fail to reduce cost function, J .

4.7.4 Convergence Testing

Each successful iteration of the minimisation is expected to produce a smaller increment to \mathbf{x} . The iterations are repeated until a convergence criterion is satisfied at which point the retrieval is deemed to have converged. While it is possible to continue iterations until the machine precision is reached, this is inefficient, as steps much smaller than the expected error of the retrieval are not statistically significant [Press *et al.*, 1995]. Many variations of convergence criteria from the literature have been tried. For example, the Met Office's implementation of the Levenberg-Marquardt method for 1D-VAR [Collard, 2004] tests whether the cost function, J , has decreased by less than 1% and that the gradient is less than J^2 . This is because if γ is large it is possible that the cost function will not change between iterations simply because of the small step size in the minimisation. But this is quite a loose criterion.

Finally, the method proposed by Rodgers [2000] is adopted here: namely to check the size of the step in measurement space after each iteration by the following χ^2 test:

$$\left[(H(\mathbf{x}_{i+1}) - H(\mathbf{x}_i)) \right]^T \mathbf{S}_{\delta\mathbf{y}}^{-1} \left[(H(\mathbf{x}_{i+1}) - H(\mathbf{x}_i)) \right] \ll m \quad (4.22)$$

where m is the dimension of \mathbf{y} ($m=15$ in this case) and $\mathbf{S}_{\delta\mathbf{y}}$ is the covariance matrix between \mathbf{y}^o and $H(\mathbf{x}_i)$, given by:

$$\mathbf{S}_{\delta\mathbf{y}} = \mathbf{R} \left(\mathbf{H}\mathbf{B}\mathbf{H}^T + \mathbf{R} \right)^{-1} \mathbf{R} \quad (4.23)$$

The threshold was then tuned to optimise the radiometer retrievals. On the first iterations, $\ll m$ is taken to mean $< m/2$, however this is relaxed to $< m$ after 10 iterations. Typically it takes 3-10 iterations to reach convergence, each requiring ~ 0.25 s of CPU time on a 2.4 GHz Pentium IV using the *Fast Absorption Predictor* model.

4.7.5 Testing for Correct convergence

Once nonlinear problems have converged it is necessary to test for consistency of the retrieved solution with the observations and/or prior to check whether the minimisation has converged to a spurious minimum. This can be achieved as a χ^2 test, which compares a residual with the expected distribution using a theoretical χ^2 pdf to assess its statistical significance. In this case, the difference of the forward modelled analysis state, $H(\mathbf{x}^a)$ and observation vector, \mathbf{y}^o , are compared with the observation error covariance matrix, \mathbf{R} , to check if the retrieval is consistent with the observation by calculating:

$$\chi^2 = \left[H(\mathbf{x}^a) - \mathbf{y}^o \right]^T \mathbf{R}^{-1} \left[H(\mathbf{x}^a) - \mathbf{y}^o \right] \quad (4.24)$$

This consistency check was compared to other tests by comparing their performance when filtering retrievals of Integrated Water Vapour (IWV). This formed the basis for the selection of the χ^2 threshold, which was fine-tuned to optimise the convergence rate, while ensuring the standard deviation of the IWV retrievals did not increase significantly above the nominal

accuracy of 0.88 kg/m². This test was found to be more efficient than others proposed in Rodgers [2000].

χ^2 of solutions should follow a χ^2 distribution with m degrees of freedom. This would imply a conservative χ^2 threshold of 31 should reject less than 1% of a population from a pure χ^2 distribution with 15 degrees of freedom. However, observations in these extreme conditions may be of great interest, so it is important not to reject them. The observed distribution of this χ^2 was found to have a large number of cases outside this limit, which nevertheless yielded good retrievals. Hence retrievals with $\chi^2 > 100$ were rejected. The choice of χ^2 threshold was found to have little influence on the statistics of the retrieved profiles.

4.8 Pre-processing Classification and Quality Control

Examination of the performance of the retrieval scheme showed there were often problems when profiles approach the threshold of cloud formation – the residuals often oscillate without reaching convergence because the minimisation was not sufficiently linear. Although this was partially improved by the implementation of the Levenberg-Marquardt method of minimisation, it still persisted. Convergence problems where $\ln q_t$ approaches the cloud threshold can also be caused by the error characteristics of T_{ir} , which can be highly non-Gaussian.

4.8.1 Cloud Classification

These convergence problems have been addressed by introducing a cloud classification as a pre-processing step to the retrieval, based on a threshold of the infrared brightness temperature, T_{ir} . If T_{ir} was within 40 K of the ambient temperature the profile was classified as *cloudy* and the retrieval proceeds as described above. (This threshold was derived empirically.) Otherwise, the profile was classified as *clear* and the control variable changed from $\ln q_t$ to the logarithm of specific humidity, $\ln q$, and an addition term, J_{sat} , is added to the cost function to prevent saturation [Phalippou, 1996]:

$$J_{sat}(\mathbf{x}) = \sum_i c_i (x_i - x_i^{sat})^{p_{sat}} = \sum_i c_i (\ln q_i - \ln q_i^{sat})^{p_{sat}} \text{ if } q_i \geq q_i^{sat} \quad (4.25)$$

where q_i^{sat} is the specific humidity at saturation at level, i , p_{sat} and c_i are constants ($c_i=10^6$).

The cubic form, $p_{sat}=3$, has been selected because it is not highly nonlinear, while permitting a nonlinear gradient. Note that the first and second derivatives (gradient and Hessian) are continuous at the junction point $x_i=x_i^{sat}$. When this super-saturation term is included in the cost function, the iterations required for the Levenberg-Marquardt method (4.21) become:

$$\mathbf{x}_{i+1} = \mathbf{x}_i + \left((1 + \gamma) \mathbf{B}^{-1} + \mathbf{H}_i^T \mathbf{R}^{-1} \mathbf{H}_i + (p_{sat} - 1) p_{sat} \sum_i c_i (\ln q_i - \ln q_i)^{p_{sat}-2} \right)^{-1} \cdot \left[\mathbf{H}_i^T \mathbf{R}^{-1} (\mathbf{y}^o - H(\mathbf{x}_i)) - \mathbf{B}^{-1} (\mathbf{x}_i - \mathbf{x}^b) - p_{sat} \sum_i c_i (\ln q_i - \ln q_i)^{p_{sat}-1} \right] \quad (4.26)$$

In *clear* cases, the representativeness term in \mathbf{R} can be reduced by re-evaluating it in only clear sky conditions to allow more accurate retrievals in clear conditions.

4.8.2 Modifying the first guess to linearise problem

Although switching the humidity control variable, observation errors and super-saturation constraint in the cloud classification scheme works well for 1D-VAR retrievals, it is likely to prove problematic when applied to higher dimensional data assimilation. This is because the switching may not be consistent with the evolution of the surrounding background and may result in stability problems due to lack of thermodynamic balance. A preferable approach is to modify the first guess from the background. Until now, the background has been used as the first guess to initialise the minimisation, but it is possible to start the iterations from any first guess. This can be beneficial if the cost function is more linear between the first guess and the solution.

In clear conditions, the first guess can be modified from the background to ensure no level's humidity exceeds the cloud threshold. It is not so straightforward to modify the background in cloudy cases – a simple scheme evaluated here saturates the humidity at the level at which the temperature corresponds to T_{ir} , whilst ensuring that lower levels are not saturated.

It may also be necessary to ensure that the background profile's IWV and ILW are conserved when modifying it, although these constraints are likely to be implicit in the 1D-VAR retrievals, given the linear relationship between them and the observations.

4.8.3 Rain Detection

It is necessary to exclude observations in moderate or heavy rainfall (§2.4). Rain-affected observations are rejected, based on the combined output of the rain sensor and cases where $ILW > 1000 \text{ g/m}^2$. This was empirically found to be associated with precipitation, whether or not it reached the ground (in §6.5). Following the analysis presented in §5.5.2, cases where the retrieved IWV differs by more than 3 kg/m^2 from that measured from a co-located GPS sensor could be rejected as likely to be contaminated by rain. Lopez *et al.* [2006] suggested rain could be filtered by rejecting all cases where the observations differ by more than 3 K from the forward-modelled background as part of their assimilation of data from a two-channel ground-based microwave radiometer. However, this seems too strict, given the large changes of T_b caused by cloud variability.

4.8.4 Additional terms in cost function for physical constraints

It is possible to add extra terms in the cost function to constrain the retrievals to obey physical or empirical rules. These would be advantageous if the retrieved profiles were found to have undesirable characteristics – either consistently or occasionally. For example Peckham and Grippa [2000] used regularisation to stabilise the retrieval of temperature profiles by constraining them towards the adiabatic lapse rate.

Another example of this would be to inhibit the retrieval of super-adiabatic layers, where the temperature lapse rate is greater than the dry adiabatic lapse rate (9.8 K/km). These are often found at the surface in the presence of strong solar heating prior to convective thermals. However, elevated super-adiabatic layers thicker than 40 hPa are rare [Slonaker *et al.*, 1996] as mixing would cause them to disappear rapidly. Although there is no explicit constraint in the cost function to penalise the retrieval of super-adiabatic layers, the correlations in **B** effectively inhibit their formation and they are only retrieved in the lowest 1 km. In fact, the climatology of super-adiabatic layers in the retrieved profiles more closely matches that of the radiosondes, than the mesoscale model background, where they are limited to the lowest 80 m. Based on this, no additional super-adiabatic constraint is needed in the cost function.

Another possible constraint would be to prevent the formation of cloud liquid water content greater than $\sim 3 \text{ g/m}^3$. This corresponds to the maximum value typically found in cumulonimbus clouds [Sand *et al.*, 1984], although these are usually associated with heavy precipitation making retrievals impossible. However, analysis of the background and analysis profiles shows this value is never approached in practise, at the total water control variable provides an effective constraint implicitly.

4.9 Results of 1D-VAR retrieval experiments

This section discusses the results of 1D-VAR retrieval experiments based on different combinations of observations, backgrounds, first guesses and their errors and retrieval configurations suggested in the previous section. Each experiment was run on 706 cases, covering a year of data from Camborne from April 2004 to March 2005. Approx 65% of these were classified as *clear*. (n.b. This is much higher than the 38% fraction suggested by T_{ir} observations.) The retrievals were validated against temperature and humidity profiles measured by operational Vaisala RS80 radiosondes launched twice daily during this period. The experiments' configurations are summarised in Table 4-2.

Table 4-2 Configuration of 1D-VAR Retrieval Experiments

Section	Observations	Observation Error Covariance	Background	Background Error Covar.	First Guess	Truth for validation	Minimisation
§	\mathbf{y}	\mathbf{R}	\mathbf{x}^b	\mathbf{B}	\mathbf{x}^g	\mathbf{x}^t	Min
4.9.1.1	Synth	Fixed	Synth	$\mathbf{B}_{\text{ATOVS}}$	\mathbf{x}^b	Sonde	LM
4.9.1.2	Synth	Clear/cloudy	Synth	$\mathbf{B}_{\text{ATOVS}}$	\mathbf{x}^b	Sonde	LM
4.9.1.3	Synth	Clear/cloudy	Synth	$\mathbf{B}_{\text{ATOVS}}$	\mathbf{x}^b	Sonde	GN
4.9.2	Synth	Fixed	Meso.	\mathbf{B}_{NMC}	\mathbf{x}^b	Sonde	LM
4.9.3	Synth	Fixed	Meso.	\mathbf{B}_{NMC}	\mathbf{x}^b mod	Sonde	LM
4.9.4	Real	Fixed	Meso.	\mathbf{B}_{NMC}	\mathbf{x}^b	Sonde	LM
4.9.5	Real	Dynamic rep.	Meso.	\mathbf{B}_{NMC}	\mathbf{x}^b	Sonde	LM

4.9.1 1D-VAR Results with synthetic observations and background

Initially, synthetic observations, \mathbf{y}^o , and backgrounds, \mathbf{x}^b , were generated to be consistent with \mathbf{R} and $\mathbf{B}_{\text{ATOVS}}$, respectively, for each radiosonde launched from Camborne during the year. As the errors in these covariance matrices are correlated, this synthesis must be performed in eigenvector space, where they are diagonal. Kleespies and Crosby [2001] described a method of modelling a vector of correlated noise, $\varepsilon\mathbf{y}$ by:

$$\varepsilon\mathbf{y} = \mathbf{V}\sqrt{\mathbf{\Lambda}}\mathbf{z} \quad (4.27)$$

where \mathbf{V} is a matrix of the eigenvectors of \mathbf{R} ,

$\mathbf{\Lambda}$ is a matrix, with a diagonal composed of the square root of the eigenvalues of \mathbf{R} ,

\mathbf{z} is a vector on Gaussian random numbers with mean $\mu=0$, standard deviation $\sigma=1$.

The use of synthetic observations and background was intended to ensure the retrieval scheme was functioning correctly, independently of the accuracy of the assumptions about their error characteristics. In this case the radiosonde profiles are taken as the *truth* and do not contribute any errors to the validation, as illustrated schematically in Figure 4-15.

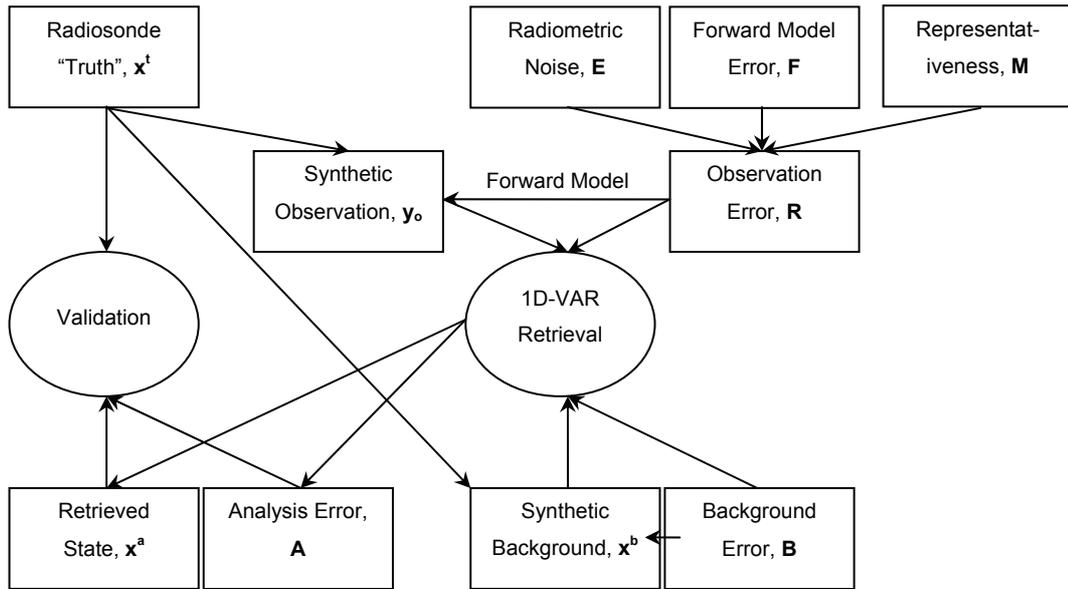


Figure 4-15 Schematic 1D-VAR with synthetic observations and backgrounds

4.9.1.1 Without Cloud Classification

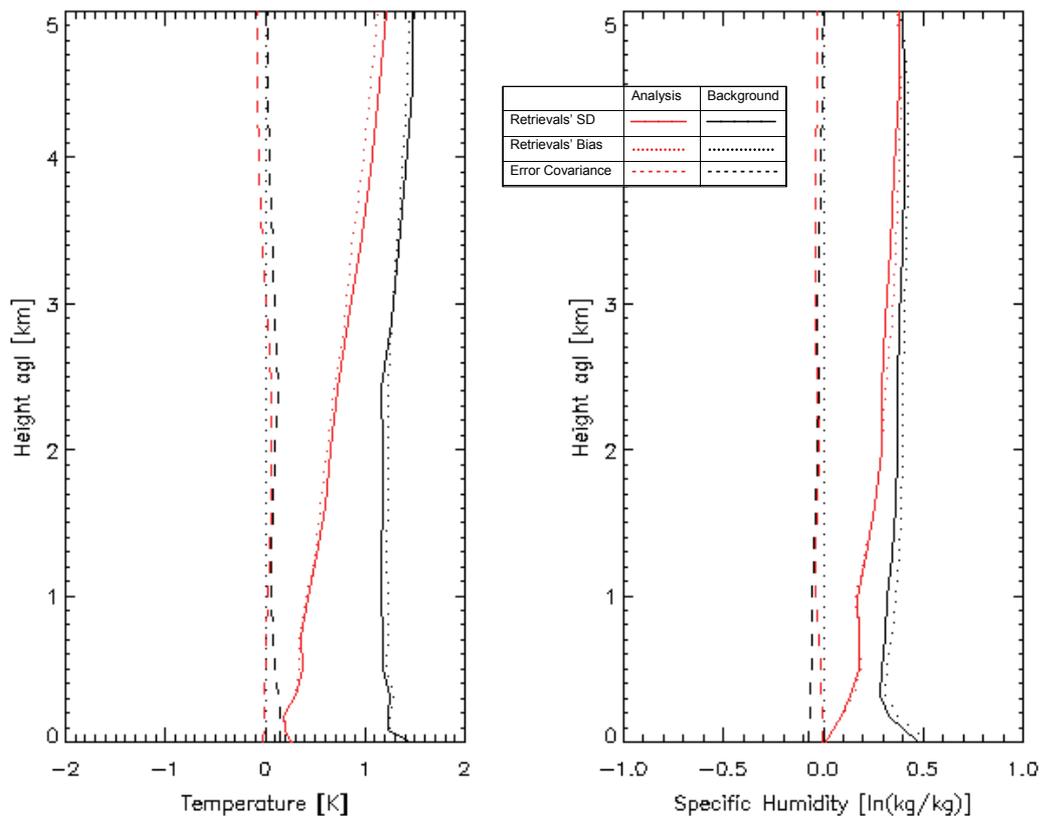


Figure 4-16 Statistics of 1D-VAR retrievals using synthetic observations and background. 484 cases from Camborne over 1 year. Dashed lines show the mean bias. Solid lines show standard deviation of difference between retrieved profiles and *truth*. Dotted lines show the diagonal terms of error covariances, $\sqrt{\text{diag}(B_{\text{ATOVs}})}$ and $\sqrt{\text{diag}(A)}$. Black lines show the background. Red lines show the statistics of the analysis.

Initially, all the retrievals were processed with no cloud classification. 75% of cases converged in an average of 8.1 iterations. The standard deviations of the background and analysis profiles, shown respectively as the solid black and red lines in Figure 4-16, follow closely the theoretical expectations from $\sqrt{\text{diag}(\mathbf{B}_{\text{AToVS}})}$ and $\sqrt{\text{diag}(\mathbf{A})}$ (shown as the dotted lines). Both temperature and humidity retrievals show improvements on the background up to 5 km. Surprisingly, the background profiles for both temperature and humidity show a small bias. This is attributed to those backgrounds which are slightly biased with respect to the truth being more likely to converge and pass the quality control χ^2 test, due to nonlinearity in the retrieval. This bias is corrected in the analysis by the retrievals. For example, if some levels in the true profile almost saturated, clear background profiles are more likely to converge than cloudy ones.

Closer investigation revealed that most of the cases that did not converge were affected by cloud. These were subsequently analysed separately from the *clear* cases.

4.9.1.2 With Cloud Classification

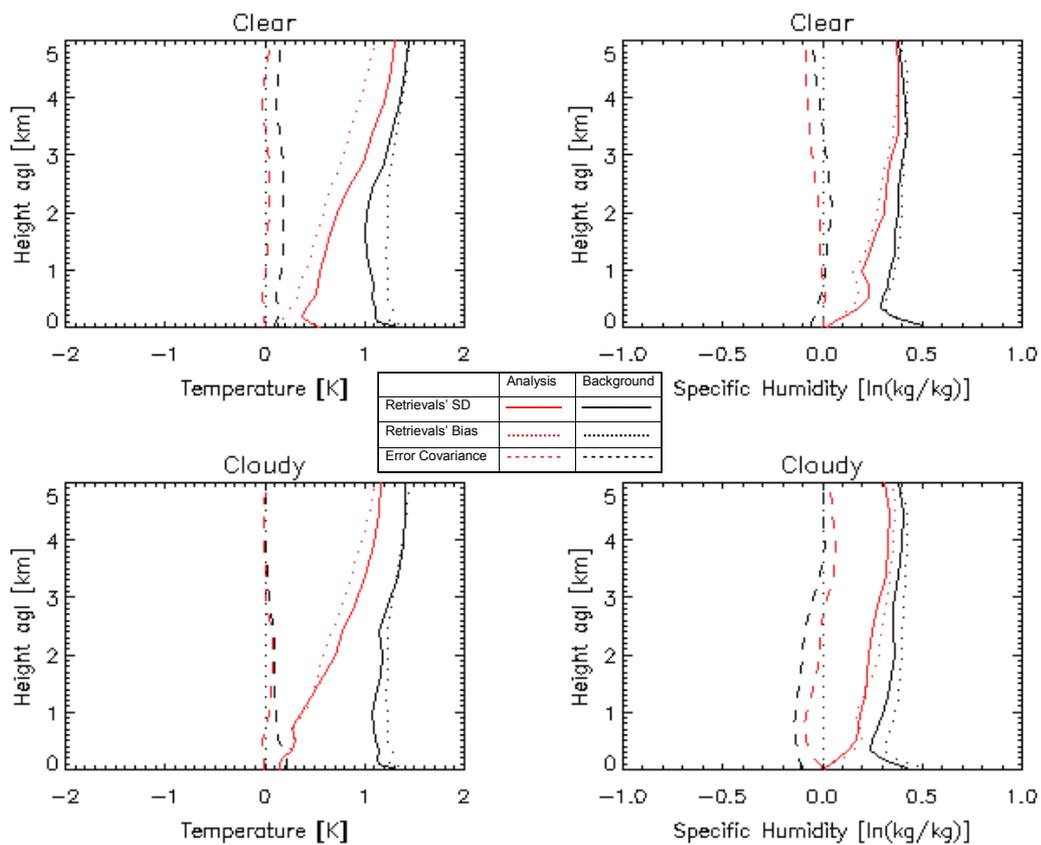


Figure 4-17 As Figure 4-16, but for 334 clear (upper) and 141 cloudy (lower) cases

Figure 4-17 shows the statistics of the profiles retrieved from the same data set and technique as applied in the previous section, but with the addition of the cloud classification scheme. In

cases classified as *clear* this tightens the observation error covariance, switches the humidity control variable to $\ln q$ and includes an explicit super-saturation constraint, described in §4.8.1.

Of the clear cases 74% converge and satisfy the χ^2 test in an average of 4.7 iterations. The standard deviations of their analysis and background profiles show less contrast than expected from **A** and **B**, respectively. i.e. There is less difference between the red and black solid lines in Figure 4-17 than between the dotted lines. This shows the backgrounds are not as diverse as expected from **B** and that the retrievals are not adding as much information as expected by the difference of **A-B**. It is likely that this is due to a deficiency in the super-saturation constraint. The analyses can be improved towards $\sqrt{\text{diag}(\mathbf{A})}$ by reducing the χ^2 threshold at the expense of reducing the number of cases. But this does not affect the temperature statistics of \mathbf{x}^p , which remain smaller than $\sqrt{\text{diag}(\mathbf{B}_{\text{ATOVs}})}$. Given the problems encountered with the super-saturation constraint, it was not used in the other experiments.

Although the cloudy cases take longer to reach convergence (10.3 iterations on average), and only 56% of cases converge and pass the χ^2 test, the standard deviations almost exactly match those expected from **A** and **B**. Again, the background shows a small, but consistent bias with respect to the truth, which is corrected in the analysis.

4.9.1.3 Comparison with Gauss-Newton method

The same data set of synthetic observations and backgrounds was re-processed with the Gauss-Newton minimisation method, instead of Levenberg-Marquardt. This resulted in very similar statistics to the Levenberg-Marquardt method, and a slightly higher convergence rate in the same number of iterations for clear conditions. However, in cloudy conditions, the convergence rate dropped to an unacceptably low 15% (albeit taking 2 fewer iterations on average) and produced similar standard deviations, albeit with slightly larger biases. This shows the benefit of the Levenberg-Marquardt method in cloudy cases, thanks to its better performance in nonlinear problems, so it is used for the remainder of the 1D-VAR retrieval experiments.

4.9.2 Synthetic observations with mesoscale background

When using \mathbf{B}_{NMC} instead of $\mathbf{B}_{\text{ATOVs}}$, which has much larger humidity errors at high level, it became apparent that the method of synthesising \mathbf{x}^p could generate unphysical, super-saturated background profiles, which degraded the retrievals due to the excessive nonlinearity. This is not the case when real backgrounds are taken from an NWP model (§4.3.1). These backgrounds were corrected for the mean bias found with respect to co-located radiosondes and the same retrieval technique described in the previous section (§4.9.1.2) was applied, using \mathbf{B}_{NMC} instead of $\mathbf{B}_{\text{ATOVs}}$, but with no clear sky super-saturation constraint, as illustrated in Figure 4-18.

This results in a much better convergence rate (98%), requiring fewer iterations: an average of 4.7 iterations in clear cases and 8.2 in cloudy cases. This clearly shows the increased nonlinearity in cloudy cases requires more iterations. However, the retrievals' statistics in Figure 4-19 show they have a much smaller impact – improving the background only in the lowest 3 km for temperature and only marginally for humidity. This is due, in part, to an excessive error variance in the temperature profiles of the background and analysis at ~1.5 km, corresponding to the typical height of the boundary layer top, due to their limited vertical resolution. The disappointing humidity retrievals show how difficult this is given their large assumed background errors. This highlights the sensitivity of the conclusions to assumptions about the errors in the background.

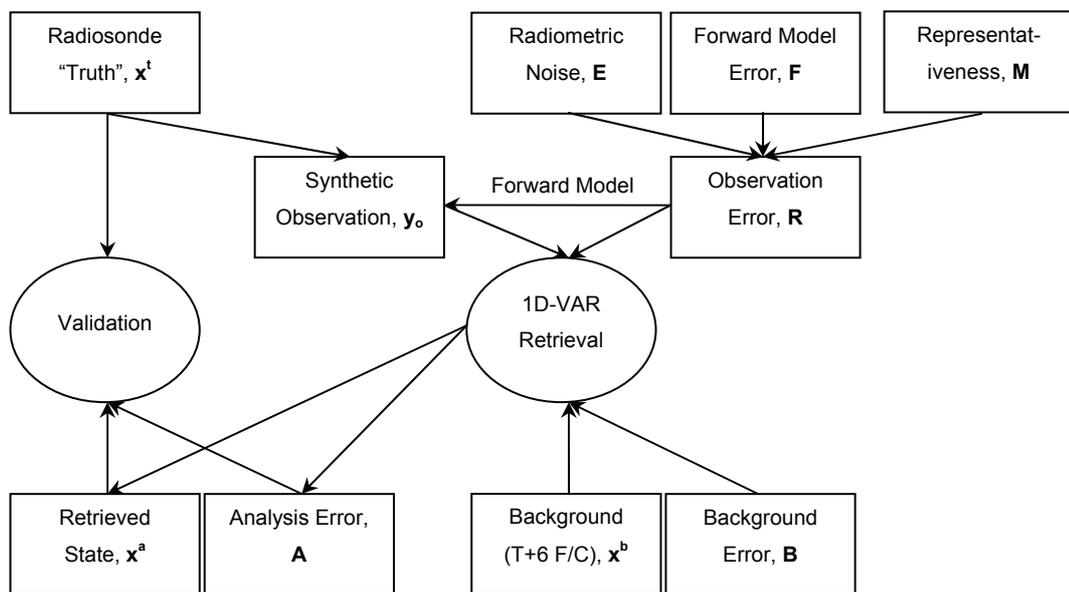


Figure 4-18 Schematic 1D-VAR with synthetic observations and background from T+6 hr forecast from NWP model

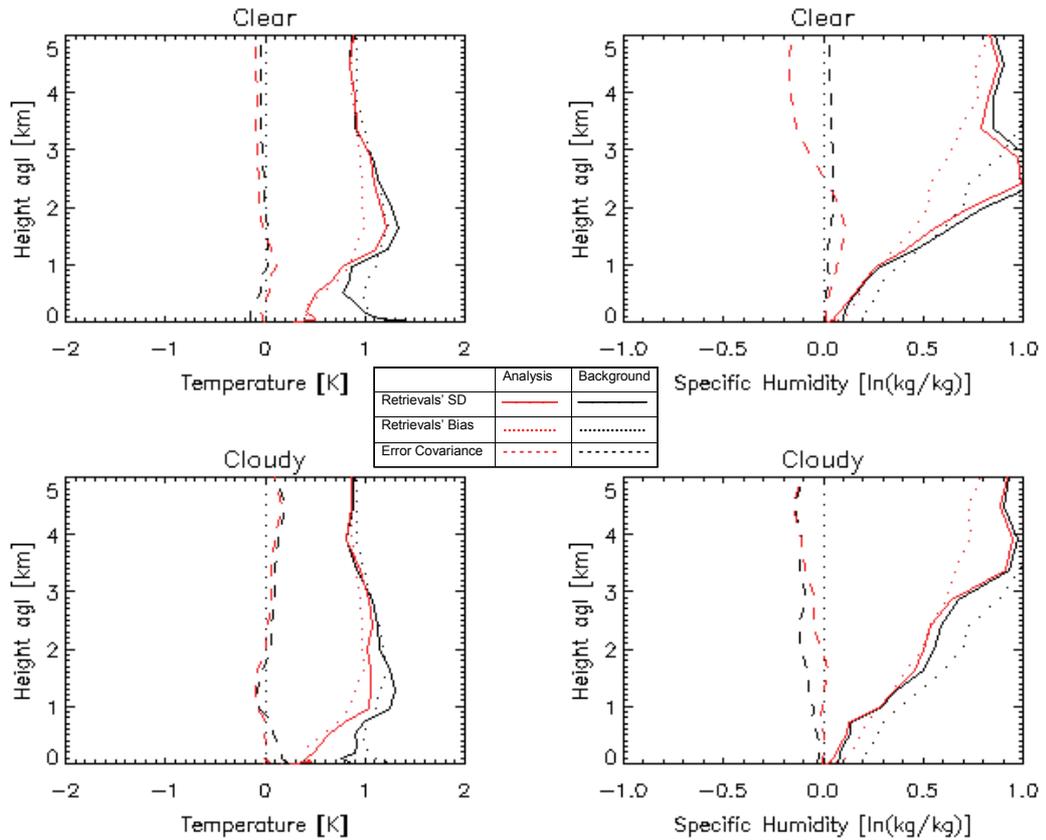


Figure 4-19 As Figure 4-17, but with real backgrounds from mesoscale model and B_{NMC} . 449 clear (upper panels) and 242 cloudy (lower panels) cases.

4.9.3 Results with modified first guess

The same retrieval technique was repeated using first guesses based on mesoscale model backgrounds modified according to the cloud classification (see §4.8.2). This gave almost identical results and convergence rates. This suggests that the retrievals are rapidly adjusting the profiles to balance the humidity modifications, as shown by the statistics of the IWV retrievals, which remained unchanged ($SD=0.80 \text{ kg/m}^2$). However, their efficiency may be improved by increasing the sophistication of the first guess modification scheme. In the results (not shown), the bias in the clear humidity retrievals is worse than with the unmodified background, however, their SD is a little better in the cloudy cases. The temperature retrievals are little changed, and still have high SD (and bias in cloudy cases) at the typical level of the inversion ($\sim 1.5 \text{ km}$).

4.9.4 Results with real observations

Radiometer observations unaffected by rain were available at Camborne within 300 s of the validation time for approximately 54% of the cases during the 2004/05 period. This reduces the number of cases available to validate the retrievals and adds noise to the statistics. When applying the same retrieval technique described in §4.9.2 to real observations, it was clear from the residuals that the observations were not consistent with the background. Before the

retrievals can be meaningfully interpreted the observations have been empirically corrected by the mean bias found in the O-B analysis of clear cases (see §4.6.3). Now the radiosonde is no longer regarded as the *truth* but an independent observation used to validate the retrievals and includes errors, R_{rs} , as shown in Figure 4-20.

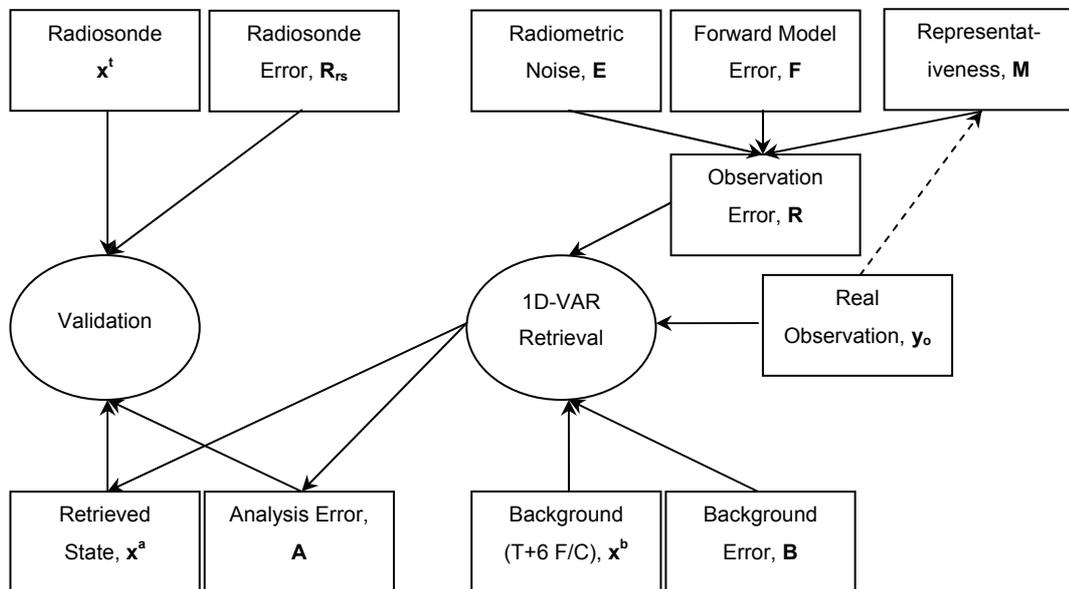
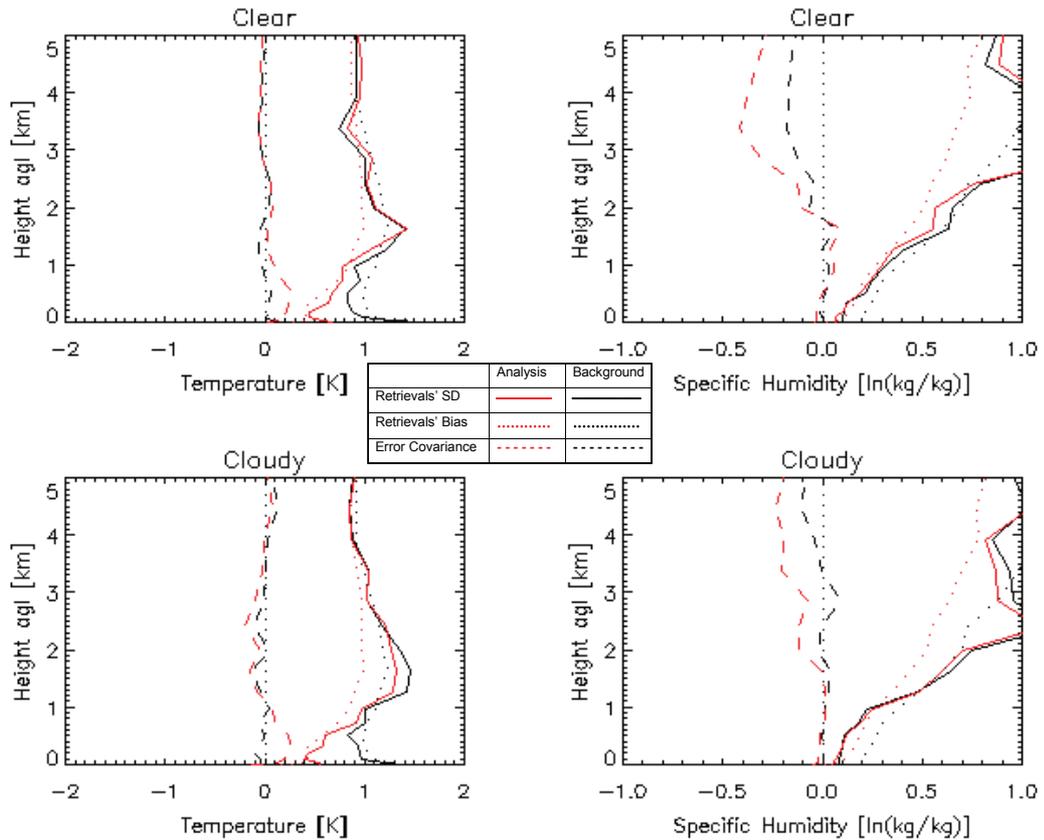


Figure 4-20 Schematic 1D-VAR with real observations and NWP background.
(Dashed arrow indicates time series of observations used to calculate M dynamically.)



**Figure 4-21 - As Figure 4-19, but with real, bias corrected observations
142 clear (upper panels) and 206 cloudy (lower panels) cases.**

These retrievals with real observations have a slightly lower converge rate (91%) than with synthetic observations, and take an average of 5.8 iterations to converge. Figure 4-21 shows the SD of the temperature profiles retrieved from real observations are close to those from synthetic data, as are the humidity profiles in clear cases. However, above 2 km there is a bias and very large variance in the retrieved humidity profiles. Both the background and analysis again show excessive variance compared to theoretical expectations at 1-2 km for temperature, but also at 2-4 km for humidity.

When retrievals were attempted with (bias-corrected) real observations using $\mathbf{B}_{\text{ATOVs}}$, although a very high convergence rate was achieved, the analyses were no better than the background. This again highlights the importance of using an appropriate background error covariance.

4.9.5 Results Dynamic Representativeness

Although radiometer observations are available at high time-resolution ($\sim 1/\text{min}$), only spot measurements closest to the time of the radiosondes used for validation have been used in

the retrievals so far. It is possible to exploit the high time-resolution to some extent by using the variability of the time-series of observations within 1 hr of the observation to calculate the representativeness error dynamically, as described in §4.4.1.3. This was implemented for the same data set and retrieval technique used in §4.9.4, still using only instantaneous observations.

It was hoped that the use of dynamic representativeness errors would allow the retrievals to be more accurate in situations where the atmospheric conditions are changing slowly. However, the results using the dynamic representativeness error (not shown) are very similar to those using fixed values for clear and cloudy cases, although fractionally higher convergence rates were achieved in cloudy conditions. This is not necessarily bad news, as the errors can be overestimated in conditions of broken cloud by this method.

4.9.6 1D-VAR retrievals of Integrated Water Vapour (IWV)

In addition to validating the retrievals against radiosonde profiles of temperature and humidity, it is also possible to compare the Integrated Water Vapour (IWV) in the retrieved profile with that from the radiosonde. As the IWV is merely the integral of the absolute humidity over the vertical profile, it does not represent any additional information. Nevertheless comparison of the retrievals' and radiosondes' IWV provides a check on the retrievals and allows comparison with other retrieval techniques.

With synthetic observations, the profiles retrieved by 1D-VAR produced IWVs with a standard deviation of $\approx 0.8 \text{ kg/m}^2$ compared to the radiosondes taken as the truth. When real observations were used, this increased slightly to $\approx 0.9 \text{ kg/m}^2$. These figures should be compared to the standard deviation of the background profiles' IWV compared to the truth of $\approx 2.2 \text{ kg/m}^2$. As will be shown in §5.5, the 1D-VAR retrievals provide IWV with comparable accuracy to regression techniques. It is, therefore, unlikely that there would be any benefit in adding an additional term in the retrieval's cost function to constrain it to fit an *a priori* estimate of IWV, as the variational retrievals do a good job of this implicitly because if the retrieval is optimal for the humidity profile, it is also optimal for IWV.

4.9.7 Example retrievals with a misplaced background inversion

Figure 4-22 shows an example of 1D-VAR retrievals using synthetic observations, based on a real radiosonde profile for Camborne (UK) at 1121 UTC on 9/12/2004 and NWP background profile from a 5 hr forecast, valid 21 minutes earlier. This case was selected because the model had forecast the inversion $\sim 200 \text{ m}$ too low and overestimated the humidity by a factor of ~ 2 over most of the profile, while matching the observed surface conditions. The retrieval was repeated for 105 such sets of observations, all of which converged in 4 iterations.

The retrieved profiles are closely clustered, with typical standard deviations of 0.2-0.5 K in temperature and 0.05-0.10 in $\ln q$, showing they are relatively robust in the presence of observation noise. Both the background and retrieved profiles have a temperature inversion at ~500 m, where initially the humidity approaches saturation and liquid water cloud is generated. In all cases, the retrieval removes this cloud, reduces the IWW from the background value to match the radiosonde's and gives profiles closer to the truth than the background. However, the structure of the background error covariance, \mathbf{B} , makes it impossible for the retrieval to move a misplaced inversion vertically, without additional information. When the first guess was modified by increasing the temperature from 500-600 m by 4 K, the retrieval was still unable to reproduce this, as it is strongly constrained by the background. If the observations had more information on the inversion height, the background would not constrain the retrievals so strongly in this way.

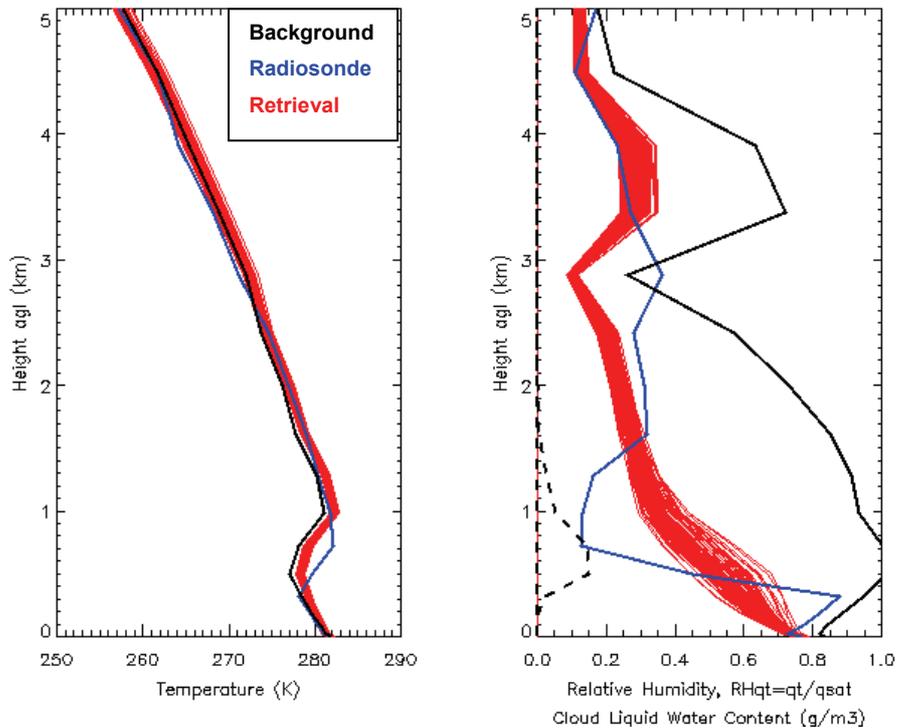


Figure 4-22 Example retrievals (red) with 105 synthetic observations, with profiles between NWP model background (black) and radiosonde (blue). Left panel shows temperature profiles. Right panel shows profiles of relative humidity ($RH_{qt}=q_t/q_{sat}$) and cloud liquid water content [g/m^3] (dotted lines).

4.9.8 Summary of 1D-VAR Retrieval Experiment Results

Retrievals were first tested with synthetic backgrounds and observations. Results matched theoretical expectations in clear skies, but poorer convergence was experienced in cloudy conditions, due to greater nonlinearity. The Levenberg-Marquardt method was found to improve on the performance of the Gauss-Newton method in these cases. Two methods were attempted to further improve the convergence in cloudy cases by making the forward model more linear between the first guess and the solution. However, there was an unidentified problem with the super-saturation constraint used for clear cases in the cloud classification scheme. An alternative approach, to modify the humidity in the first guess to be consistent with the infrared observations resolved this.

Synthetically generated background states were found to be prone to being unphysical. However, taking the background data from NWP models reduced the impact of the synthetic observations on temperature to the lowest 3 km, even after bias correction. Real observations also need bias correction, after which they gave results broadly consistent with synthetic observations. The temperature profiles retrieved from real observations were in line with theoretical expectations, but the humidity retrievals did not improve the background state much. This is a common problem encountered when trying to assimilate observations of humidity. Atmospheric water vapour is highly variable in space and time and not well represented in current NWP models, which have to parametrize the processes responsible for its sources and sinks. This means they are unable to provide accurate background data for humidity profiles and it is more difficult to accurately describe their error covariance, so it is more difficult for the non-linear humidity retrievals to converge on the truth. It is hoped that this situation will improve somewhat with the advent of convective-scale NWP models.

The observations in channels sensitive to cloud are limited by representativeness errors, which can be calculated dynamically from time-series of observations to tighten up the retrievals' errors in stable conditions. However, this did not change their statistics on average.

Although the retrievals rely on the correlations in the background error covariance to resolve the *ill-posed* nature of the retrieval problem, these correlations make it impossible for a misplaced feature in the prior profile to be moved in the vertical. This is inherent in the way the 1D-VAR retrievals are defined on a fixed grid. However, the retrieval problem can also be defined as minimising the position (height, temperature) of significant points in the profile [Berkels *et al.*, 2004]. The positioning of features in the retrieved profile could also be improved by the addition of other constraints in the cost function, based on other *a priori* knowledge. However, this cannot be achieved by modifying the first guess alone. For example, we could use the height of the maximum signal from a UHF wind profiling radar, which is related to the gradient of the refractive index, indicating levels with strong humidity changes [Gaffard and Nash, 2006]. These concepts are discussed further in Chapter 7.

Chapter 5 Neural Network Retrievals

This chapter investigates the application of artificial neural networks to retrieve temperature and humidity profiles from ground-based microwave radiometers. This is a fundamentally different technique from the variational methods discussed in the previous chapter and offers different advantages and disadvantages which are reviewed in this chapter, along with a statistical assessment of the neural network developed by Solheim *et al.* [1996].

5.1 Introduction to Neural Networks

The neural network is essentially a nonlinear statistical regression between a set of predictors – in this case the observation vector, \mathbf{y} – and a set of predictands – in this case profiles of atmospheric temperature, humidity and liquid water content, \mathbf{x} . This regression of $\mathbf{x} = f(\mathbf{y})$ is implemented as a series of empirical nonlinear functions, $\mathbf{g}(\mathbf{y})$ in any number of layers of *hidden nodes* or *neurons* used to combine each of the predictors (*input nodes*) to each of the predictands (*output nodes*), as illustrated in Figure 5-1. For example, $f(\mathbf{y})$ can be a nonlinear weighted sum of the form

$$f(\mathbf{y}) = K \left[\sum_i w_i g_i(\mathbf{y}) \right] \quad (5.1)$$

where K is a pre-defined function, such as a hyperbolic tangent.

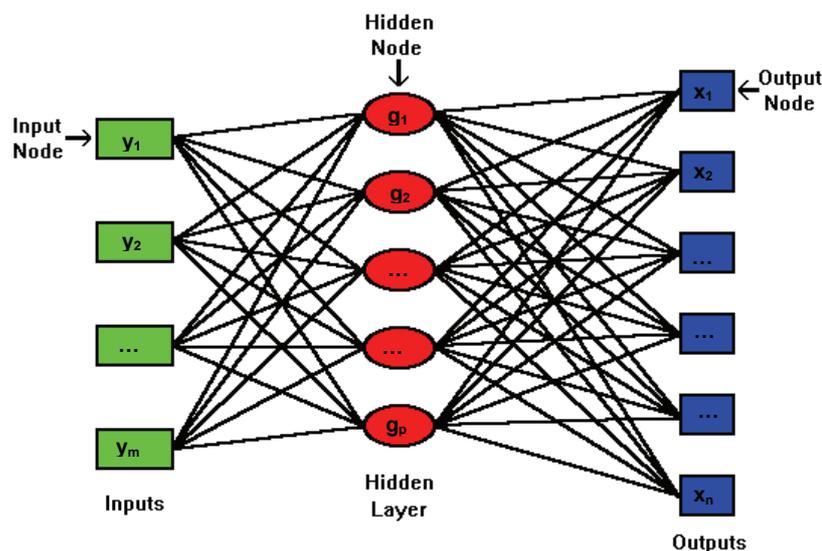


Figure 5-1 Schematic diagram showing a Neural Network with 1 hidden layer fully connected between a set of input nodes, \mathbf{y} , and output nodes, \mathbf{x} .

The coefficients of these functions are determined *a priori* by a process of *training* using a separate data set of input and output nodes. There is some debate about whether neural networks outperform other regression based methods – for example, see [Solheim *et al.*,

1996], [Rose *et al.*, 2005], [Cimini *et al.*, 2006b]. However, using neural networks has the advantage of requiring none of theoretical assumptions such as only moderate nonlinearity or the Gaussian error character of the variables.

In this case, the neural network was trained using a data set of radiosonde profiles and observations simulated from them. This data set provides the *a priori* background data in the form of a climatological ensemble used to resolve the *ill-posed* nature of the retrievals. This is another key difference from the variational method of the previous chapter, which took their background data from an NWP model. The fact that a radiative transfer model is used to generate the synthetic observations in the training data set means these neural network retrievals are also *physical*.

Although the neural network is slow to train, it is very fast to perform retrievals. And since these retrievals are independent of NWP model data, they can be used to provide independent validation of forecasts or be deployed in stand-alone nowcasting applications where fast response is more important than absolute accuracy of the retrieved profiles. However, as with any nonlinear regression, neural networks are prone to generate erroneous results if applied outside the range of the training data set. This potentially serious problem is discussed in §5.3.4.

5.2 Training Neural Networks

Radiometrics [2001] implemented the neural network retrievals using the Stuttgart Neural Network Simulator [Zell *et al.*, 1993] to construct a standard feed forward network with one hidden layer and full connection between this and the input and output layers. The coefficients of the hidden nodes were trained using a standard back propagation method. This is a very computationally expensive process, but it only need be performed once for each neural network.

Separate neural networks have been trained using microwave radiometer observations in the zenith view and a range of elevation angles. In both configurations the radiometer data is combined with observations of ambient pressure, temperature, humidity and the brightness temperature from a zenith viewing infrared radiometer. The results of these are described in the following sections.

5.2.1 State space definition for neural network

The neural network is trained to output profiles of temperature [K], absolute humidity (water vapour density) [g/m^3] and liquid water content [g/m^3], as well as the integrated amounts of water vapour [kg/m^2] and liquid water [g/m^2]. The measurements of surface temperature, relative humidity and pressure are assumed to be accurate and are fed directly through the

neural network to provide the values at the lowest level of the profile. Although the absolute humidity is not commonly used in meteorological applications, it is the variable used in many absorption models and can be readily converted to relative humidity or specific humidity, which has the advantage of being conserved in adiabatic ascent/descent.

The profiles are output at discrete levels every 100 m up to 1 km, then every 250 m from 1 to 10 km. Although the number of independent measurements is less than the 47 levels output (as will be shown in Chapter 7), this sampling ensures the retrieved profiles can accurately be represented on fixed levels.

5.2.2 Training data sets

The first neural network to be tested here was trained on a data set of 10 years of radiosonde profiles from Camborne (UK). During some of this period radiosondes were launched 4 times daily (at 0515, 1115, 1715 and 2315 UTC). The launch frequency later reduced to 3 and then 2 times daily. This resulted in approximately 10,000 soundings being used in the training data set. The profiles had been discretised at the standard and significant levels of the WMO *TEMP* code (see §2.6.3). Other versions of this neural network have been trained using high resolution (~10 m) radiosonde profiles for elevation scanning retrievals, which are discussed later in this chapter.

The radiosondes provided pressure, temperature and humidity profiles. Liquid water cloud was added synthetically to the profiles, based on a simple humidity threshold up to a maximum column amount of 1500 g/m² [Randolph Ware, personal communication, 21/10/03]. Higher values will be outside the range of the training data set and liable to produce retrieval errors. The profiles were then input to a radiative transfer model to synthesise the observations of a ground-based microwave radiometer at the elevation angles and centre frequencies of the channels used in the retrieval. The use of centre frequency introduces small biases compared to observations over a finite bandwidth as described in §2.5.2. This model used the water vapour absorption from Liebe and Layton [1987] and the oxygen absorption from Rosenkranz [1993]. The input and output nodes corresponding to the surface temperature and humidity are directly connected. The brightness temperature of the infrared radiometer, T_{ir} , was calculated simply as the temperature at the lowest level of the profile with cloud.

The simulated observations in the training data set were duplicated 4 times each with uncorrelated Gaussian noise added to simulated observations [Solheim *et al.*, 1998], with standard deviations of 0.5 K for microwave brightness temperatures, 0.5 K for surface temperature, 1% for RH, 3 hPa for pressure, and 2 K for the T_{ir} . This extends the training data set slightly and reduces the sensitivity of the network to noise in the data and should represent all the errors affecting the observations. However, the analysis presented in

§4.4.1.4 shows these figures underestimate the expected variance in the channels sensitive to water vapour and cloud, due to representativeness errors when used for NWP models with ~10 km grids, and overestimate expected variance in the high frequency channels (54-59 GHz) near the centre oxygen absorption band, which are dominated by radiometric noise.

5.2.3 Training statistics

Figure 5-2 show the mean profiles of temperature, vapour density and liquid water content of the 10 year data set of radiosondes from Camborne used to train the neural network. The mean profile is featureless, showing that inversions do not form predominately at any one level during over the annual cycle (unlike wintertime in Payerne in the case of the TUC data set). It approaches the dry adiabatic lapse rate at high levels, and the saturated adiabatic lapse rate in the 1 km near the surface.

Also shown in Figure 5-2 are the standard deviation of the training data set and the r.m.s. error of the neural network retrievals applied to this data set. This shows that the retrievals are expected to improve dramatically on the background information available from climatology for temperature, reducing its uncertainty near the surface from 3.75 K to 0.4 K and from 5.4 K to 1.1 K at 2 km. The humidity retrievals reduce the relative uncertainty from 26% to 4% at the surface and from 51% to 21% at 2 km. It is also clear that the error characteristics of the cloud profile are not Gaussian, as their standard deviation is greater than the average value. However, the retrievals still improve the cloud in the background by a factor of ~2 below ~3 km.

This performance is comparable with the analysis error covariance of the 1D-VAR retrievals, shown in Figure 4-10 and Figure 4-11 up to ~2 km. However, at higher levels the 1D-VAR retrievals are expected to be more accurate than the neural network because they rely more heavily on the NWP background here, which is more accurate than the climatological one used by the neural network.

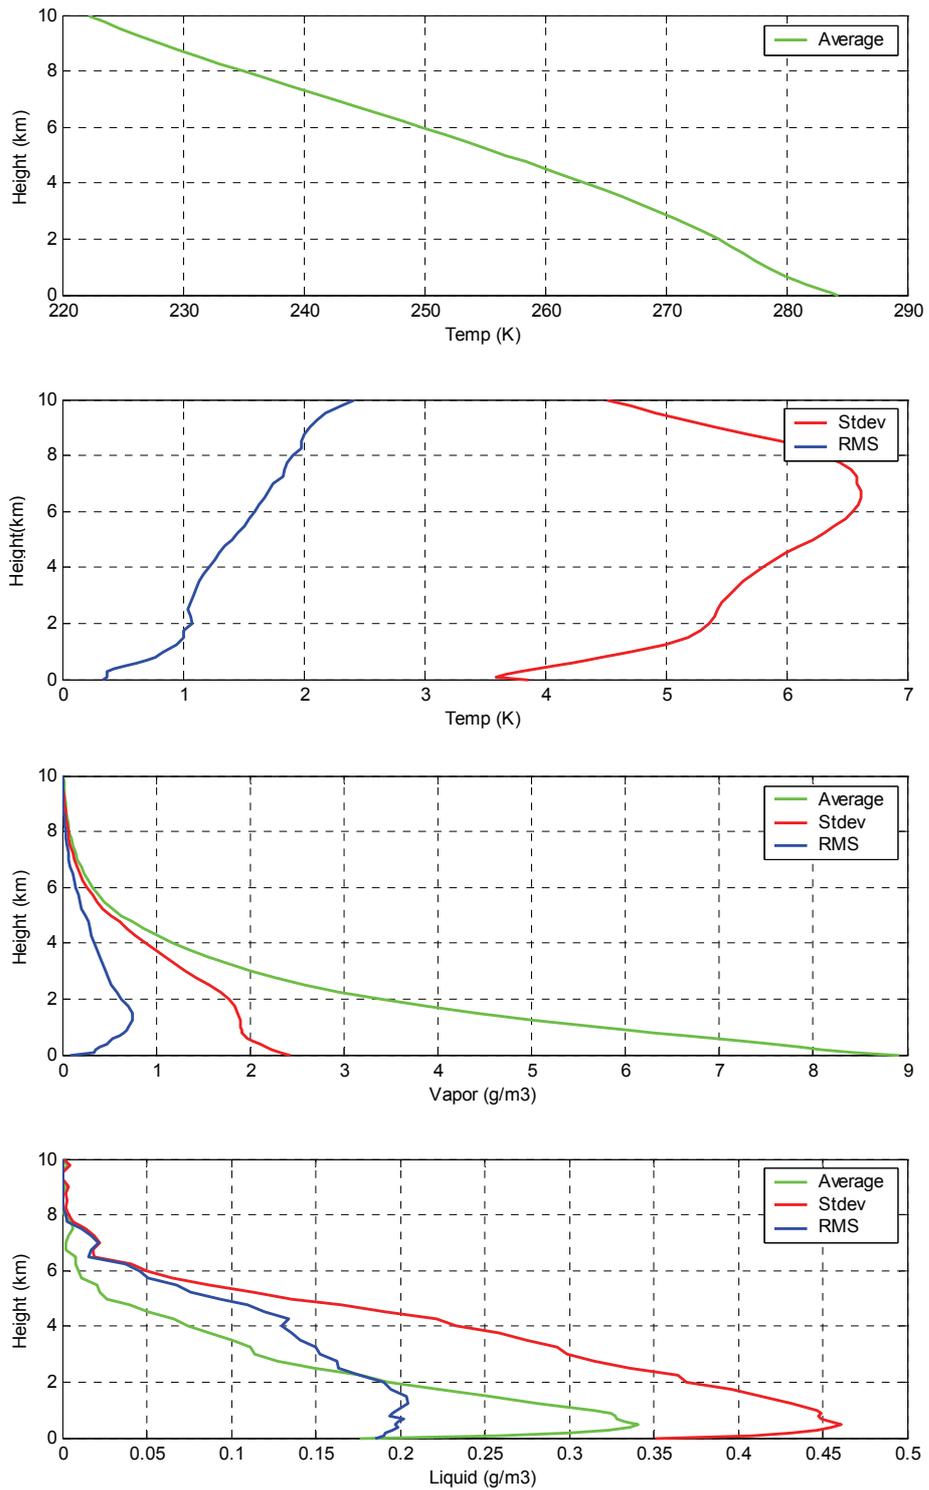


Figure 5-2 Average and standard deviation (Stdev) of 10 year radiosondes from Camborne used to train the neural network and r.m.s. error (RMS) of its retrievals. [Courtesy of Radiometrics]

5.3 Neural Network Retrievals from Zenith Observations

Originally the neural network only retrieved profiles from the zenith view of the Radiometrics TP/WVP-3000. These retrievals have been validated against radiosondes, launched from Camborne during the radiometer's first year of operation there. The radiometer retrieval closest to the time of launch of each radiosonde was compared with its temperature and humidity profile. In all cases, this resulted in a time difference of less than 9 minutes. This approach was chosen to provide the best match in time between the radiometer retrievals and the radiosondes' measurements of the lowest part of the profile, where most of the radiometer's information is concentrated.

Cases were rejected when the radiometer's rain flag was set, or when the retrieved Integrated Liquid Water (ILW) was greater than 1000 g/m^2 , on the basis of empirical evidence that these conditions were liable to produce rain (§6.5). This filtering left 703 cases over the period 20/2/2002 – 18/3/2003, for which the statistics are shown in Figure 5-3 for temperature and vapour density profiles. The humidity statistics are also shown after conversion to dew point temperature and relative humidity, assuming no errors in the temperature profile.

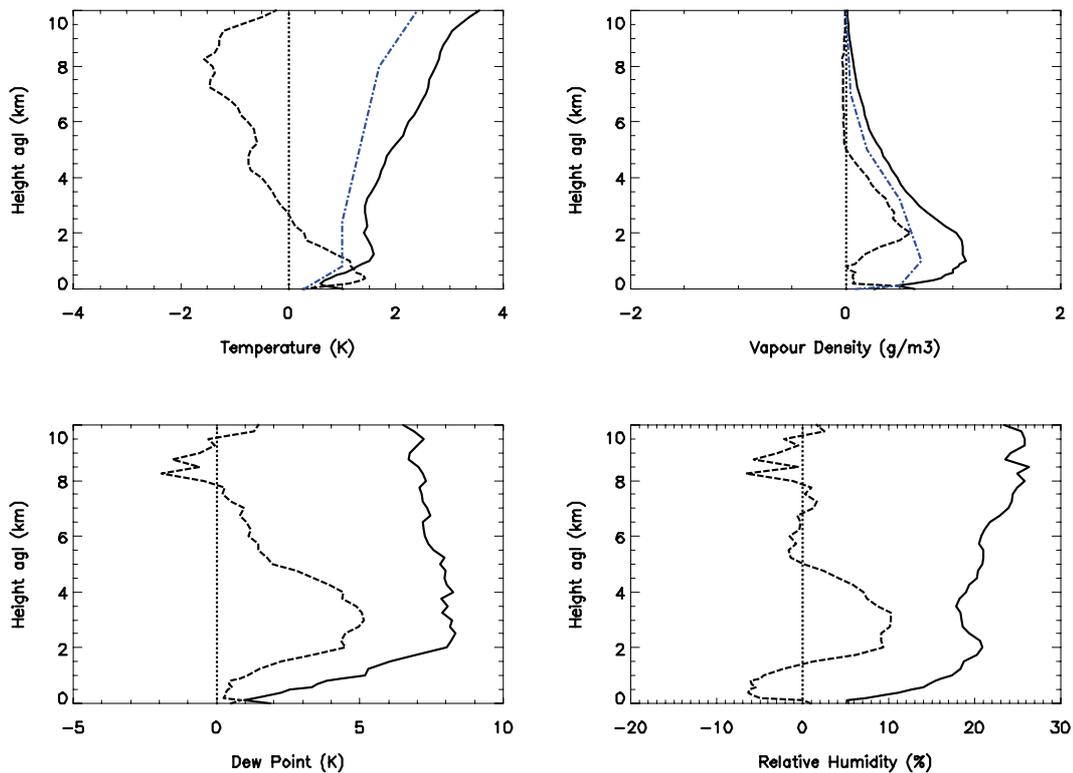


Figure 5-3 Statistics of neural net retrievals at Camborne 20/2/2002-18/3/2003.

Black solid lines show Standard Deviation of retrieved-radiosonde profiles,

Black dashed lines show mean bias of retrieved-radiosonde profiles,

Blue dot-dash lines show r.m.s. residuals of neural network training. 703 cases.

Figure 5-3 shows the standard deviation of the difference between the temperature profiles retrieved by the neural network and measured by the radiosonde increases from ~0.6 K just above the surface to ~1.5 K by 3.3 km. However, their variance is larger at the surface, and also ~1.0-1.5 km, corresponding to typical levels of subsidence inversions frequently found in anticyclonic conditions. Investigation showed the errors are non-Gaussian at this level due to the retrievals' poor resolution of these strong inversions. For vapour density the SD increases from ~0.5 g/m³ near the surface to ~1.1 g/m³ between 1-2 km, before decreasing again at higher levels as the average absolute humidity decreases in the colder air.

These results are substantially higher than the r.m.s. residuals of the training data set. This is partly due to the optimistic assumptions made about the observation error in the low frequency channels when training the neural network, but can also be caused by errors in the radiosondes' measurements associated with instrumental noise and representativeness. The later accounts for the atmospheric variability between the atmosphere seen by the radiometer and the radiosondes.

5.3.1 Bias in neural network retrievals

Where there are consistent biases between the profiles retrieved by the neural network and those measured by co-located radiosondes, this suggests an inconsistency between the radiosondes, the radiometer, the radiative transfer model or the retrievals themselves. Whilst consistent biases may be corrected empirically, it is beneficial to investigate their causes to understand the mechanisms responsible and, possibly, improve them.

Figure 5-3 also shows the mean bias in the temperature and humidity profiles retrieved by the neural network relative to the radiosondes. The neural network temperature profiles have a consistent and highly significant positive bias, peaking at +1.4 K at 400 m. (The uncertainty on the bias can be estimated as $SD/\sqrt{n} \sim 0.02$ K near the surface to 0.12 K by 5 km, where $n=703$.) This bias may be largely explained by a bias in the calibration, affecting the highest frequency channels (55-59 GHz) (§3.4.2), which give most of the information here. This bias reduces to zero by 2.5 km and becomes significantly negative at higher levels. However, the temperature bias at higher levels was found to depend on the atmospheric conditions – it was larger in cold, dry conditions. This is believed to be related to the bias found in comparison of observed and forward-modelled brightness temperatures of the 51-53 GHz channels in §3.4.2.

There is also a significant positive bias in the humidity retrievals between 1-5 km. Again, this was found to depend on the atmospheric situation – larger in warm, humid conditions – as did the observed brightness temperatures compared to forward-modelled radiosondes in §3.4.2. This is believed to be related to the known dry bias of the Vaisala RS80 radiosondes used in this study [Nash *et al.*, 1995] – see also §5.5.

If the retrievals are considered as weighted averages of a *priori* background data and the radiometer's observations, they will be biased if the climatological background is biased with respect to the average profile of a given ensemble. When considering the statistics of an experiment lasting one season or less, this can introduce significant bias with respect to the truth, $(\mathbf{x}^a - \mathbf{x}^t)$, for which the magnitude can be estimated as [Eyre, 1987]:

$$\overline{(\mathbf{x}^a - \mathbf{x}^t)} = (\mathbf{I} - \mathbf{MRM}) \cdot \overline{(\mathbf{x}^b - \mathbf{x}^t)} \quad (5.2)$$

where $\overline{(\mathbf{x}^b - \mathbf{x}^t)}$ is the mean bias of the background state with respect to the truth, \mathbf{I} is the identity matrix and $\mathbf{MRM} = \mathbf{KH}$, is the Model Resolution Matrix [Rabier *et al.*, 2002], where $\mathbf{K} = \mathbf{AH}^T \mathbf{R}^{-1}$ is the Kalman gain matrix, which can be interpreted as a generalised inverse of \mathbf{H} .

MRM has been calculated by perturbing each element of a reference profile in turn, calculating a synthetic observation, \mathbf{y} , and using this in the neural network retrieval to retrieve a perturbed profile, \mathbf{x}' . This was evaluated for the average profile retrieved from radiometer observations and measured by radiosondes from the winter-time TUC campaign (§1.5), where a neural network was specifically trained on a full year's high-resolution radiosonde data from Payerne. This results in a positive bias expected in the lowest 1 km of the retrieved profiles, peaking at ~1 K near 500 m. This may explain part of the positive bias of ~1.5 K at 500 m found in the temperature profiles retrieved during this campaign [Hewison *et al.*, 2004]. This is not expected to produce significant bias in the results shown in Figure 5-3 because the data set covered a full year.

It is also possible that a bias in the retrievals can be introduced as a result of applying their finite vertical resolution to retrieve a nonlinear feature, such as a consistent inversion at the top of the boundary layer. However, this is unlikely to be a significant factor in this data set – these inversions are typically rather small features at Camborne and the height of the top of boundary layer varies dramatically over 1 year (or even 1 day). Even in the winter-time climate of Payerne, Switzerland the neural network showed no sign of a dipole in the bias that might be expected from a poorly resolved consistent inversion [Hewison *et al.*, 2004].

5.3.1.1 Bias correction to brightness temperatures

In the assimilation of data into NWP it is common practice to correct observations for the mean bias found with respect to the forward-modelled background (§2.7.5). This approach can also be adopted to profiles retrieved by the neural network – the mean bias found in the comparison of observations to a forward-modelled reference can be subtracted from the observations prior to their input into the neural network. This has been applied to the radiometer's brightness temperatures by comparison with forward-modelled radiosondes in clear air [Hewison *et al.*, 2006] and also to its surface observations using reference observations from an operational Stevenson screen.

The bias in the high frequency channels (55-59 GHz) of ~1 K was larger than expected as the brightness temperature of these channels is typically quite close to that of the ambient black body reference target. However, it was confirmed independently by comparison with another radiometer [Rose *et al.*, 2005] and has since been corrected in the radiometer's control software. Prior to applying this correction, the statistics of the profiles retrieved during TUC closely resembled those shown in Figure 5-3 [Hewison *et al.*, 2004]. The bias corrections resulted in the ~+1.5 K bias in the retrieved temperature near 500 m being reduced and actually becoming slightly negative. The negative temperature bias above 2 km is also much reduced. However, the humidity profiles and IWV are largely unaffected by the small corrections applied to the channels in the water vapour band.

5.3.2 Neural network's sensitivity to noise

A retrieval system can be designed to minimise sensitivity to noise or to optimise vertical resolution by effectively changing the correlations in the background error covariance matrix [Rodgers, 2000]. The sensitivity of the neural network to radiometric noise can be investigated by constructing an ensemble of randomly perturbed observations and examining the variance of the profiles retrieved from them. When an observation was perturbed by uncorrelated Gaussian noise equivalent to the radiometric noise, the standard deviation of the temperature profiles was ~0.2 K and ~1% for the humidity profile. These results are comparable to those of a similar experiment with the 1D-VAR retrievals (§4.9.7) and much smaller than the expected variance of the retrievals. This shows that the neural network does not amplify observation noise and that it appears to follow the minimum variance solution as the variational method.

It is also useful to examine the impact of averaging retrieved profiles prior to comparison with radiosondes. The comparison with radiosondes shown in Figure 5-3 was repeated using the average of all profiles retrieved from unflagged observations within 30 minutes of the radiosonde launch time. This increased the number of cases available. The results in Figure 5-4 show that averaging retrieved profiles reduces the variance of the temperature and humidity above ~1 km compared to the radiosondes. Below 1 km, the averaged temperature profile has slightly higher variance than the instantaneous profiles. This confirms that some of the excess variance in the validation of instantaneous retrievals with radiosondes compared to expectations from the statistics of the training data set was indeed due to underestimated observation errors in the form of modelling and representativeness. However, the retrievals still have greater variance from the radiosondes than the variance of the training data set, as this comparison includes contributions from the radiosonde errors.

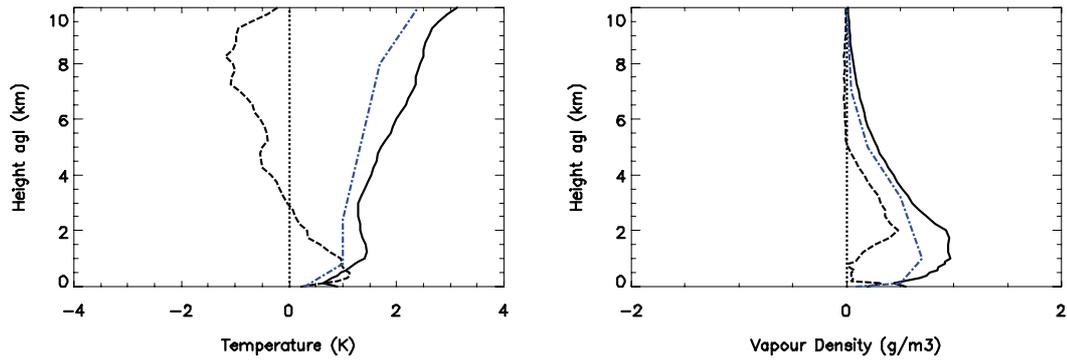


Figure 5-4 Statistics of neural net retrievals at Camborne 20/2/2002-18/3/2003.

Retrieved profiles averaged within 30 min of radiosonde launch time.

Black solid lines show Standard Deviation of retrieved-radiosonde profiles,

Black dashed lines show mean bias of retrieved-radiosonde profiles,

Blue dot-dash lines show r.m.s. residuals of neural network training. 1038 cases.

Later in this study (§7.1) the sensitivity of the retrieval accuracy to observation errors are investigated by averaging the observations prior to using them in the retrievals.

5.3.3 Linearity of neural network retrievals

The characteristics of an ensemble of neural network retrievals can also provide insight on the linearity of the retrieval problem. The mean temperature profile retrieved by the neural network over a period of 100 days at Camborne was found to match the temperature profile retrieved from the mean of all unflagged observations over this period with a SD of 0.27 K. This implies the temperature retrieval process is approximately linear. However, this is not the case for the humidity profiles, where there was an mean difference of 0.12 g/m³ and a SD of 0.15 g/m³ (8% of the mean surface humidity), which is a significant fraction of the expected retrieval errors.

Different conclusions could be drawn from applying the same analysis to the data from Payerne during TUC. The difference between the average profile retrieved by its neural network and that retrieved from the average observations had a SD of 2.3 K in temperature and 0.75 g/m³ (33%) in humidity. These are larger than the expected SD of the retrievals, showing that both temperature and humidity retrievals have significant nonlinearity for this data set because of its frequent, strong inversions.

5.3.4 Application outside range of training data set

As with all regression methods, neural networks are prone to amplify errors when applied outside the range of the training data set. However, this can be more serious in the case of neural networks because of their capability of producing a nonlinear response to inputs. For

this reason great care should be taken to ensure the training data set is not only representative of the climate of the radiometer's operating site, but also that it encompasses the full range of conditions ever likely to be observed at that site. However, extreme events will always be problematic with neural networks.

This turns out to be difficult to achieve in practice when using a set of radiosonde profiles to train the neural network. All the operational radiosonde stations in the UK have until recently been at coastal sites. Together with the limited sampling available from radiosondes launched at only midday and midnight, this means the extremes of the diurnal cycle experienced at inland sites are inadequately represented by this method. Furthermore, it is not easy to estimate likelihood of extending beyond the range of the training data set statistically as the extremes of the distribution are unlikely to be normally distributed.

The neural network trained on Camborne radiosondes was applied to radiometer observations taken from Linkenholt, an inland site, during CSIP (§1.5) in summer 2005. The 1.5 m air temperature here often exceeded 25 °C, which is the highest value in the training data set. When the retrievals from this neural network were validated against co-located radiosondes, a strong negative bias (-1.6 g/m^3) was found in the humidity profile at $\sim 1 \text{ km}$ [Hewison and Gaffard, 2006]. This bias was found to be stronger when surface temperatures were higher and is indicative of the neural network being applied outside the range of its training data set. The magnitude of this bias was much reduced (to -0.5 g/m^3) by repeating the retrievals using the coefficients for another version of the neural network trained using radiosonde data from Lindenberg, Germany. This training data set includes more extreme cases because of the continental climate of Lindenberg, compared to the maritime one of Camborne. The temperature and humidity profiles retrieved using this neural network were better than those from the Camborne neural network to the radiosondes – both in terms of bias and standard deviation [Hewison and Gaffard, 2006]. This highlights the need to ensure the data set used for training the neural network covers the full expected range of the observations, allowing for possible future climate change. This may be achieved by including profiles from other radiosonde stations, artificially extending the training data set, or by normalising its inputs and scaling the neural network so these only fall in the range 0.10-0.90 following Churnside *et al.* [1994]. However, as the coefficients of the neural network become less site-specific, there is a danger that its retrievals will be biased if the background is no longer representative of the site's climate.

It would also be possible to train the neural network using a data set of profiles from an NWP model. This has the advantage of including cloud properties and producing a full diurnal cycle and should be more representative of the climate of each station, assuming this is well represented in the model. This can be validated using other observations, such as radiosondes, or those available from *CloudNet* [Hogan and O'Connor, 2006].

5.3.4.1 Altitude of training data set

The altitude of the station used to supply the training data set should be consistent with that of the observations as emission near the water vapour line at ~22 GHz and in the wings of the oxygen band (e.g. 51-53 GHz) is sensitive to pressure. The altitude of Lindenberg is 112 m above sea level. Although this is not perfect for application at Linkenholt (altitude 185 m), it is marginally better than Camborne (altitude 88 m), and did not prove to be problematic. However, this was another motivation for training a new neural network specifically for Payerne (altitude 491 m), as those trained for lower altitude stations were expected to produce serious errors.

5.4 Elevation Scanning in Neural Network Retrievals

To improve the vertical resolution of the retrievals needs extra observations. Viewing the atmosphere at lower elevation angles increases the channels' sensitivity because of the extra path length. This effect alone has been used to retrieve temperature profiles in the lowest 600 m of the atmosphere using a simple radiometer with a single channel near the centre of the 60 GHz oxygen band observing a range of elevation angles [Troitsky *et al.*, 1993]. This technique can also be combined with observations at a range of frequency channels and has the potential to improve the vertical resolution of their retrievals under the assumption that the atmosphere is horizontally homogeneous. During the course of this study, elevation scanning has been included in a new neural network, which can be applied to data from the Radiometrics TP/WVP-3000.

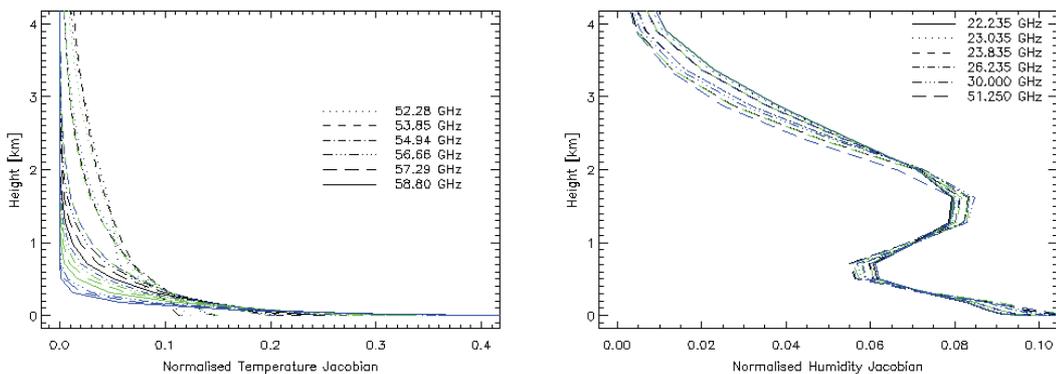


Figure 5-5 Normalised Jacobians showing sensitivity of Radiometrics TP/WVP-3000 channels viewing at zenith angles, $\theta=0^\circ$ (black), 60° (green) and 75.5° (blue) for temperature (left) and humidity (right).

The sensitivity of the radiometer's channels viewing different elevation angles to temperature and humidity changes at different levels are shown by their normalised Jacobians, plotted in Figure 5-5. These differ from the Jacobians in Figure 4-9, as they have been divided by the

integrated Jacobian for each channel to show the *relative* sensitivities. This shows that the highest frequency channels become proportionally more sensitive to temperature changes near the surface at lower elevation angles. However, the normalised humidity Jacobians have the same shape for all elevation angles, showing there is no extra information available, only a reduction in noise.

A new neural network was developed with inputs corresponding to the 12 brightness temperatures measured at 4 elevation angles (15, 20, 30 and 90°) plus the surface pressure, temperature and humidity and the infrared brightness temperature, measured at zenith. Two training data sets used high-resolution radiosondes from Camborne and Payerne, respectively, sampled every 2 s (equivalent to ~10 m) to improve the representation of small features near the surface.

These neural networks were applied to observations sequentially sampling 7 elevation angles, symmetric about zenith. Observations in opposite angles were averaged before being input to the neural network. Although this observing sequence can take 60-300 s depending on the instrument's configuration, it has the advantages of cancelling out small angular misalignments and allowing better quality control as described below. The observation errors are approximately the same as would be achieved using a single sided scan with double the integration times, as the only difference is the overhead required to move the scan mirror, which is small (~1 s).

5.4.1 Asymmetry in elevation scans

The assumption of horizontal homogeneity can be checked by comparing radiances observed in symmetric views. These were found to differ by less than the observation error in more than 90% of cases. The remainder mostly correspond to cases with broken cloud. This is a fundamental restriction to the application of the elevation scanning technique to channels at frequencies lower than ~54 GHz, which are sensitive to cloud. Another quality control check, based on a threshold of the standard deviation of the observations in the 30 GHz channel mapped back to the equivalent zenith brightness temperature was developed by Cimini *et al.* [2006b]. Retrievals using elevation scanning should not be attempted for cases failing these quality control checks. However, these scans are still of potential interest as they can provide information on the horizontal gradient of temperature and/or humidity, although asymmetries are most often caused by inhomogeneous cloud fields.

It was found that brightness temperatures measured in opposite elevation angles differed by more than $\sqrt{2}$ times the corresponding term of the observation error covariance, $\sqrt{\text{diag}(\mathbf{R})}$, on 2-11% of scans (6% on average). This is less than would be expected if the errors between different views were independent (16%) because there is a significant correlation between the components associated with the forward model, calibration noise and representativeness.

This shows the potential benefit of elevation scanning retrievals needs to be assessed using observation error covariances including all elevation angles being considered. These can be evaluated by calculating the covariance of the full observation vector including multiple views, as described in §7.1.3.

5.4.2 Results of elevation scanning retrievals

The Radiometrics TP/WVP-3000 was configured to scan seven elevation angles and its internal black body reference target in a 300 s observing sequence during most of the TUC experiment in Payerne, Switzerland during the winter of 2003/04. These observations were bias corrected, as described in §5.3.1.1, and input into two versions of the neural network – one using zenith views only and another using full elevation scans. The profiles retrieved by these neural networks from cases meeting the above homogeneity condition were compared with those from 165 co-located radiosondes and the statistics shown in Figure 5-6. Although all the profiles retrieved by the neural networks within 30 minutes of the launch time of each radiosonde were averaged, this was found to have no significant impact on the resulting statistics. This contrasts with the case of the Camborne data set, and shows impact of atmospheric variability was much less during the TUC experiment because of the more stable conditions.

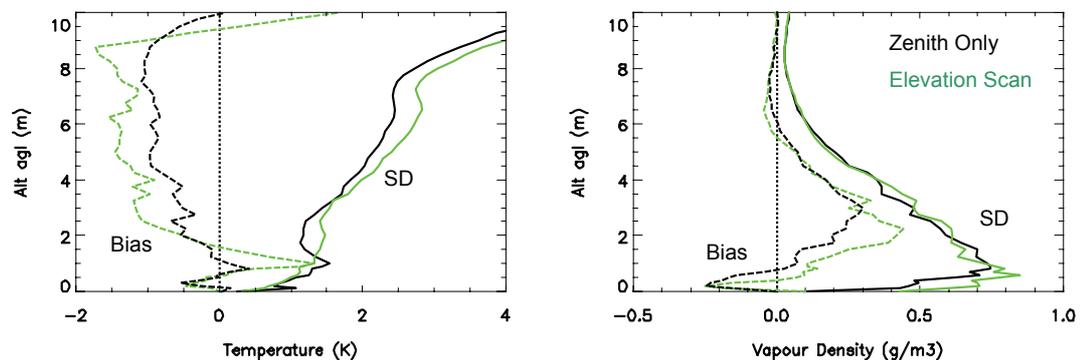


Figure 5-6 Statistics of neural net retrievals at Payerne 6/11/03-13/2/04, 165 cases.

Solid lines show standard deviation of retrieved-radiosonde profiles.

Dashed lines show mean bias of retrieved-radiosonde profiles.

Retrievals from zenith views only (black) and elevation scans (green).

Figure 5-6 shows the standard deviation of the temperature profiles retrieved from zenith observations are similar to those found at Camborne in Figure 5-4, while the humidity profiles have less variance than at Camborne, because of the lower absolute humidity found during TUC. The bias of the zenith retrievals is also smaller than found at Camborne, both for temperature – because of the bias corrections applied to the observations, and for humidity – because of the lower absolute humidity. However, the basic pattern of a negative temperature

bias above 2 km, and a positive humidity bias in mid levels remains the same, which suggests this may be due to a residual inconsistency between the forward model and the observations.

The retrievals from elevation scans only reduced the variance of the temperature profiles slightly in the lowest 1 km. At higher levels, and throughout most of the humidity profile, the elevation scan retrievals produced worse results than the retrievals using only zenith views, both in terms of variance and bias. This was not expected. If nothing else, these retrievals use more observations, so the noise on them should have reduced by a corresponding factor. However, the validation of these retrievals has already been shown not to be noise limited. This may point to an error in the retrieval process due to the lack of correlation in the observation error used to train the neural network. It could also be explained by observation errors, described below.

5.4.3 Technical limitations of elevation scans

In addition to the assumption of horizontal homogeneity, which was checked by testing the symmetry of observations at elevation angles on opposite sides of zenith, other factors may degrade observations at low elevation angles, similar to those influencing the tip curve calibrations (§3.3.4). Microwave radiometers do not measure radiation in infinitesimal beams, but have a finite beamwidth, limited by diffraction in the optical components of the antenna. Their response to emission over a range of angles becomes an important factor in using observations in the optically thin channels of the water vapour band at low elevation angles, where the scene changes view rapidly with angle. This effect can be accounted for by using effective zenith angles slightly higher than their true values, following Han and Westwater [2000]. However, this correction is not currently performed as part of the elevation scanning neural network.

Furthermore, the radiometer's field of view is not purely Gaussian and has some sensitivity to radiation emitted outside the *main beam* of the antenna. At low elevation angles, this results in sensitivity to emission from the ground. This is more likely to be problematic in the lowest frequency channels used for humidity profiling, as these have the widest beams in the near field, which are more easily truncated by components of the antenna's optics. Land surfaces typically have an emissivity of ~ 0.95 at microwave frequencies [Hewison, 1999], so will have a typical brightness temperature of ~ 285 K. Hence, the contrast between the brightness temperature of the sky and the ground is also greatest in these channels. Together with the lack of extra information available from the 20-30 GHz water vapour channels from multiple elevation angles, these factors severely limit the application of the elevation scan technique for humidity profiling.

5.5 Integrated Water Vapour Retrievals by Neural Network

In addition to humidity profiles, the neural network retrieves the column Integrated Water Vapour (IWV). This was found to be consistent with the vapour density profile integrated over 0-10 km (within $\pm 0.8 \text{ kg/m}^2$) as expected because there is very little water vapour at higher levels. Data from high precision GPS (Global Positioning System) receivers at fixed positions can also be used to estimate IWV by measuring the phase delay of signals transmitted from multiple GPS satellites [Bevis *et al.*, 1992]. These IWV estimates have been compared to those calculated from radiosonde profiles.

Figure 5-7 shows the difference between the IWV measured by radiosondes and retrieved by the neural network using radiometer observations nearest the launch time over its first year of operations at Camborne. The radiometer retrievals gave $\sim 8\%$ higher IWV than the Vaisala RS80 radiosondes, which have been widely reported to have a dry bias [e.g. Nash *et al.*, 1995, Cimini *et al.*, 2003a]. This figure also shows that the agreement between the retrievals and the radiosondes is better in clear conditions, because the channels that are used to retrieve IWV are also sensitive to the variability of cloud liquid water. During most of this period, the radiometer sampled only once every 14 minutes, which will be shown to limit the accuracy of the comparison.

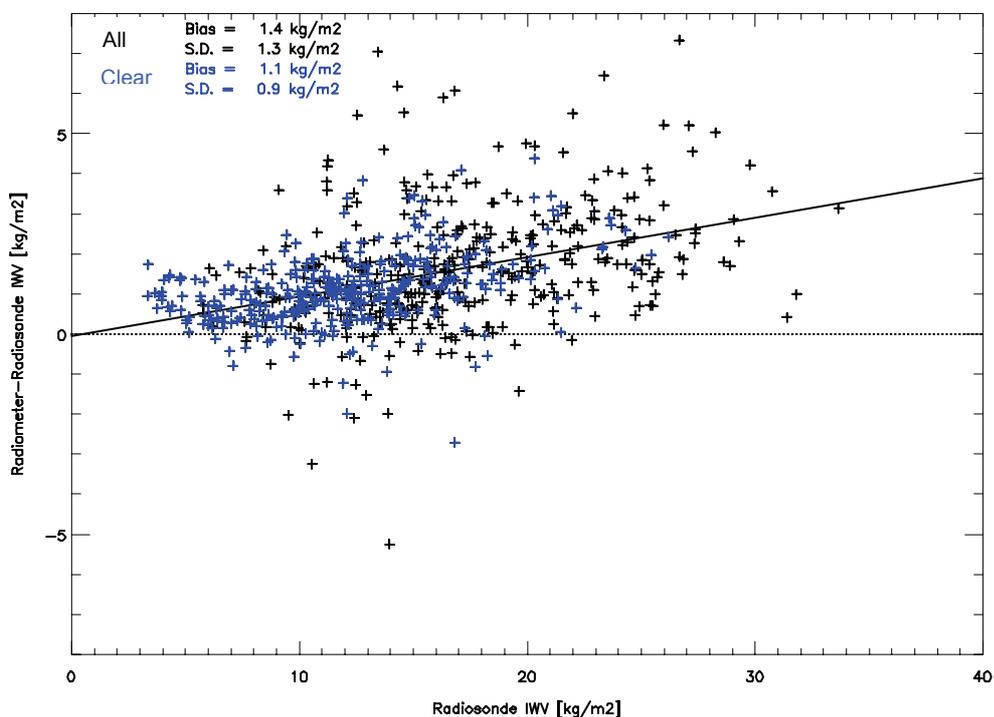


Figure 5-7 Difference between IWV retrieved using neural network and calculated from radiosonde profiles at Camborne 20/2/02-18/3/03. Solid line shows best fit by linear regression. Black crosses are all cases, Blue crosses are the subset of clear cases.

Table 5-1 summarises the results of various comparisons between IWV retrieved from microwave radiometers, measured by radiosondes and GPS systems during this study. These results are typical of those from other experiments, some of which are shown in the last two rows. The initial results from Camborne shown in Figure 5-7 have the largest standard deviations as they cover a full year of operations, and have the slowest observation cycle, resulting in the largest time differences between radiometer and radiosonde observations. Later studies showed the radiometer retrievals could identify a diurnal cycle in the dry bias of the radiosondes, and that this was larger for those using Vaisala RS92 sensors due to inadequate correction for solar heating of the sensors [Smout *et al.*, 2005]. There was no significant diurnal changes in the difference between the radiometer and GPS, but their variance was much larger in day-time, due to the greater variability of IWV. Typically, the radiometer retrievals had lower variance than GPS when compared to radiosondes (not shown). These examples show how robust IWV retrievals from microwave radiometers can be used to validate, and, ultimately, improve humidity observations from other systems.

Table 5-1 Summary of Results of IWV comparisons between Microwave Radiometers, Radiosonde and GPS

Reference	Radiometer - Radiosonde kg/m ²		Radiometer- GPS kg/m ²		Ave IWV kg/m ²	Sonde Type	Retrieval	Data set
	Bias	S.D.	Bias	S.D.				
Gaffard and Hewison [2003]	+1.4 +1.1	1.2 0.9	-0.3	1.4	16.5	RS80	Camborne NN (clear only)	Camborne 2002
Smout <i>et al.</i> [2005]	+0.4 +0.4	1.0 0.9	-1.2	1.4	22.3	RS80 RS92	Camborne NN	Camborne May 04
	+0.4 -0.1	0.7 0.6	-1.2 -1.1	1.6 0.8	Day Night	RS80	Camborne NN	
	+1.5 -0.4	0.9 0.6			Day Night	RS92	Camborne NN	
Martin <i>et al.</i> [2006]	+0.1	0.9	+0.6	0.7	10.2	Srs400	Payerne NN	Payerne, Winter 03/04
Hewison and Gaffard [2006]	+0.9 +0.4	0.6 0.7				RS92	Camborne NN Lindenberg NN	CSIP (UK) Summer 05
Mattioli <i>et al.</i> [2005]	+0.1	0.7	+0.6	0.6	~14	RS90	2 channel regression (ave)	ARM SGP, Spring 03
Van Baelen <i>et al.</i> , [2005]	+0.7	1.3	+1.0	2.7	Day	RS90	Toulouse NN	Toulouse, Sept 02
	+0.4	1.4	+0.2	2.7	Night			

5.5.1 Atmospheric variability of IWV

To understand the differences between the results of IWV intercomparisons it is useful to consider its spatial and temporal variability. This is calculated in this section following the same approach used to estimate representativeness errors in §4.4.1.3.

Treuhart and Lanyi [1987] constructed a statistical model of water vapour fluctuations by approximating the spatial structure of refractivity fluctuations by Kolmogorov turbulence theory and assuming temporal fluctuations are caused by spatial patterns advected by wind. This is known as the persistence of *frozen* spatial structure. They expressed this model in terms of spatial and temporal *structure functions* (§3.5.4) of the *Wet Delay*, which is proportional to IWV. For time scales, $\tau \ll h/v$, where h is the scale height of the wet troposphere and v is the typical wind speed at this height, they found the Time Structure Functions, $TSF \propto \tau^{5/3}$, whilst on longer times, $\tau \gg h/v$, $TSF \propto \tau^{2/3}$. At Camborne, the annual mean $h=0.9$ km, where $v=11.2$ m/s, giving the scale break at $\tau \sim 80$ s. They describe the two scales as being “consistent with the intuitive picture that many small irregularities contribute to the short distance structure, while a small number of larger irregularities dominate the long distance structure”. Their model is compared to the observed TSF of IWV in Figure 5-8.

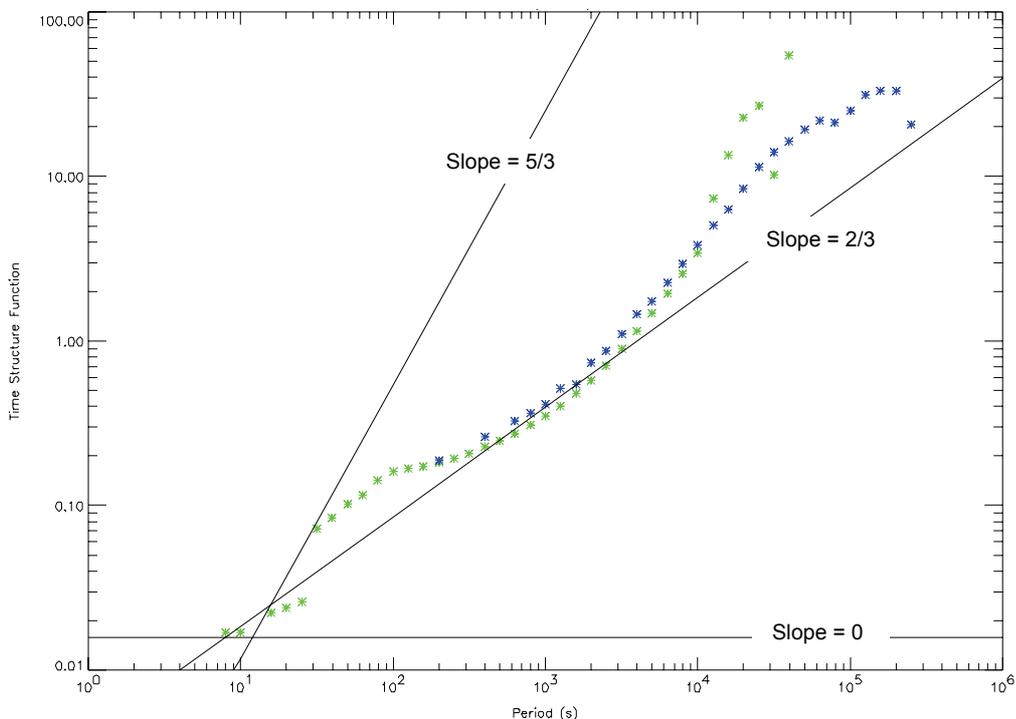


Figure 5-8 Time Structure Functions of IWV (kg/m^2) retrieved from Radiometrics TP/WVP-3000 (blue) and WVR1100 (green) during Cardington trial of Spring 2003. Lines show theoretic slopes of 2/3 and 5/3, and 0 (instrument noise).

Figure 5-8 shows close agreement between the Time Structure Function (*TSF*) of IWV retrieved from two co-located radiometers, on timescales of ~2 minutes to ~3 hours. The variability of IWV continues to decrease down to timescales of ~10 s, although the WVR1100 radiometer shows larger variance in the region of ~80 s, corresponding to the typical scale break for Camborne. The data set used to generate this figure covered 22 days including several convective days, but the *TSF* changed vastly from day-to-day and is typically larger in daytime than at night. The form of the structure functions matches the theoretical predicted slopes of Treuhaft and Lanyi [1987]. However, this could be more simply interpreted as humidity fluctuations being driven by diurnal cycles and convective eddies, causing local maxima in the *TSF* on time scales of $\tau \sim 43200$ s and ~100s respectively.

The horizontal line in Figure 5-8 indicates the white noise spectrum associated with WVR1100 instrument noise. This is equivalent to an r.m.s. uncertainty in the retrieved IWV of 0.09 kg/m^2 , which is consistent with the expected performance of a simple bilinear regression retrieval (1.1), based on a radiometric noise, $\Delta T_{\min} = 0.1$ K. Structure in the IWV is evident above the instrument noise level on time scales longer than ~30 s. If the instrument noise were reduced, it would improve the sensitivity to IWV on shorter time scales.

Figure 5-8 also shows that increasing the time separation between IWV observations, increases the expected variance between them, due to atmospheric variability. For example, if observations are separated by 15 minutes, such as the typical interval between GPS and radiosonde observations, atmospheric variability will increase the standard deviation of the IWV difference by 0.35 kg/m^2 . This is equivalent to a horizontal separation between observations of ~10 km, assuming the advection of frozen spatial structures. This explains a lot of the differences between the statistics shown in Table 5-1.

5.5.2 Comparing IWV from radiometer and GPS in precipitation

Because GPS measurements of IWV are not expected to be substantially biased in rain rates less than ~50 mm/hr [Solheim *et al.*, 1999], they can be compared with those from microwave radiometers to investigate the effects of rain on the radiometer's performance. The difference between these estimates of IWV have been compared with the precipitation rate during TUC in Figure 5-9, which shows the radiometer tends to overestimate IWV in the presence of precipitation. This bias increases with rain rate and becomes statistically significant at precipitation rates of >1 mm/hr.

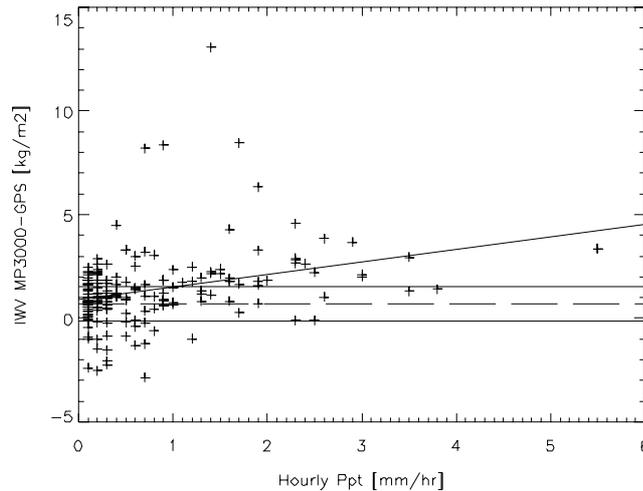


Figure 5-9 Bias in IWP retrieved with Radiometrics TP/WVP-3000 neural network relative to GPS compared to precipitation rate for all wet hours during TUC. Horizontal lines show mean and standard deviation of the difference in IWP during dry conditions. Sloping line indicates linear regression of observations.

5.6 Integrated Liquid Water retrievals by Neural Network

Figure 1-1 shows that liquid water emission follows a broad continuum, which increases with frequency and dominates the signal in the 30-50 GHz region. So microwave radiometer channels here are sensitive to the column Integrated Liquid Water (ILW), as described in §1.3.1. Although the neural network outputs profiles of liquid water content, there is very little information on its vertical distribution – this is only provided by the cloud base information available from the infrared radiometer’s observations and through its correlation with humidity.

Crewell and Löhnert [2003] analysed the expected performance of ILW retrievals from microwave radiometers alone and in combination with other instruments. They found a two channel microwave radiometer with noise characteristics similar to the Radiometrics TP/WVP-3000 should retrieve ILW with an r.m.s. accuracy of 35 g/m². These retrievals were most improved by the addition of a 90 GHz channel, which has greater sensitivity to liquid water, but adding a 50 GHz channel improved the r.m.s. to 20 g/m². Further improvements came from the inclusion of information on the cloud base height or temperature from a ceilometer or infrared radiometer. Their results suggest the neural network should be able to retrieve ILW with an accuracy of ~20 g/m².

5.6.1 Validation of ILW retrievals and NWP models

The absence of routine *in-situ* measurements of cloud liquid water content makes it very difficult to validate ILW retrievals. However, they have been validated against aircraft

measurements on a campaign basis. For example, Wood and Taylor [2001] found radiometer retrievals of ILW closely matched the adiabatic liquid water content for maritime stratocumulus cloud, except where drizzle was present. In fact, microwave radiometry is widely regarded as providing the most accurate retrievals of ILW [Crewell *et al.*, 2002] and used as a reference for other systems and to validate NWP models. For example, van Meijgaard and Crewell [2004] described a comparison of liquid water cloud predicted by different NWP models with high-resolution ground-based observations.

It is, however, possible to validate the ILW retrievals in clear skies, when there is known to be no liquid water. This can be determined independently using a threshold of infrared brightness temperature, as described in §4.8.1. Examining the ILW retrieved in clear conditions revealed a standard deviation of 17 g/m^2 – in line with expectations. However, the ILW retrievals had an asymmetric distribution caused by the neural network's inability to return negative values of ILW [Hewison, 2003]. Although these would be unphysical, it is expected that half of all clear sky cases should return slightly negative values of ILW because of noise on the radiometer observations. Furthermore, the clear sky ILW was found to have a positive bias in humid conditions of high IWV because of the infrared radiometer's sensitivity to emission by atmospheric water vapour (§2.8). As infrared absorption is not accounted for in the forward model used to train the neural network it is misinterpreted as high level cloud and biases the ILW retrievals. This technique has been further developed to adjust the calibration of microwave radiometers to ensure $ILW \approx 0$ in clear sky conditions [Gaussiat *et al.*, 2006]. In patchy cloud the variability of the infrared brightness temperature also generates errors in the neural network ILW [Guldner and Leps, 2005].

5.6.2 Problems with ILW retrievals

There are several potential sources of error in the ILW retrievals from the neural network – a full review of which is given in Crewell and Löhnert [2003]. The retrievals are sensitive to errors in the absorption models used to train the neural network, including a great uncertainty regarding the absorption of super-cooled liquid water, which may introduce a bias of $15\text{--}30 \text{ g/m}^2$ in ILW retrievals [Marchand *et al.*, 2003]. Typically, ILW shows a correlation length of about 5 min, determined by a high resolution radiometer (1 s sampling at 1° beamwidth) [Rose *et al.*, 2005]. Using broader beams and/or longer integration times may not resolve the full variability of ILW. Furthermore, time and beamwidth differences between sampling each channel may introduce additional errors in the ILW retrieval. For the Radiometrics TP/WVP-3000, with a 4 s delay between sampling the 22.235 and 51.25 GHz channels that provide the ILW information and a beamwidth $\sim 4^\circ$, this is estimated to introduce errors of up to 20% for clouds with base $\sim 1 \text{ km}$ following the analysis of Crewell *et al.* [2001]. This will dominate over the random error in ILW due to radiometric noise for thick cloud layers $> 100 \text{ g/m}^2$ (the global average).

Scattering by rain drops or wet snow is not included in the forward model or the state space of the neural network retrievals. In situations where this becomes significant it is expected to bias the retrievals. Czekala *et al.* [2001] pointed out that the amount of liquid water in rain can approach that in the cloud and that microwave radiometers are sensitive to emission from both sources. Löhnert and Crewell [2003] calculated that ILW retrievals deteriorate when the integrated total of cloud and rain liquid water amount exceeded $\sim 700 \text{ g/m}^2$ and that for $\text{ILW} > 400 \text{ g/m}^2$ droplets often become large enough to produce significant scattering at 90 GHz and introduce a sensitivity to drop size distribution.

5.6.3 Linearity of cloud fraction averaging

As cloud can introduce large errors of representatives due to its great variability, it may be beneficial to average observations in cloudy conditions prior to using them in retrievals. A simple experiment was conducted to investigate the implications of doing this.

The linearity of the cloud retrievals has been investigated by comparing the retrieval of a uniform layer of thin cloud with that for an equivalent layer of thicker, broken cloud to ascertain the optimum sampling strategy. A radiosonde profile was selected from Camborne with a $\sim 1 \text{ km}$ thick layer near saturation. Initially, the forward model was run on this profile with no cloud to generate a synthetic observation vector. Then cloud was added, uniformly distributed over this layer in increments of 50 g/m^2 up to 500 g/m^2 and the forward model repeated for each modified profile. To investigate whether 100% cloud cover with thin cloud (50 g/m^2) is equivalent to 10% cloud cover with thicker cloud (500 g/m^2), another case was created representing the linear average of 9 sets of synthetic observations in clear skies and 1 in dense cloud (500 g/m^2). The mean brightness temperatures, $\langle T_b \rangle$, in this case were found to be identical with those for the case of uniform thin cloud within their errors (r.m.s. difference = 0.15 K), as illustrated in Figure 5-10.

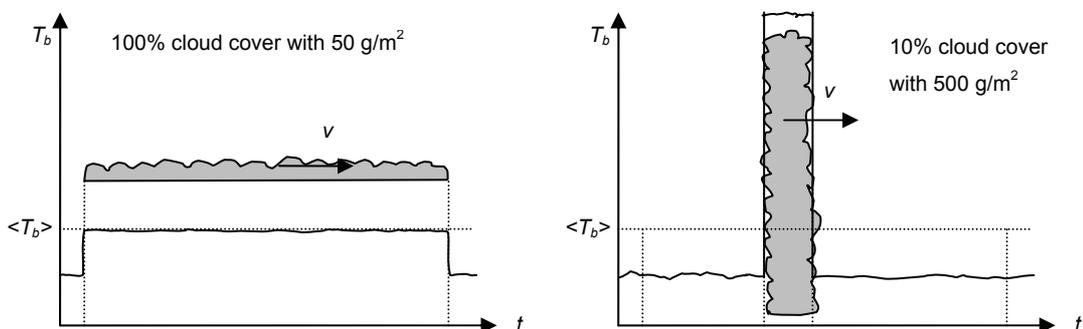


Figure 5-10 Schematic time series of brightness temperatures viewing synthetic clouds

The neural network was applied to each set of synthetic observations and the retrieved profiles compared with the profile retrieved from the clear sky case. The case of 100% cloud cover with 50 g/m^2 was found to give the same small biases as the case of 10% cloud cover with 500 g/m^2 . The r.m.s. differences of their temperature and humidity profiles were 0.05 K and 0.3%RH, respectively. This is much smaller than their difference from the clear sky case (0.28 K and 3.6%RH). The neural network's ILW retrievals were also identical in both cases.

The result of this experiment shows the effect of cloud liquid on the radiative transfer is linear. So, in situations with broken cloud, observations can be averaged in time before being used in the retrievals without introducing bias in the retrieved average profile or ILW. However, this may not be the most efficient use of the observing cycle, as will be discussed in Chapter 7.

5.7 Summary of Performance of Neural Network Retrievals

The statistics of the profiles retrieved by the neural network reflect the limitations in the background data: they have a large variance to cover the annual range of conditions, and corresponding biases as annual average conditions do not represent typical conditions in most cases.

Although the neural network retrievals were originally designed for application to zenith views only, they have recently been extended to apply to observations from a range of elevation angles (§5.4). This offers the potential to improve their vertical resolution. However, a number of practical problems were encountered, including the sensitivity to the alignment of the instrument and the homogeneity of the atmosphere. This effectively limits the application of the elevation scanning retrievals including low frequency channels to clear sky conditions. However, absorption by oxygen at frequencies above $\sim 54 \text{ GHz}$ is so strong that clouds do not contribute significantly to the observations here, so there is scope for elevation scanning to be applied to improve the vertical resolution of temperature profiles in the lowest 1 km.

The performance of the neural network retrievals was validated against profiles from co-located radiosondes (§5.3). The standard deviation of these retrievals was slightly higher than expected from the statistics of the training data set, due to their underestimation of the observation errors. There were also significant biases found in the retrievals. The temperature below $\sim 2 \text{ km}$ was consistently too warm, while it was often too cool at higher levels. The former was found to be due to a bias in the calibration of the high frequency channels of the radiometer, and has been corrected empirically. The latter may be related to a bias found in the validation of the forward model (§2.7.2). The humidity profiles were also found to have state-dependent biases, which may be due to inconsistencies between the observations and the forward model. Although retrievals using the elevation scan were found to be beneficial in some cases, the statistics showed no significant improvement to the temperature profile and

substantial degradation to the humidity profile, probably due to the sensitivity of the low frequency channels at low elevation angles to surface emissions.

The neural networks have also been trained to retrieve the integrated water vapour (IWV) and liquid water (ILW) amounts. The IWV has also been validated against co-located radiosondes and found to agree with a standard deviation of $\sim 0.9 \text{ kg/m}^2$ and a small bias. It has also been compared with IWV derived from co-located GPS sensors, with similar statistics. This raises the question of the benefit of the low frequency channels over and above the information available from cheaper GPS and surface sensors, which will be addressed in §7.1.4.

5.7.1 Pros and Cons of neural network retrievals

The key advantages of the neural network are that it allows profiles to be retrieved from observations in very little computer time and that they are independent of NWP models, so can be considered as independent validation of their representation of the evolution of the atmosphere.

The independence of NWP can also be considered a disadvantage of this neural network, as it has to rely on a statistical climatology to provide a background profile, which is not as accurate as a short-range forecast. This causes the neural network retrievals to be less accurate than 1D-VAR taking their background from NWP, especially above 1 km [Cimini *et al.*, 2006b]. However, there are a number of other disadvantages of the current implementation of the neural networks. The training process is computationally expensive and requires a very large data set of profiles. Because of this it is difficult for the neural network to adapt to situations where any of the channels become unusable because of instrument faults or interference. It also introduces errors if it is applied to locations other than the site of the training data set. The retrievals are prone to generate unpredictable results if applied outside the limited range of the training data set.

5.7.2 Future Improvements to Neural Network

The accuracy of the neural network's retrievals could be improved by including a better forward model for the infrared brightness temperature, and accounting for the microwave channels' finite bandwidth and their finite beamwidth when using elevation scanning. The elevation scans should also include a check on the assumption of horizontal homogeneity by comparing symmetric views in opposite directions.

The neural network should also include quality control checks on whether the observations are within the range of the training data set and whether they are consistent with the retrieved profile. These could be implemented as χ^2 tests, similar to those described in §4.7.5, comparing the distribution of the residuals of the observations and the forward-modelled

retrieved profile with its covariance. Initial tests showed most of the 1% of cases with $\chi^2 > 100$ were suspicious, while all those with $\chi^2 > 1000$ were clearly erroneous.

As suggested in §4.8.2 the first guess for variational retrievals can be chosen arbitrarily, although it is not always beneficial to do so. For example, Cadeddu *et al.* [2006] suggested using the profile retrieved from the neural network as a first guess for a variational retrieval. They claimed it is necessary for the first guess humidity profiles to be within 20% of the truth for the variational retrievals to converge. However, the results above show the profiles retrieved from the neural network are no better than the NWP backgrounds above ~1 km, so this is unlikely to be beneficial where NWP data is available.

It is also possible to improve the background used by the neural network to make it more representative of the profile expected in each season and time of day. This could be achieved by training a new neural network including the date (as day number of year) and time as additional inputs. The range of the training data set could also be extended by including profiles from more radiosonde stations at similar altitudes, taking care to maintain an average climate representative of the radiometer's site. It could also be extended artificially, by adding noise to the training data. Another approach would be to use NWP model profiles to train the neural network, as this should include a wider range of conditions, although it will also cause the retrievals to reproduce any biases in the training data set. However, there is a trade-off when selecting the training data set. If it is more specific to a site or season the neural network may give better retrievals, but only over a more limited range of conditions. On the other hand, the validity range can be improved by using a broader training data set at the expense of producing blander retrievals.

Chapter 6 Case Studies

This chapter uses several case studies to illustrate the benefit of retrievals from ground-based microwave radiometers for nowcasting and research applications. These examples are case studies and necessarily somewhat qualitative in nature compared to the statistical analysis of the neural network's performance in Chapter 5. Retrievals from observations of the Radiometrics TP/WVP-3000 microwave radiometer by the neural network have been used for speed. Variational retrievals may show different benefits due to their better background data.

6.1 Retrievals in Convective Conditions

With the advent of convective-scale NWP models that can explicitly resolve convective processes instead of parametrizing them, it is pertinent to examine in this section the radiometer's ability to provide observations on these scales.

6.1.1 Example of Convection – Cardington 8 April 2003

During Spring 2003, the Radiometrics TP/WVP-3000 was operated at Cardington – an inland site – and radiosondes launched frequently on 5 case study days [Hewison, 2003]. Profiles retrieved from radiometer observations by the neural network for a 12 hr period on 8/4/03 are shown in Figure 6-1 as time-height plots from 0-2 km. Although the retrievals have a poor vertical resolution, their high time-resolution helps reveal fine structure and identify trends in the data.

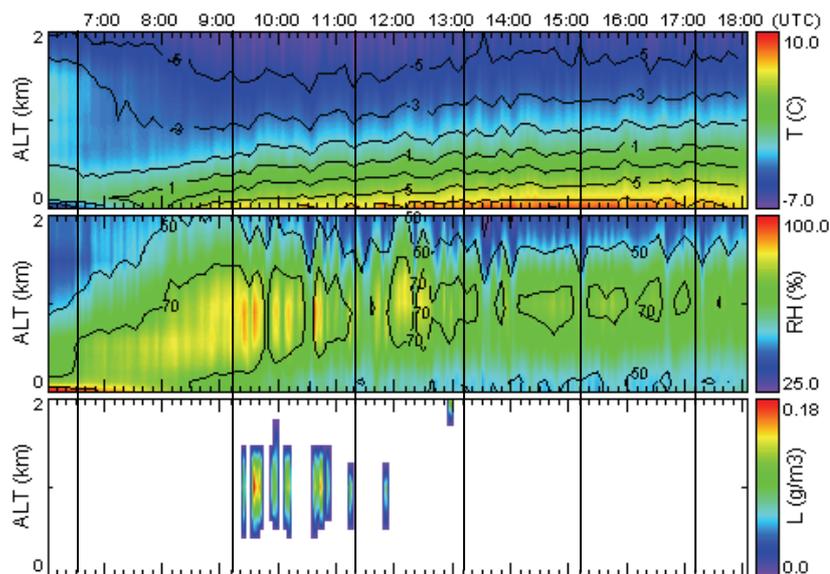


Figure 6-1 Time-height plots of profiles retrieved from radiometer observations at Cardington over 12 hours on 8/4/03. Upper panel shows temperature (°C), middle panel shows relative humidity (%RH), lower panel shows liquid water content (g/m³).

Vertical lines indicate times of radiosondes plotted in Figure 6-2.

This day started with a surface inversion. Solar heating warmed the surface, resulting in an adiabatic profile (below 1 km) for most of the day, with a very shallow super-adiabatic layer at the surface during the early afternoon. The upper panel of Figure 6-1 shows the lapse rate increasing rapidly below 1 km during the morning. This initiated convection, causing the relative humidity to increase around 1 km and decrease near the surface, forming small cumulus clouds after 0917 UTC. The cloud base then increased slowly from 1.2 to 1.7 km. The humidity and cloud retrievals show high temporal variability.

On this day, the retrievals were validated against hourly radiosondes, a selection of which are shown in Figure 6-2. The neural network retrieved the initial surface inversion quite well. It then correctly retrieved adiabatic profiles in the lowest 1 km for the remainder of the day. However, the retrievals failed to resolve the capping inversion at the top of the boundary layer (~1.5 km), which limited the vertical extent of the cloud development. For example, considering the 1122 UTC case, the radiosonde profile shows a 5 K temperature inversion and strong hydrolapse at the top of the boundary layer at 1.6 km due to anticyclonic subsidence. This is poorly resolved by the radiometer retrievals. They correctly identify a well mixed layer near the surface with temperature decreasing with height following the dry adiabatic lapse rate and the dew point decreasing more slowly with height as the RH increases and the specific humidity remains fairly constant. However, above 1 km the radiometer retrievals fail to resolve the inversion and produce less steep lapse rates for both temperature and dew point, resulting in errors of up to 4 K and 10 K respectively. The retrievals also spread the super-adiabatic layer at the surface (e.g. at 1514 UTC) too deeply in the vertical as a result of the correlation between these levels in the background data. This shows how poorly the neural network retrievals perform when there are *interesting* features in the profile.

Profiles forecast from the mesoscale NWP model run initiated at 0000 UTC are also plotted on Figure 6-2. These show generally excellent agreement with the radiosonde profiles, although the humidity in the lowest 1 km was initially too low, but too high later in the day once convection was established. The forecast temperatures also showed small errors at 0917 UTC and 1713 UTC, due to the model seriously under-estimating the magnitude of the nocturnal surface inversion (see Figure 6-4). The mesoscale model forecasts represent the inversion at ~1.5 km well as this is a synoptically forced feature. It is expected, therefore, that variational retrievals using the model for background would perform better than the neural network in this case.

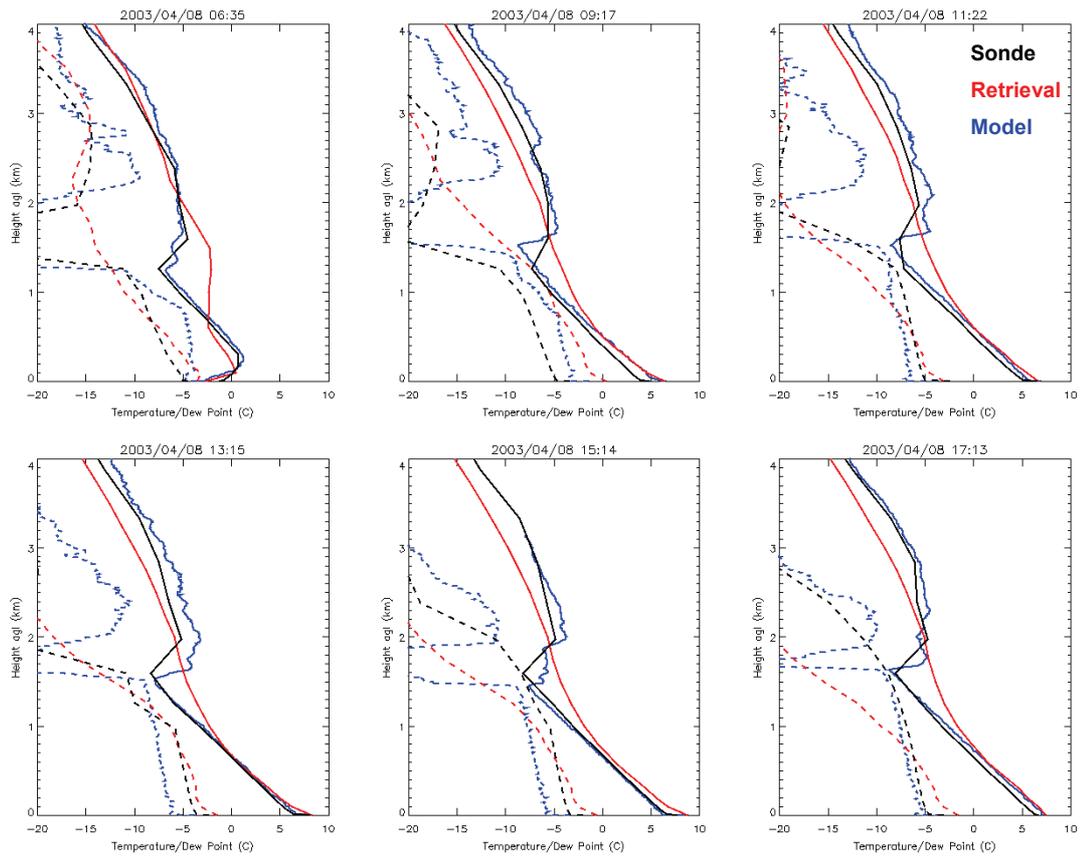


Figure 6-2 Profiles of temperature (solid lines) and dew point (dotted lines) measured by 6 radiosondes (blue), forecast from mesoscale NWP model (black) and retrieved from radiometer observations (red) at Cardington on 8/4/03.

6.1.2 Static stability

The potential of the neural network to retrieve information on the static stability is investigated in this section. The static stability describes the buoyancy of a parcel of air when displaced relative to its environment. In unsaturated air, it can be determined by examining the vertical gradient of the *potential temperature* profile, $\partial\theta/\partial z$. More generally, the gradient of the *virtual potential temperature* should be considered, as the virtual temperature is proportional to the density of air and its gradient determines the stability. Here the *wet-bulb potential temperature*, $\partial\theta_w/\partial z$, has been examined, as θ_w is conserved in reversible moist adiabatic processes. Its calculation requires profiles of both temperature and humidity, and can be derived from profiles retrieved from the radiometer's observations. However, these gradients represent differences of noisy retrievals, and, hence they are prone to larger errors.

Figure 6-3 shows an example of θ_w evaluated over the same 12 hr period shown in Figure 6-1. Initially, the profiles were stable with θ_w increasing with height. But as the surface warmed, $\partial\theta_w/\partial z$ became negative in the lowest 500 m around 0800 UTC, showing convective instability. During the rest of the day, there were short periods of convective instability

extending to above 1 km, corresponding to the build up and release of Convective Available Potential Energy (CAPE). This has also been found in single column modelling of convective processes [Xie and Zhang, 2000]. However, it is difficult to accurately estimate the CAPE in a retrieved profile, because this is very sensitive to fine-scale features, which are not well represented by the retrievals' poor vertical resolution. (The CAPE was small in this case due to the strong capping inversion.) Profiles of θ_w calculated from the hourly radiosondes, shown in the right-hand panel of Figure 6-3, follow a similar pattern to those calculated from the radiometer retrievals. Although they detect the second minima at ~ 1.6 km, corresponding to the subsidence inversion, they do not resolve the fine temporal structure of the convection.

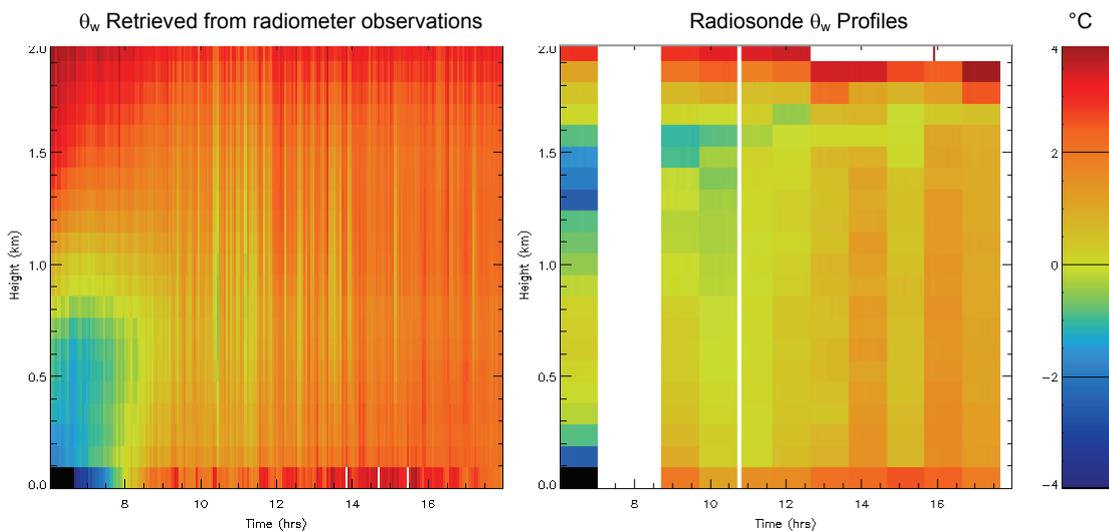


Figure 6-3 Time-height plots of wet-bulb potential temperature from profiles retrieved from radiometer observations (left) and radiosondes (right) at Cardington on 8/4/03.

More simply, the lapse rate near the surface can be estimated as the difference between the surface temperature and that retrieved at 1 km. This is evaluated over the 24 hr period of the same example in Figure 6-4. This shows the retrieved lapse rate became negative soon after sunrise as the nocturnal surface inversion was overcome by solar heating of the surface. For most of the daylight hours the boundary layer approximately followed the dry adiabatic lapse rate (-9.8 K/km). But note there are small oscillations ~ 1 K with a period of ~ 30 min, which corresponds to the time-scale of thermal overturning of the convective boundary layer. Solar heating of the surface generated shallow super-adiabatic layers, which resulted in lapse rates temporarily lower than the dry adiabatic lapse rate. The retrieved lapse rate closely followed that measured by radiosondes. The mesoscale model also followed the same dry adiabatic lapse rate (allowing for surface heating) while the boundary layer was convectively unstable. The model's resolution (12 km) does not allow the convective features to be represented explicitly, so its lapse rate is much smoother than the radiometer retrievals during this period. However, the model seriously over-estimates the surface temperatures at night, resulting in an under-estimation of the lapse rate from 00-08 UTC and again from 17-24 UTC.

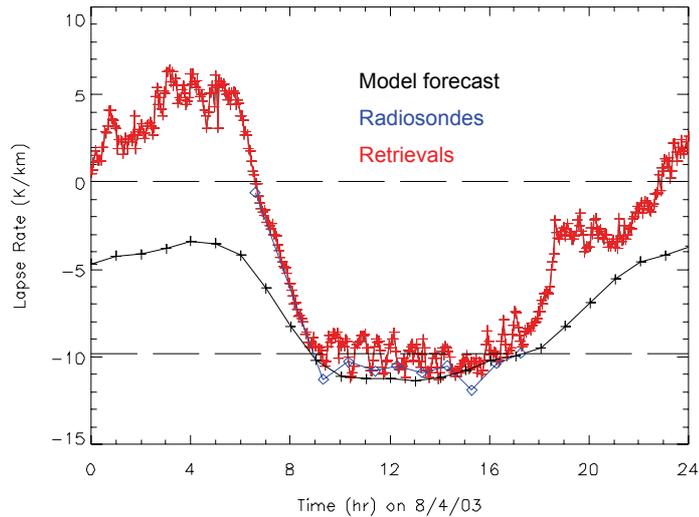


Figure 6-4 Time series of difference between the temperature at 1 km and surface forecast by mesoscale model, measured by radiosonde and retrieved from radiometer observations by neural network over 24 hr period at Cardington 8/4/03.

6.1.3 Resolving thermals

In Summer 2005, the Radiometrics TP/WVP-3000 was deployed at Linkenholt, Hampshire (UK) for the duration of the Convective Storm Initiation Project (CSIP) [Hewison and Gaffard, 2006]. A wind profiling radar was also operated at the same site configured to measure the vertical velocity and spectral width at high resolution, allowing the identification of individual thermals. In this example, one short period of data on 29/6/05 is examined to illustrate the ability to resolve the structure of strong convective thermals from the radiometer data. This convection was associated with a thunderstorm which later caused flash flooding in Oxford.

The neural network retrievals from the radiometer data are shown in Figure 6-5 for a 1 hr period on this day. The wind profiler showed a thermal plume started ascending at 1633 UTC, reaching a maximum vertical velocity of 1 m/s [Nash *et al.*, 2006]. Figure 6-5d shows the surface temperature gradually increased from 18.2 to 19.3 °C at 1628 UTC, which may have triggered convection, forming the updraught. The temperature profiles in Figure 6-5a show unstable conditions – a large temperature gradient – near the surface at 1633 UTC, before the thermal. The temperature profiles then show neutral well-mixed conditions within the thermal – a smaller temperature gradient. The relative humidity was lower within the thermal (60-75 %), due to its higher temperature. The profile remained closely adiabatic until the end of the updraught at 1639 UTC. Figure 6-5i also shows the IWV built up before the thermal, then released, although this could be associated with the presence of low clouds (Figure 6-5j). At 1631 UTC this low cloud cleared, allowing solar heating of the surface to initiate convection. The other periods of Figure 6-5 showing similar warm plumes of low RH also correspond to periods of ascent in the wind profiler data, but with weaker vertical velocities.

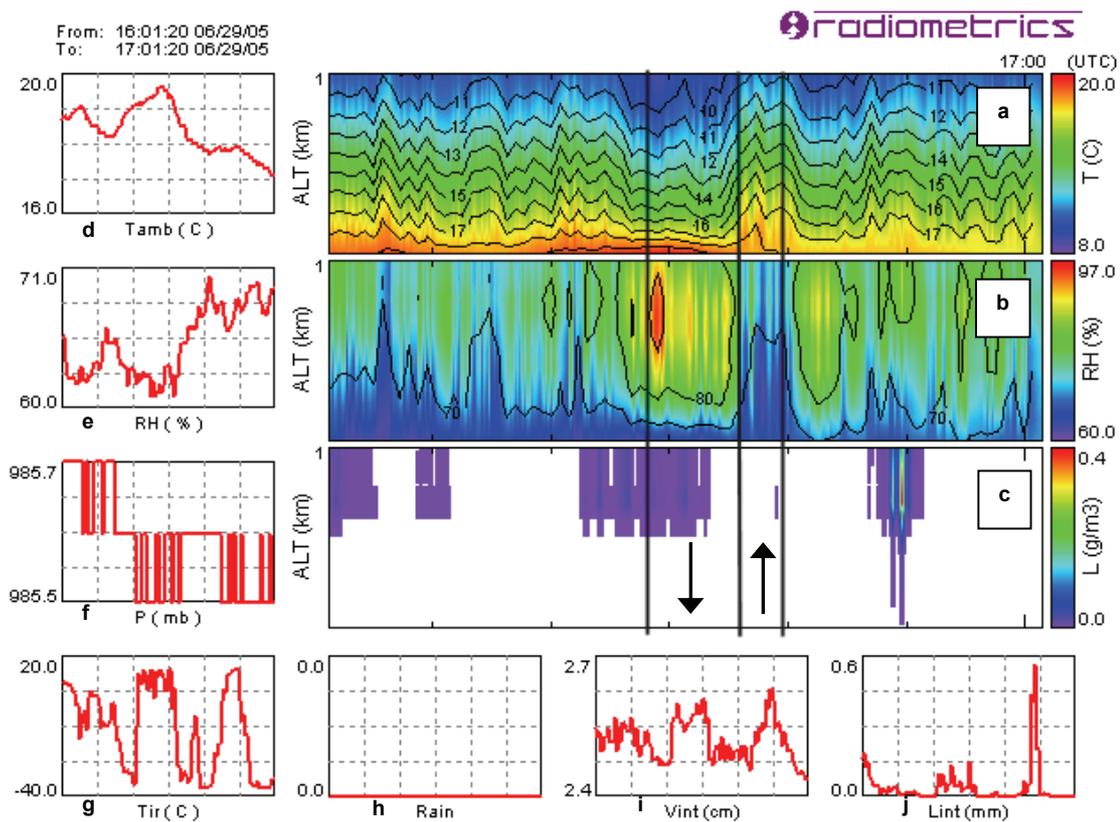


Figure 6-5 Retrievals from microwave radiometer observations on 29/6/05 16-17 UTC showing (a) temperature, (b) relative humidity and (c) liquid water time-height cross sections for 0-2 km height. Time series of scalars are also shown in red for (d) surface temperature, (e) surface relative humidity, (f) pressure, (g) infrared brightness temperature, (h) rain sensor, (i) integrated water vapour and (j) integrated liquid water. Arrows indicate updraught and downdraught.

Although radiosondes were not launched from Linkenholt on this day, one was launched at 1600 UTC from Larkhill, 38 km south-west of Linkenholt and 90 m lower altitude. The profile measured by the radiosonde, shown as a tephigram in Figure 6-6 has a temperature inversion at about 900 hPa (~0.7 km). Normand's construction shows that the surface temperature needed to exceed 19.5 °C for a parcel of air to overcome this convective inhibition and become unstable in relation to the environment, releasing convective available potential energy (CAPE) of 289 J.kg⁻¹. This would result in the development of cloud more than 7 km deep and hence thunderstorms. Figure 6-6 also shows the profiles retrieved from radiometer observations at two times, representing the extremes of the convective cycle. At 1633 UTC, before the thermal had initiated, the retrieved profile shows absolute instability in a super-adiabatic layer near the surface. By 1637 UTC, during the updraught, the profile is ~3 K warmer at 950 hPa with neutral stability over a deep well-mixed boundary layer with a weak inversion at the surface. However, these conclusions depend on the synchronisation between the wind profiler and radiometer clocks. On other days, their signals are not coincident.

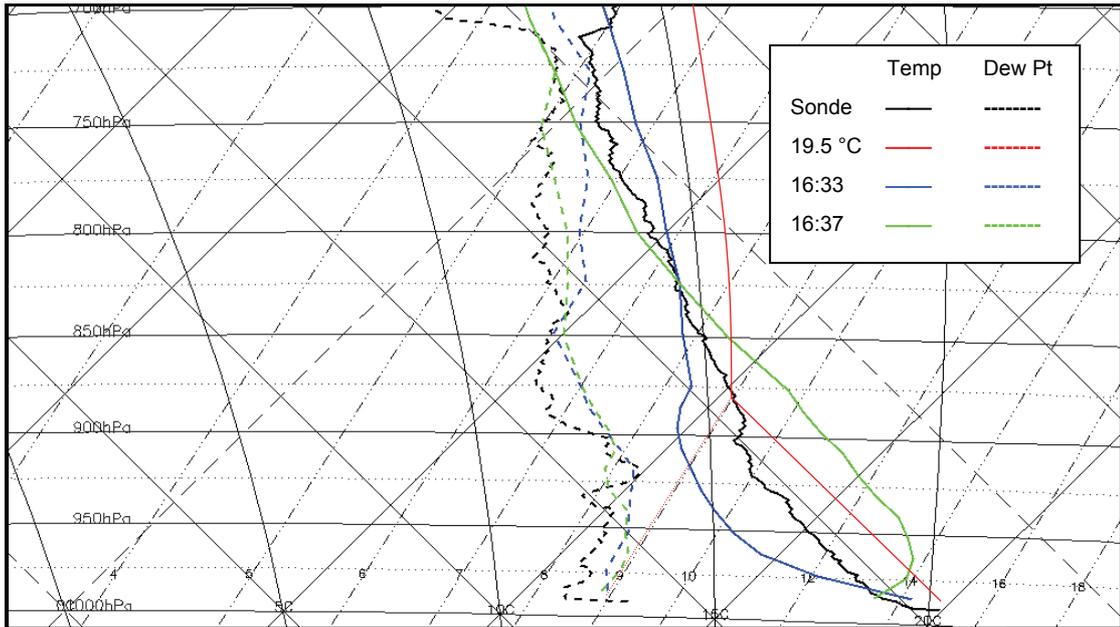


Figure 6-6 Tephigram showing profiles of temperature (solid line) and dew point (dotted lines) in lowest 3 km on 29/6/05.

Black lines show profile from radiosonde launched from Larkhill at 1600 UTC.

Blue and green lines show profiles retrieved from radiometer observations at Linkenholt, before the thermal at 1633 UTC, and within the thermal at 1637 UTC.

Red lines are Normand's construction showing the trajectory of a buoyant parcel of air from the surface at 19.5 °C / 13.5 °C temperature/dew point.

6.2 Retrievals of cloud and fog

Although this work has concentrated on the retrieval of temperature and humidity profiles, the neural network also outputs profiles of liquid water content as well as the column integrated amounts of liquid and vapour. The Integrated Liquid Water retrievals were discussed in §5.6. The microwave channels have very little sensitivity to the vertical distribution of cloud as the spectrum of emission by liquid water follows a broad continuum and only has a weak temperature dependence. Most of the skill in the neural network's retrieval of liquid water profile comes from the infrared radiometer measurement of the cloud base temperature and is inferred from the humidity profile, as liquid water cloud cannot persist in very dry air. It is difficult to validate liquid water retrievals because of the absence of suitable sensors on radiosondes, but it is useful to examine the performance of the retrievals in a foggy case.

Figure 6-7 shows time-height plots of the profiles retrieved from zenith radiometer observations at Payerne over 24 hr on 5/12/03. In the morning, the retrievals show a surface inversion - the temperature increases with height to 1 km – trapping a humid layer near the surface, forming fog in the lowest 200m. Solar heating warmed the surface during the

morning, and subsequently the lowest 1 km, causing the inversion to weaken and the relative humidity to reduce. This allowed the fog to lift briefly from 1300-1600 UTC, before it reformed after sunset when the surface cooled again. Although the cloud retrievals do not resolve a clear fog top, they suggest it was deeper in the evening when the inversion was weaker than in the morning. This was confirmed by a co-located cloud radar to be 200-300 m during the morning and ~500 m in the evening. The cloud retrievals can be improved by applying a minimum threshold to the relative humidity where clouds are placed in the retrievals, but this is unlikely to have a strong impact due to the retrieved humidity profiles' poor vertical resolution.

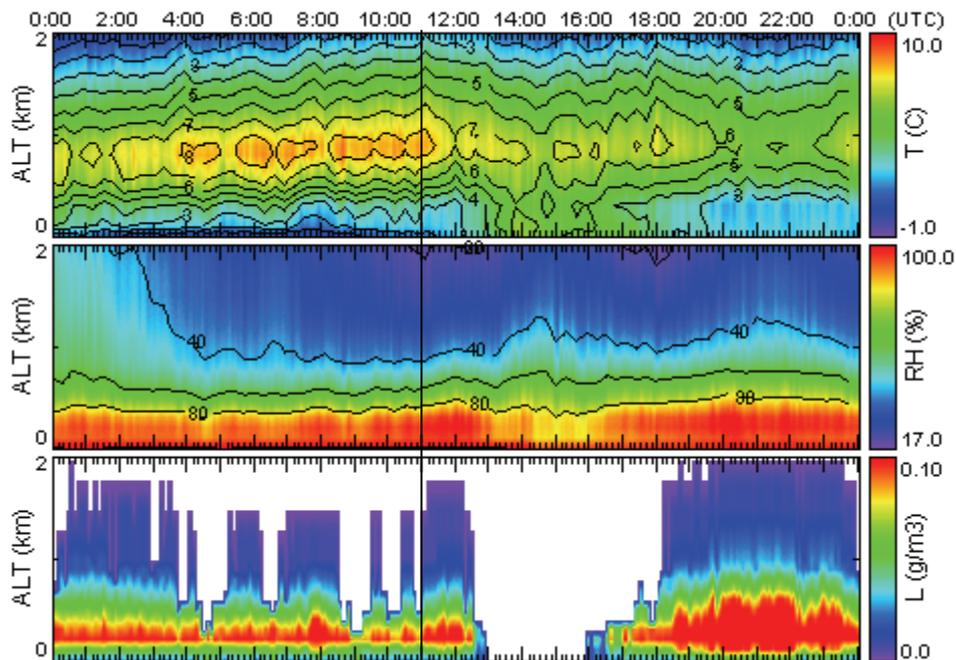


Figure 6-7 Time-height plots of profiles retrieved from radiometer observations at Payerne for 0-2 km over 24 hours on 5/12/03. Upper panel shows temperature, middle panel shows relative humidity, lower panel shows liquid water content. Vertical line shows time of radiosonde profile.

The profiles retrieved by the neural networks using zenith and elevation scanning observations from the Radiometrics TP/WVP-3000 are compared with the co-located radiosonde in Figure 6-8. The radiosonde shows a saturated, isothermal layer 300 m deep at the surface and a strong inversion above this, with further stratified layers extending to 770 hPa (~2.2 km), above which it is approximately adiabatic, but very dry. Both retrievals identified the temperature inversion and general humidity structure, though were unable to resolve the structure of the layers above. The retrieval using the elevation scan is slightly closer to the radiosonde's temperature near the surface, though still has temperature errors of up to 3 K at some levels and is unable to accurately represent the shallow isothermal layer. In

this example the atmosphere was sufficiently homogeneous to pass the quality control test described in §5.4.1 and condensation had not formed on the radiometer, so there is no reason to reject the retrieval on these grounds. The poor performance of the elevation scan retrievals can only be explained by their inherent poor resolution, combined with the problems of sensitivity to surface emissions described in §3.2.1 for the humidity profile.

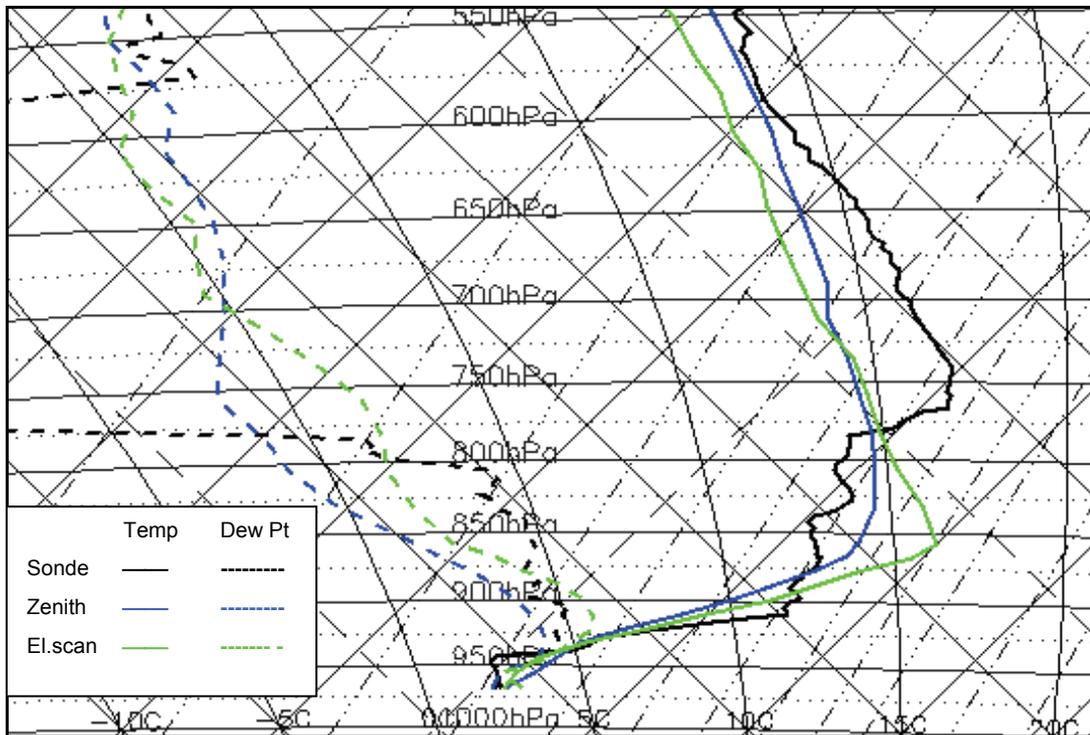


Figure 6-8 Tephigram showing profiles of temperature (solid line) and dew point (dashed lines) in lowest 5 km at 1100 UTC on 5/12/03 at Payerne, Switzerland.

Black lines show profile from radiosonde.

Blue and green lines show profiles retrieved by neural network using zenith only and elevation scanning radiometer observations, respectively.

6.3 Ability to retrieve temperature inversions

The statistical assessments of the retrievals in Chapter 5 do not explicitly test their ability to retrieve inversions, which are of great interest when forecasting boundary layer phenomena such as the evolution of cloud or fog. It seems from inspection of Figure 6-2 and other cases that the neural network is incapable of fully resolving elevated inversions with a base above 1 km. The retrievals' performance can be assessed by comparing the properties of the temperature inversions in the retrieved profiles with those calculated from the significant levels from coincident radiosondes.

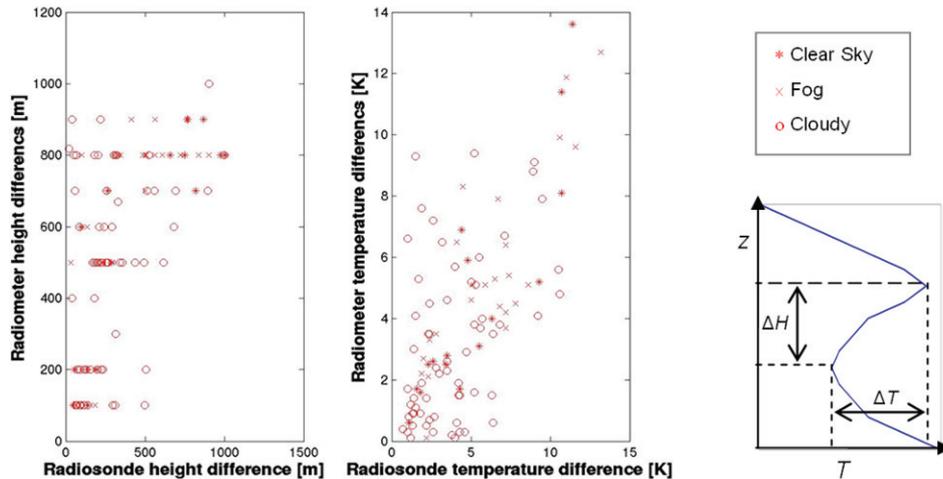


Figure 6-9 Differences between height and temperature at the top and base of inversion layers retrieved from the radiometer and measured by coincident radiosondes at Payerne 5/11/03-19/1/04. Includes all inversions – surface based and elevated.

The differences between the height and temperature at the top and base of inversions retrieved from the radiometer and measured by coincident radiosondes during the TUC experiment are compared in Figure 6-9. This included all inversions with magnitude of at least 1 K, with a base below 2.5 km. The r.m.s. differences for the subset of comparisons in clear skies, foggy and cloudy conditions are given in Table 6-1. This shows the ability to retrieve inversions is reduced by the presence of fog or cloud due to the variability they introduced into the scene radiances. It is also evident from Table 6-1 that the retrieval of surface-based inversions is less challenging, as the surface sensor provides a strong constraint of the temperature at its base.

Table 6-1 Statistics of neural network retrievals of inversion height, ΔH , and temperature differences, ΔT

Inversions $\Delta T > 1$ K, $z_{base} < 2.5$ km	r.m.s. ΔH (m)	r.m.s. ΔT (K)	r.m.s. ΔH (m)	r.m.s. ΔT (K)
	All Inversions		Surface inversions	
Clear Sky *	204	1.9	123	1.1
Foggy x	329	1.9	208	1.2
Cloudy o	340	2.8	375	2.8

6.4 Average layer temperatures

The geopotential thickness of atmospheric layers between levels of fixed pressure (*thickness*), which is commonly used by forecasters for a number of applications, is directly proportional to the average temperature of that layer. The microwave radiometer's channels in the oxygen band are sensitive to emission over relatively thick atmospheric layers and can provide accurate retrievals of the layer average temperatures. For example the 1000-850 hPa thickness has been used by forecasters in the UK as an indicator of the probability of any precipitation falling as snow and for maximum temperature forecasting. This is a linear function of the average temperature of the lowest ~1.5 km of the atmosphere, which in turn is linearly related to the brightness temperature in the 54-60 GHz channels.

It is also possible to calculate the thickness between other pressure levels. For example, the 1000-500 hPa thickness is commonly used to identify warm and cold advection on synoptic scales. A time series of this thickness over the 2 months of the radiometer's deployment at Cardington in Spring 2003 is shown in Figure 6-10. This shows diurnal cycles on some days as well as larger changes on longer timescales related to frontal passages over the radiometer. The 1000-500 hPa thickness has been calculated here from the average temperature retrieved by the neural network over the lowest 5.5 km of the atmosphere. It is evident that the resolution of the thickness calculated in this way is better than 1 dm, corresponding to a <0.5 K in average temperature. This approaches the accuracy achieved by radiosondes (~0.2 K r.m.s.) [John Nash, personal communication, 2006] and is comparable to short-range forecasts by NWP models as seen by the background error covariance matrices calculated by NMC method in §4.3.2.2. It would also be possible to retrieve thickness directly from the observations by re-training the neural network, but this may not be advantageous as the optimal temperature profile retrieval should also yield the optimal retrieval of layer average temperatures [Thiebaut, 1973].

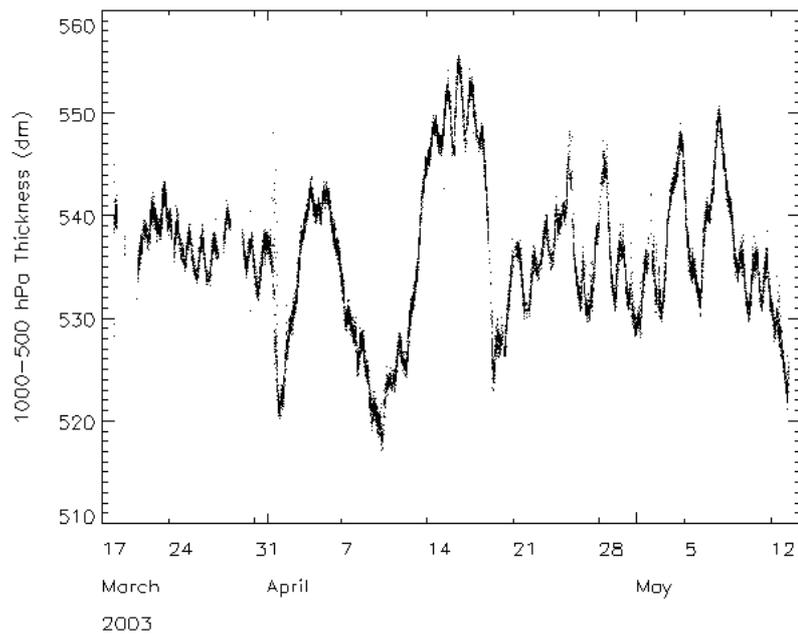


Figure 6-10 Time series of 1000-500 hPa thickness calculated from profiles retrieved by neural network from radiometer observations at Cardington in Spring 2003.

6.5 Performance in rain

Retrievals from radiometer data are expected to be degraded in rain because of the scattering of microwave radiation by raindrops as well as emission from liquid water that accumulates on the instrument's window (as discussed in §5.5.2). It is useful to compare the retrieved profiles with co-located radiosondes to assess the impact of rain on the retrievals. This is shown in Figure 6-11 for a series of radiosondes, showing an elevated temperature inversion, which lowered during the morning. The humidity increased at higher levels and the cloud thickened (LWP increased) as a cold front approached Camborne from the northwest. Although the radiometer's rain flag was set throughout this period, the retrievals closely matched the radiosonde profiles during the morning, until the Integrated Liquid Water (shown as LWP) exceeded 1 mm (1000 g/m^2) after 1515 UTC. Initially the rain biased the retrieved temperature profile above 1 km (although the bias near 2 km at 1319 UTC is likely to be partially due to the retrievals poor vertical resolution at this level). The information for this part of the profile comes from the 51-54 GHz channel, where the oxygen absorption is weaker and emission from liquid water is more significant than at higher frequencies. After analysing other examples, this was adopted as an empirical threshold for the maximum value of ILW, above which retrievals were rejected as being likely to be rain-affected due to the radiometer's calibration being biased.

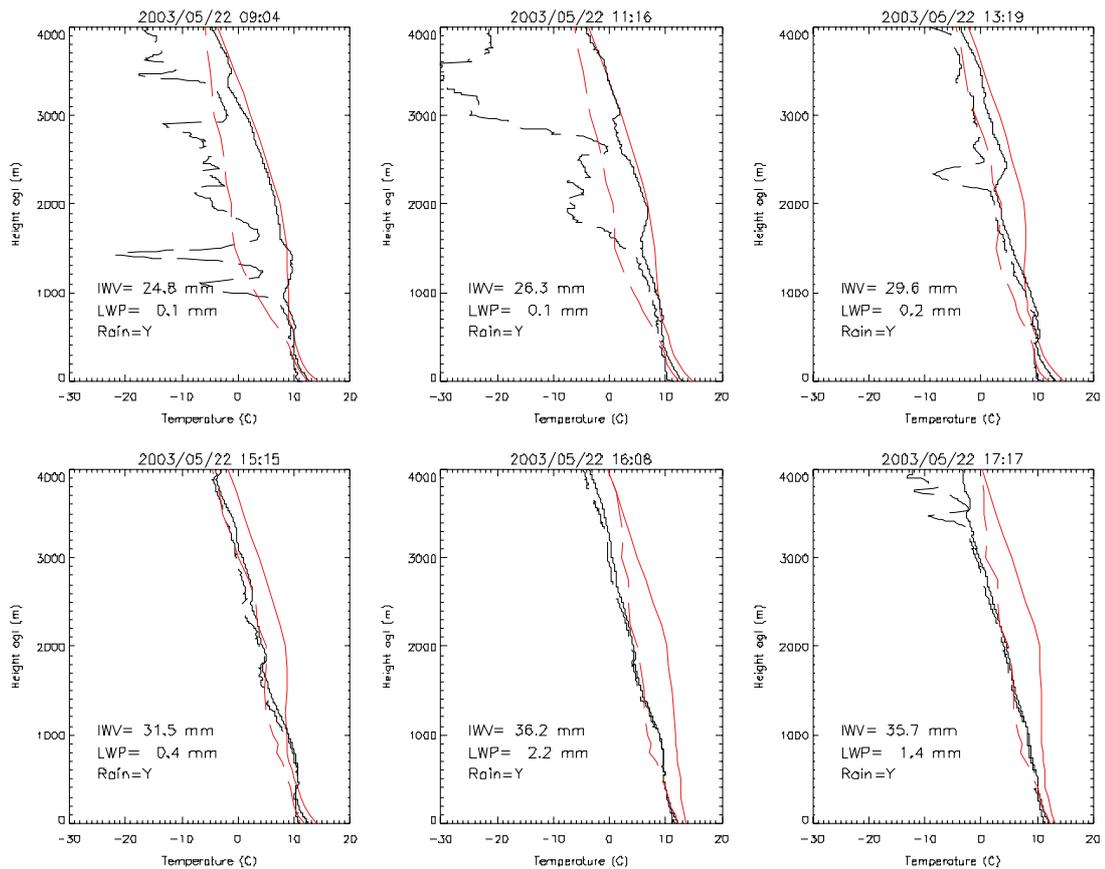


Figure 6-11 Profiles of temperature (solid lines) and dew point (dashed lines) measured by 6 radiosondes (black) and retrieved from radiometer observations (red) at Camborne on 22/5/03. Integrated water vapour (IWV) and liquid water (LWP) retrieved by the neural network are indicated in each panel.

n.b. The data in this example was taken before the radiometer's blower was upgraded and the hydrophobic coating applied to its radome. It is expected that these hardware modifications will reduce the impact of precipitation on the observations (§2.4.3)

6.6 Comparison with model forecast profiles

Retrievals from radiometer observations can also be compared with NWP model analyses and forecasts. Although it can be useful to validate a single forecast against observations, comparing a large ensemble of observations and forecasts can provide insight into the systematic shortcomings of each data set. The temperature and humidity profiles retrieved from high time-resolution (~12 s) radiometer observations are compared with hourly forecasts from an NWP model. Although the retrievals have poor vertical resolution, their high time-resolution helps reveal fine structure in the data and identify trends.

6.6.1 Clear Sky CSIP case – 12 July 2005

Both retrievals and model in Figure 6-12 show the development of a surface inversion during the early hours, trapping a shallow humid layer near the surface, which generated low stratus by 0400 UTC, then fog by 0500 UTC. However this inversion is stronger in the retrievals than the model, probably as a result of the initial humid layer at 2 km dissipating more rapidly than forecast, giving faster cloud clearance, which allowed greater radiative cooling of the surface. In this case, if radiometer observations at 0000-0200 UTC had been assimilated into the model, the fog may have been forecast more accurately because the cloud would be better represented.

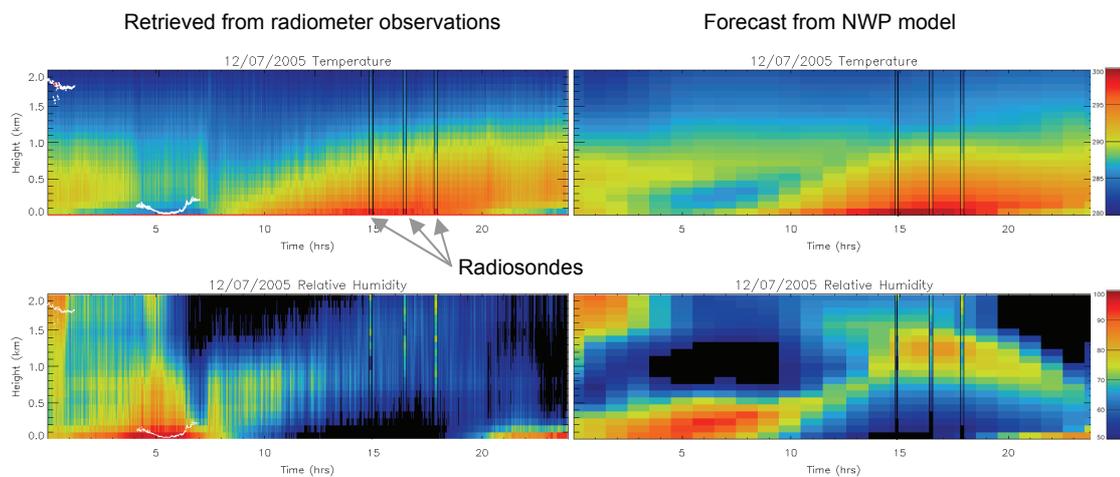


Figure 6-12 Time-height plots of temperature [K](top) and relative humidity [%](bottom) retrieved from radiometer observations (left) and hourly T+0 to T+24 hr forecasts (right) from mesoscale version of the Met Office Unified Model for Linkenholt on 12/7/05. 3 radiosonde profiles are shown for comparison, surrounded with black vertical bars. White dots show cloud base height measured by laser ceilometer (mostly clear).

The development of the boundary layer, starting with solar heating of the surface and a warm layer spreading upward was forecast and retrieved well with temperatures closely matching radiosonde observations. However, the retrievals failed to resolve the forecast elevated temperature inversion associated with the residual of the previous day's boundary layer. The relative humidity near the surface decreased as the temperature increased, while it increased at the top of the boundary layer both in the retrievals and model. However, the retrievals spread the humidity over a deeper layer, diluting it more than was observed by the radiosondes, while the model tended to over-estimate the humidity at the top of the boundary layer compared to the radiosondes.

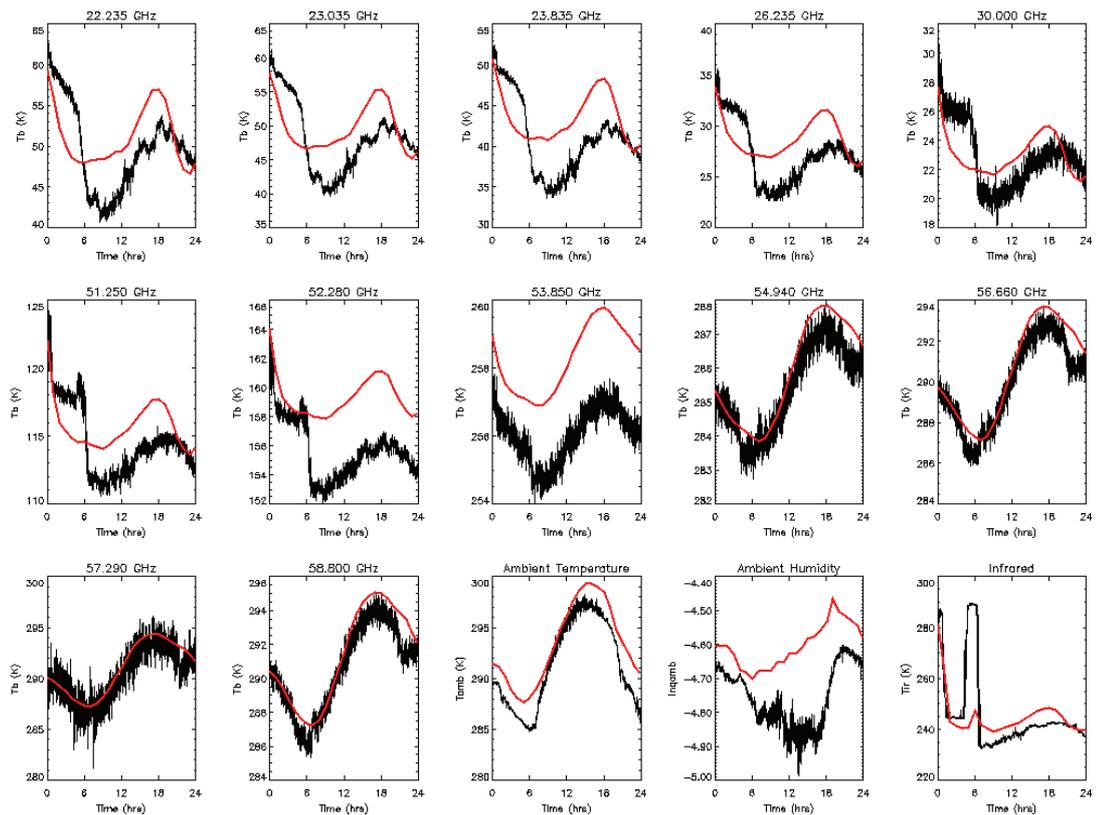


Figure 6-13 Time series of observations from Radiometrics TP/WVP-3000 (black) and forward modelled mesoscale model output (red) from 00-24 UTC on 12/7/05 at Linkenholt.

The radiometer observations are compared to the hourly forecasts from the 00 UTC run of the mesoscale model in Figure 6-13 in observation-space. This is an example of the insight available from an *O-B* analysis (§4.6.3). Observations in the low frequency channels initially show emission from cloud, which thins within the first hour. This cloud thins more slowly in the model forecasts, causing the ambient temperature at the surface to not fall as rapidly as observed. The 51-52 GHz and infrared channels show enhanced emission due to the formation of fog about 06 UTC, whereas the model only forms thin low-level cloud at this time, giving a much lower signal. Solar heating of the surface and a reduction of the boundary layer humidity (seen as a reduction in the T_b of the low frequency channels) after 06 UTC caused the fog layer to dissipate. Later in the day, the model tended to increase the boundary layer humidity too rapidly and overestimated its maximum, compared to the low frequency channels' observations. However, the temperature cycle seen in the high frequency channels was well represented by the model. n.b. The forward model shows a systematic bias compared to the observations in the 51-54 GHz channels, which may be due to a bias in the absorption model (§3.4.2).

6.6.2 Cumulus capped boundary layer CSIP case – 13 July 2005

A similar comparison is shown for 12 hours of the following day in Figure 6-14. In contrast to the previous example, the weather on this day was partially cloudy, with more unstable conditions and cumulus congestus developing in the area as a result of convergence between a weak northerly flow and a sea breeze from the south during the afternoon. However, these clouds were not thick enough to produce rain at Linkenholt. The convective boundary layer developed rapidly during the morning and cloud bases increased from 1.0 to 1.5 km later.

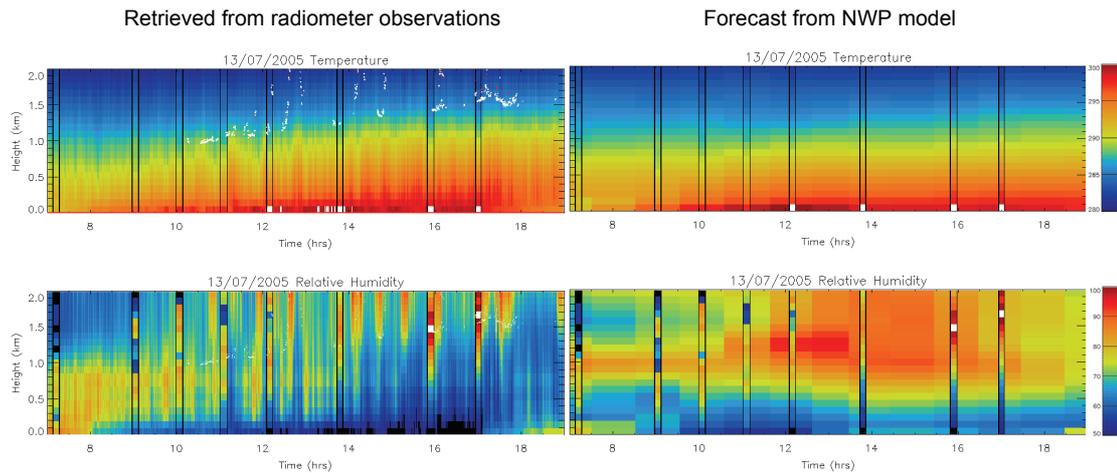


Figure 6-14 Time-height plots of temperature [K](top) and relative humidity [%](bottom) retrieved from radiometer observations (left) and hourly T+6 to T+18 hr forecasts (right) from mesoscale version of the Met Office Unified Model for Linkenholt on 13/7/05. 8 radiosonde profiles are shown for comparison, surrounded with black vertical bars. White dots show cloud base height measured by laser ceilometer (cumulus).

The profiles retrieved from the radiometer observations show oscillations in temperature and relative humidity as a result of convective thermals and colder downdraughts. It remains possible that cloud could be partially misinterpreted by the retrievals as changes in humidity and this could explain some of the strong oscillations seen in the RH profiles. Nevertheless, they are broadly consistent with the radiosonde observations, although exact comparisons are made difficult by the atmospheric variability in these conditions. This highlights the problems of representativeness errors even between two systems, although they are both effectively spot measurements, making it difficult to assimilate observations in an NWP model. In this case the hourly forecasts from the NWP model do not reflect the variability, either in temperature or humidity, as its 12 km grid is not able to represent convective scales and only provides area averages over the model grid-boxes. Hence it will be impossible to fully exploit the observations' full temporal resolution until it is possible to assimilate them into convective-scale NWP models.

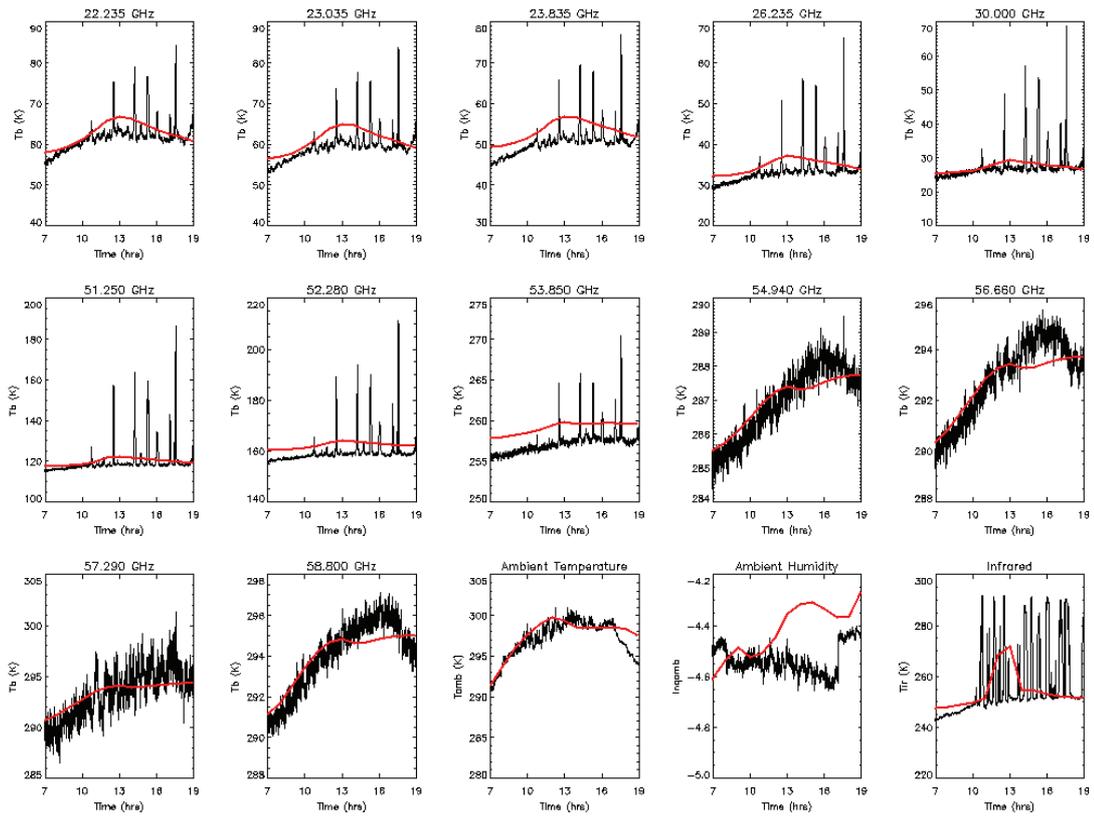


Figure 6-15 Time series of observations from Radiometrics TP/WVP-3000 (black) and forward modelled mesoscale model output (red) from 07-19 UTC on 13/7/05 at Linkenholt.

Figure 6-15 shows the time series for this case of observations in the period 07-19 UTC compared to forecasts from the 00 UTC run of the mesoscale model. The observations in the low frequency channels show a weak maximum around 13 UTC due to the increase in boundary layer humidity, but with large spikes, corresponding to emission from liquid water cloud passing over the radiometer. Although the general increase is represented by the mesoscale model, its magnitude tends to be overestimated, probably because it is incapable of resolving the individual clouds that generate the strongest signal in the observations. The infrared radiometer clearly shows emission from the individual clouds at the temperature of their base, while the model only generates thin cloud from 11-14 UTC, which has a lower maximum brightness temperature of 273 K. The temperature cycle of the boundary layer is initially well represented by the model, although there is a discrepancy with the observations in the high frequency channels between 13-18 UTC, corresponding to the period when the model did not generate any cloud. This resulted in the model overestimating the surface temperature after 17 UTC.

6.7 Summary of Case Studies

The microwave radiometer is sensitive to emission over relatively thick layers of the atmosphere. This makes its data ideal for retrieving geopotential thickness, a measure of the layer average temperature, and identifying trends in this. However, it can also be used to retrieve profiles of temperature and humidity, albeit with limited vertical resolution. Its ability to retrieve temperature inversions near the surface has been demonstrated, although the retrievals were unable to fully resolve elevated inversions with a base above ~1 km. The poor vertical resolution can generate errors of several kelvin when comparing retrieved profiles with radiosondes which show strong inversions. This is a fundamental limitation of the radiometer's observations and can only be improved by the addition of other data in the retrieval – for example, an NWP background or observations from other instruments. Retrievals using elevation scanning radiometer observations improved the temperature profile near the surface but degraded it at higher levels, as well as the humidity profile compared to zenith viewing observations.

The static stability has been analysed from the retrieved profiles and found to be capable of representing the basic features of the convective cycle of the boundary layer. It is also possible to derive stability indices from the retrieved profiles, either as a simple lapse rate for dry processes or more sophisticated potential temperature gradients for moist processes. These have been shown to be able to identify periods of static instability in time series of observations.

Although the microwave radiometer observations provide very little information on the vertical distribution of liquid water, some information on this is provided by the infrared radiometer and can be inferred from the retrieved humidity profile. Together with the retrieval of Integrated Liquid Water and temperature inversions capping fog layers, these can be related to the evolution of the fog.

As well as being validated against radiosondes, time series of profiles retrieved from radiometer observations have been compared with forecasts from an NWP model. This can be used to investigate shortcomings of both the retrievals and forecasts. It also highlights the difficulty faced when trying to assimilate high time-resolution observations in highly variable conditions because of the problems of using spot measurements to represent model grid-box values. This is fundamentally limited by the resolution of the NWP model, but further confounded by the inability to exploit the temporal variability in the 3D variational assimilation (3D-VAR) method currently used in the mesoscale NWP model. Ultimately, with the development of 4D-VAR in convective-scale NWP models, the observations' high time-resolution could address the deficiency of the retrieval's poor vertical resolution and the necessary sparsity of any network of radiometers.

Chapter 7 Optimising Observations for Operations

This chapter extends the analysis of the previous chapters to draw together some recommendations for how the radiometer observations could be optimised for operational use. It starts by examining the information content of the observations, allowing a trade-off between different observing strategies. From this recommendations are made for the optimum observing sequence of the Radiometrics TP/WVP-3000 in its current configuration and for the specification of future radiometers suitable for operational use. The integration of observations from other instruments is also considered and the potential applications of their data reviewed. Finally, potential methods to exploit the high time-resolution of the observations are reviewed.

7.1 Information Content – Degrees of Freedom for Signal

The information content of the observations with respect to the background can be quantified in the Bayesian framework by introducing the concept of *Degrees of Freedom for Signal*, *DFS*. This allows a trade-off between different channel combinations and observing strategies. *DFS* represents the number of layers in the profile which are retrieved independently [Rabier *et al.*, 2002]. It can be calculated as:

$$DFS = \text{Tr}(\mathbf{I} - \mathbf{A}\mathbf{B}^{-1}) \quad (7.1)$$

where Tr is the trace operator, extracting the diagonal components of the matrix, \mathbf{I} is the identity matrix, a matrix with ones on the main diagonal and zeros elsewhere, \mathbf{A} and \mathbf{B} are the error covariances of the analysis and background respectively, as defined in Chapter 4.

DFS is somewhat akin to the signal to noise ratio of a scalar measurement, where the term $(\mathbf{I} - \mathbf{A}\mathbf{B}^{-1})$ is the *Model Resolution Matrix* [Rabier *et al.*, 2002], which describes how much the measurement and retrieval processes spreads the profile information relative to a perfect retrieval of a reference profile with perturbations applied to each level in turn.

7.1.1 Instantaneous Observations at Zenith

Initially, \mathbf{A} was evaluated from instantaneous zenith observations, as described in §4.6.1. *DFS* was then calculated using this \mathbf{A} and \mathbf{B} defined over the full state vector and for subset of \mathbf{A} and \mathbf{B} only covering the temperature profile, DFS_T . The *DFS* for the humidity profile can then be calculated as $DFS_q = DFS - DFS_T$. This calculation has been repeated for the radiometers and surface sensors (\mathbf{A}), surface sensors only (\mathbf{A}_s) and radiosondes (\mathbf{A}_{rs}) with \mathbf{B}_{ATOVS} and \mathbf{B}_{NMC} in clear and cloudy reference profiles, as shown in Table 7-1.

Table 7-1 Degrees of Freedom (DFS) for temperature and humidity profiles using instantaneous observations from Radiometrics TP/WVP-3000

Back-ground Error	Profile	Radiometers, A		Surface Sensors, A_s		Radiosondes, A_{rs}	
		Temp. <i>DFS_T</i>	Humidity <i>DFS_q</i>	Temp. <i>DFS_T</i>	Humidity <i>DFS_q</i>	Temp. <i>DFS_T</i>	Humidity <i>DFS_q</i>
B_{ATOVS}	Clear	2.8	1.8	0.9	1.0	8.6	7.1
	Cloudy	2.9	3.0	1.0	1.0	8.6	7.1
B_{NMC}	Clear	3.0	1.8	1.0	1.0	11.7	11.4
	Cloudy	3.1	3.2	1.0	1.0	11.7	11.4

For **B_{ATOVS}** in the clear US standard atmosphere, the microwave radiometer (including surface sensors) can provide nearly 3 degrees of freedom for temperature, but only ~2 for humidity. The surface sensors alone provide approx 1 *DFS* for both temperature and humidity, but these are not linearly independent of the radiometer's information. The Integrated Water Vapour (IWV) can be retrieved accurately from 2 channel microwave radiometers with $DFS_q \approx 0.7/1.0$ in clear/cloudy conditions, respectively. So the 2 *DFS* for humidity available from the radiometer observations represents little additional information on the vertical distribution of the humidity profile in clear cases over what would be available from this or a GPS combined with a surface sensor. It should also be noted that the correlations in **B** limit the *DFS* from radiosondes to *only* 8.6 for temperature and ~7 for humidity, despite their high vertical resolution. However, this is still much more information than is available from a single 1D-VAR retrieval from radiometer data.

The *DFS* calculation was repeated for a different reference profile – one taken from a mesoscale model profile from Camborne with low cloud. This showed little difference in temperature information, but increased DFS_q by ~1.0, which highlights the benefit of the use of microwave radiometers in cloudy conditions. Although this is a typical case, DFS_q was found to vary by a factor of ~2 depending on the amount of cloud, through its sensitivity in the Jacobian, **H**. The *DFS* calculation was also repeated including T_{ir} in the observations. As expected, this made little difference in clear conditions, but was found to have a small impact on the humidity profile in cloudy cases.

The sensitivity of *DFS* to the assumed background errors was investigated by repeating the calculations in Table 7-1 using **B_{NMC}** as well as **B_{ATOVS}**. This was found to produce slightly more *DFS* for temperature and humidity.

7.1.2 Averaging Observations to Reduce Radiometric Noise

The observation errors of the highest frequency channels were found to be dominated by radiometric noise (§4.4.1) – if this could be reduced better temperature profiles could be retrieved in the lowest 1 km. This may be achieved by applying a rolling average of their

observations over ~5 minutes, as they are less sensitive to short-term atmospheric variability, which dominates the observation error in other channels. Examination of the time structure functions in Figure 4-6 showed the 55-59 GHz channels are noise-limited for periods up to ~1000 s, so if they are averaged over this period, their variance is expected to reduce linearly with the number of observations averaged. The radiometric noise of the four highest frequency channels was reduced by a factor of 11 and the information content analysis repeated. This increased DFS by 0.8 (a 28% improvement) for the temperature profile with B_{NMC} and 0.5 with B_{ATOVs} , showing a large benefit in both clear and cloudy cases, but no benefit to the humidity profile. This was accompanied by a big improvement (~50%) in the vertical resolution of the temperature retrievals at all levels above 300 m. The drawback of this technique is that applying a symmetric rolling average to real-time data obviously introduces a delay before the data is available for retrievals and reduces the ability to resolve atmospheric structures on time-scales shorter than the averaging period.

Lower frequency channels are dominated by representativeness errors due to atmospheric variability. The experiment described in §5.6.3 showed that observations in cloudy conditions can safely be averaged without biasing the retrievals which use them. However, the time structure functions in Figure 4-6 show the variance of these channels would not be much reduced by averaging in the same way. So it is recommended that they should be used as instantaneous observations in the retrievals, and the retrieved profiles co-averaged to represent the average conditions over the period. This technique can also yield useful information on the variability of cloud not available by other methods (§5.6).

To summarise: although current instruments can be made fast and stable with sufficiently low noise to produce beneficial retrievals, reductions in the random noise of the channels near the centre of the oxygen band would produce further improvements.

7.1.3 Extra information from elevation scanning

Viewing the atmosphere at different elevation angles can bring additional information to the retrievals, due to the increased path length, as described in §5.4. The emission seen by a channel viewing the atmosphere at different elevation angles can be approximated by multiplying the absorption coefficient at its centre frequency by the secant of the zenith angle, $\sec \theta$, (known as *airmass*), to account for the increased atmospheric path length. For example, an extra view at $\theta=60^\circ$ is equivalent to an additional channel with the absorption coefficient double that of its zenith viewing counterpart. Figure 7-1 shows the absorption at 60 GHz is equivalent to that at 56 GHz through a doubled atmospheric path length. Therefore, observation at low elevation angles with any additional channels at or below 56 GHz will be redundant if they are also observed at an angle with half the value of $\sec \theta$. Given the default channel set of the Radiometrics TP/WVP-3000, it is recommended that elevation angles 14.5° , 30° , 90° , 150° and 165.5° would provide the optimum set, but that only the three

highest frequency channels in the oxygen band need to be observed away from zenith. The elevation scan angles need to be symmetric about zenith to check the assumption of horizontal homogeneity, which also cancels out small misalignments of the antennae (§5.4).

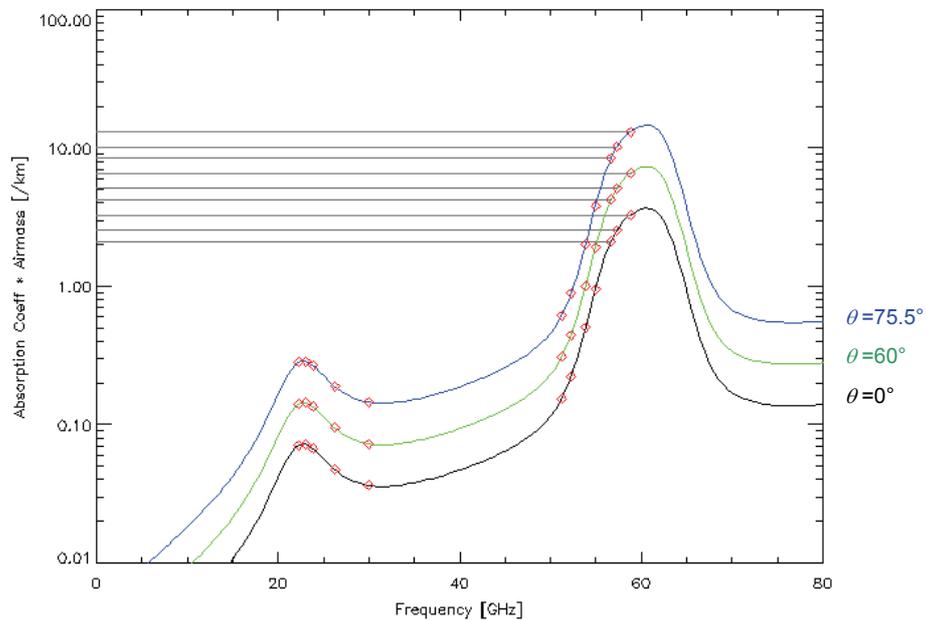


Figure 7-1 Spectrum of absorption coefficient times the secant of the zenith angle, θ , for typical conditions (1013 hPa, 288 K, 75 %RH) viewed at $\theta=0^\circ$, 60° , 75.5° ($\sec \theta=1$, 2 and 4). Red points show centre frequencies of Radiometrics TP/WVP-3000 channels. Grey lines show absorption coefficients at highest 3 frequencies at 3 zenith angles.

Viewing low elevation angles also increases the sensitivity of the 22-30 GHz channels used to profile humidity. However, their relative sensitivities remain unchanged, as shown in Figure 5-5, because the atmosphere remains optically thin ($\tau \cdot \sec \theta \ll 1$). So no extra information is available by sampling these channels at multiple elevation angles, other than by virtue of their lower noise. There are other practical considerations which limit the application of elevation scans with these channels (§5.4).

Retrievals exploiting observations at different elevation angles were implemented in the neural network and found to improve the accuracy in the lowest 1 km of the temperature profiles, but degraded the accuracy of the humidity profiles (§5.4).

The potential benefit of adding more elevation angles has been assessed by calculating the *Degrees of Freedom for Signal (DFS)* for elevation angles corresponding to $\sec \theta=1$, 2, 3 and 4. This showed an improvement over retrievals using only zenith views of ~ 1.2 in *DFS* for temperature using $\mathbf{B}_{\text{ATOVs}}$ in clear or cloudy conditions. The humidity components of *DFS* were improved by 0.8 and 2.0 in clear and cloudy conditions, respectively. But this only

reflects a reduction in the observation noise due to extra observations. The vertical resolution of the temperature profile is improved dramatically in the lowest 300 m, and significantly above 3 km. With \mathbf{B}_{NMC} the DFS improvements are ~50% larger than with $\mathbf{B}_{\text{ATOVs}}$. These represent vast improvements in the information content and vertical resolution of the profiles retrieved from radiometer data by including additional elevation angles. However, it should be remembered that this comes at the price of a longer observing cycle, and that there are practical problems associated with the use of elevation scanning, particularly with the low frequency channels. The trade-off with observing cycle is reviewed in §7.1.5.

7.1.4 Channel selection

The implications of adding or removing channels from a given configuration of a microwave radiometer can also be analysed using the concept of *Degrees of Freedom for Signal (DFS)*.

7.1.4.1 Benefit of 22-30 GHz channels

Given the low number of DFS for humidity raises the question of the value the 22-30 GHz channels to the retrievals. These channels provide the information on the humidity profile and the integrated liquid water content. However, the 51-52 GHz channels are also sensitive to humidity and cloud, so how much extra information are the low frequency channels adding? To investigate this, they were replaced with a single channel at 23.8 GHz. This channel is least sensitive to the vertical distribution of water vapour [Mätzler *et al.*, 2006], so its response is most linear with IWV. When using $\mathbf{B}_{\text{ATOVs}}$, this was found to reduce the DFS for humidity by only 0.1 for both clear and cloudy profiles with and without an infrared radiometer. However, when calculated using \mathbf{B}_{NMC} , DFS_q reduces more - by 0.2 in clear and 0.3 in cloudy profiles. This represents a relatively small change (~5-10%) in the information available on humidity profile, because of the large correlation between the low frequency channels as a result the representativeness and modelling errors. It also reflects the fact that there is a considerable degree of correlation between the IWV and the specific humidity measured at the surface. (This was calculated over 1 year's observations at Camborne as $r^2=0.54$.) So it may be possible to provide the information on the IWV from a co-located GPS sensor at a much lower cost than the extra cost of the low frequency band of the radiometer. However, again the conclusions depend critically on the assumed \mathbf{B} matrix and on the atmospheric situation: more humidity information is available in more humid atmospheres, where the absorption is stronger. This simplified analysis would need to be extended over a full range of conditions in practice before drawing firm conclusions regarding the configuration of future radiometers.

7.1.4.2 Influence of excess noise on 57.29 GHz channel

Similarly, the effect of the excess noise experienced by the 57.29 GHz channel was investigated by replacing its elements in \mathbf{R} with those of the adjacent 58.8 GHz channel (0.03 K^2) and zero off-diagonals. This improved DFS_T slightly in both clear and cloudy cases

(by 0.03-0.04 for $\mathbf{B}_{\text{ATOVs}}$ and 0.02 for \mathbf{B}_{NMC}), but otherwise made little difference. This highlights the redundancy in these channels because of their overlapping weighting functions.

7.1.4.3 Possible extra channels

The same method can be applied to quantify the benefit of adding extra channels. However, it is difficult to estimate the observation error covariance without observations from an existing system. Although the radiometric performance of proposed channels can be predicted with some accuracy and safely assumed to be diagonal, and the forward model errors can also be calculated using the methods described in §4.4.1.2, it is not possible to estimate the representativeness errors by the method described in §4.4.1.3 without a time-series of observations. As this term dominates the error budget of channels sensitive to humidity and cloud, this limits the confidence we can place in the error analysis of the performance of hypothetical channels to those that can be reliably interpolated from existing channels. One approach is to assume the representativeness error covariance of these channels is similar to that of existing channels, scaled according to the sensitivity of their Jacobians. Alternatively it may be possible to estimate it by applying a forward model to a high-resolution cloud resolving model [Löhnert *et al.*, 2007].

The architecture of the Radiometrics TP/WVP-3000 makes it possible to measure emission in any number of channels within the bands of 22-30 GHz and 51-59 GHz sequentially. This feature is not currently fully exploited in the supplied configuration, although it is simple to implement by programming the frequency synthesisers used to select the centre frequency of the channels. In this way it is possible to sample 21 channels with almost contiguous frequency coverage across each of the 8 GHz bands with 400 MHz steps in centre frequency. *DFS* was recalculated with a full set of 42 channels with an error covariance constructed by interpolating the eigenvectors of the covariance for the original 12 channels by the method described in §4.3.2.1. However, this was found to produce an *ill-conditioned* error covariance matrix, resulting in unstable results when attempting to invert it [Rodgers, 2000]. Instead their error covariance was approximated as a diagonal matrix, defined to give results similar to those found in §7.1.1, and this was interpolated to the frequencies of the new channels. This showed that adding 14 extra channels in the 51-59 GHz band improved the *DFS* for temperature by ~0.4 and its vertical resolution by ~12% above 1 km in both clear and cloudy conditions. However, the extra 14 channels in the 22-30 GHz band only improved the *DFS* for humidity by ~0.15. These relatively small gains come at the price of longer overall cycle times and support the choice of default channels used in this instrument.

The 183.31 GHz water vapour line is much stronger than the 22.235 GHz line currently used by ground-based radiometers, such as the Radiometrics TP/WVP-3000, so has the potential to improve the accuracy and resolution of humidity profile retrievals. (Channels near 183 GHz are also used in satellite instruments for sounding humidity profiles.) An error analysis

conducted as above shows that adding 3 channels centred at 183 GHz (± 1 , ± 3 , ± 7 GHz) makes little difference to the temperature profile, but improves humidity in 1-3 km, adding ~ 0.26 DFS for humidity and improving its vertical resolution. At present this represents a small additional benefit ($< 5\%$ improvement in humidity information) for the additional cost of these extra channels. Ground-based radiometers operating at these frequencies are now becoming commercially available [Cadeddu *et al.*, 2006], primarily aimed at applications in polar regions, where water vapour quantities are typically very low, resulting in low sensitivity of channels near 22 GHz. As technology improves these frequencies will become increasingly affordable, while greater potential for miniaturisation and pressure on the spectrum at lower frequencies may make a move to 183 GHz more attractive in future.

7.1.5 Trade-off channel and angle selection

In choosing the optimum configuration of a particular instrument, there is a trade-off between radiometric noise, observing cycle and retrieval performance, which must consider the requirements for each particular application. For example, considering the assimilation of data from a ground-based microwave radiometer, the characteristics of the NWP model determine the time window available for observations. While mesoscale models typically have time-steps ~ 100 s, their horizontal resolution (~ 10 km) determines the time scale for advection (~ 1000 s) and the background error covariance. This may be considered as the available time window for observations (although it is likely to reduce in future as the resolution of NWP models increases). Although it is possible to average all observations over this time period prior to the retrieval, this is not recommended given the variability of cloud and humidity, which dominates the error characteristics of some channels. At the other extreme, all the profiles retrieved from instantaneous observations could be averaged, although this would not produce a large improvement as the background errors are assumed fixed. However, it is beneficial to average those channels that are sensitive only to temperature in the lowest 1 km prior to the retrieval process. It has also been shown that viewing a range of elevation angles can improve the retrievals, although this extends the observing cycle. These different methods of improving the retrievals are compared here, based on the Radiometrics TP/WVP-3000.

There are delays associated with internal switching of the reference noise diodes, changing channel and elevation angle, which also need to be accounted for in determining the optimum configuration for this instrument. However, these details are not pertinent to general discussion, which will proceed on the basis that 12 channels can be sampled in 7 views (plus the internal reference target) in 60 s. These views can be all at zenith, or at 4 symmetric elevation angles and optionally averaged over a 300 s window. Alternatively, 42 channels could be sampled in the zenith view plus the internal reference view every 60 s, and, optionally, an elevation scan of 5 angles could be sampled in a 300 s window.

This calculation was done using the observation error covariances, \mathbf{R} , defined in §4.4.1, but with the dubious spectroscopic error removed, and the elements corresponding to the noisy 57.7 GHz channel substituted with those from an adjacent channel. To simulate the observations of the high frequency channels (55-59 GHz) being averaged over this period, their components of \mathbf{R} were divided by the number of scans available and the DFS re-calculated using both $\mathbf{B}_{\text{ATOVs}}$ and \mathbf{B}_{NMC} background error covariances.

Table 7-2 Degrees of Freedom for Signal in temperature (DFS_T) and humidity (DFS_q) available from Radiometrics TP/WVP-3000 in different configurations after averaging 55-59 GHz observations over 300 s before retrievals.

	Instrument Combination and Configuration	$\mathbf{B}_{\text{ATOVs}}$				\mathbf{B}_{NMC}			
		Clear		Cloudy		Clear		Cloudy	
		DFS_T	DFS_q	DFS_T	DFS_q	DFS_T	DFS_q	DFS_T	DFS_q
(a)	Radiosonde	8.6	7.1	8.6	7.1	11.7	11.4	11.7	11.4
(b)	Surface sensors only	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
(c)	(b) + Radiometrics TP/WVP-3000	3.2	2.0	3.3	3.0	3.6	1.9	3.7	3.2
(d)	(c) + infrared radiometer (IR)	3.2	2.0	3.3	3.0	3.6	1.9	4.1	3.3
(e)	(c) at 4 elevation angles +zenith IR	4.4	2.7	4.4	5.0	5.5	3.4	5.7	6.3
(f)	(b) + 21 chs 22-30 GHz + 21 chs 51-59 GHz + IR (all zenith)	3.4	2.4	3.4	4.0	3.7	2.9	4.1	3.8

The results in Table 7-2 show that the normal zenith view of the microwave radiometer now provides 2.2 DFS for temperature and 1.0 for humidity more than the surface sensors in clear conditions. In cloudy conditions DFS_q increased by 1.0-2.3, while DFS_T remained unchanged. Although averaging 30 zenith observations in the 55-59 GHz channels increased the DFS for temperature, this improvement is not as large as that potentially available by spending the same time scanning 4 elevation angles and averaging the same channels over the same period. However, it should be remembered that these are theoretical figures only and do not account for the practical problems encountered at low elevation angles in the low frequency channels. The results for the 42 channels based on interpolated observation errors are unstable and subject to large errors.

The analysis error covariance for each observation strategy described above and listed in Table 7-2 are plotted in Figure 7-2 as $\sqrt{\text{diag}(\mathbf{A})}$ compared to the background error plotted as $\sqrt{\text{diag}(\mathbf{B}_{\text{ATOVs}})}$. In these clear conditions, d) is plotted over c) as the infrared radiometer does not add any extra information.

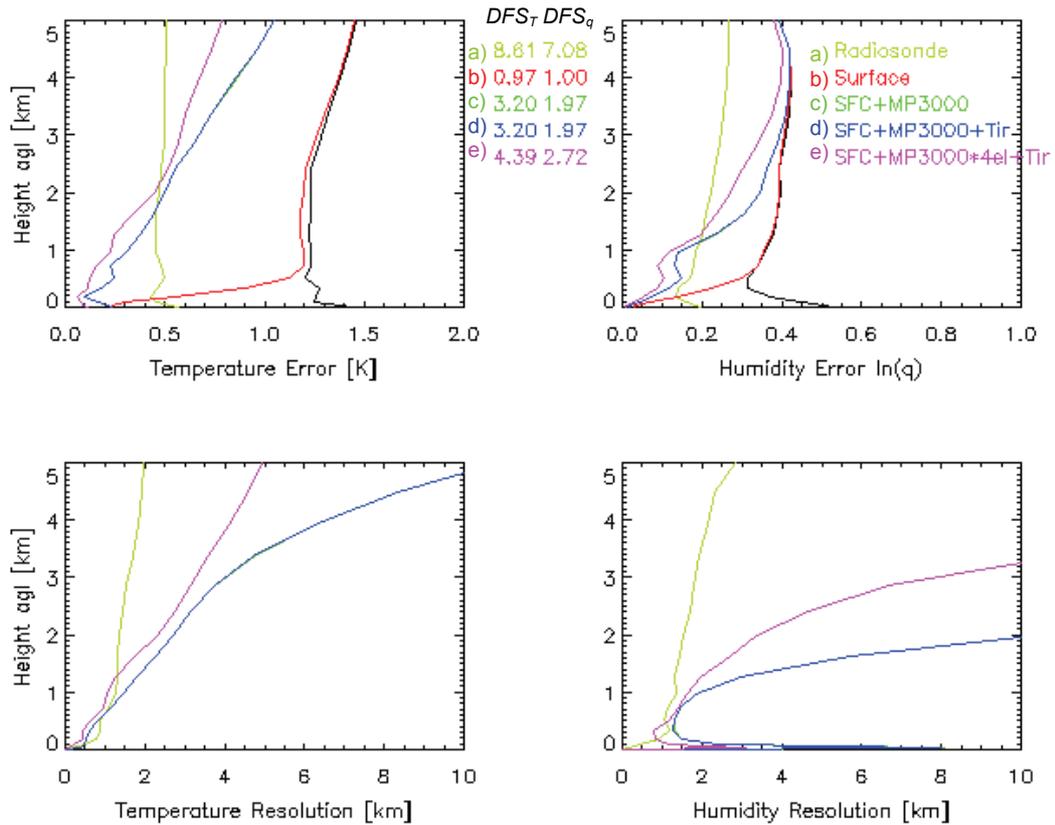


Figure 7-2 Error Analysis for 1D-VAR retrievals in clear US standard atmosphere. Background error covariance matrix from mesoscale model, $\sqrt{\text{diag}(\mathbf{B}_{\text{ATOVs}})}$ (black) and analysis error covariances matrices, $\sqrt{\text{diag}(\mathbf{A})}$, with surface sensors only (red), surface sensors and microwave and infrared radiometer in zenith view (blue), and for elevation scan of 4 angles (purple). 55-59 GHz channels averaged over 300 s window. Diagonal components plotted for lowest 5 km of temperature [K] and humidity (lnq).

More detailed channel selection trade-offs have been performed by Solheim *et al.* [1996], Cadeddu *et al.* [2002] and Rabier *et al.* [2002]. Solheim *et al.* [1996] described the eigenvalues analysis originally used to select the frequencies for the Radiometrics TP/WVP-3000 channels. Cadeddu *et al.* [2002] used a wavelet technique to quantify the vertical resolution for different configurations of microwave radiometers, which were compared in terms of their information gain. They showed a multi-channel system at a fixed viewing angle provides more information on the temperature profile above 1 km than a scanning single channel system, while the two configurations gave similar results nearer the surface. They also found that the best results were achieved by increasing the bandwidth in a multi-channel system – this is akin to the conclusions here on the benefit of reducing the noise on the high frequency channels. However, these methods have generally not attributed much benefit to placing the channels in bands protected for passive use to minimise the likelihood of interference (§3.6).

7.1.6 Recommended observing sequence

The following observing sequence is recommended for the current configuration of the Radiometrics TP/WVP-3000, based on the results of the trade-off between observing channels and view angles in §7.1.5.

A sequence should be scheduled to cycle between 6 views in turn in a 60 s period. These views should normally include the internal reference target and elevation angles 165.5°, 150°, 90°, 30° and 14.5° and each view allocated a total of 10 s. This sequence can be used for both tip curve calibration of the optically thin channels (22-30 GHz) and for elevation scanning retrievals, subject to tests on the consistency of symmetric views. However, the lowest elevation angles should not be used in the tip curve calibration and need not be observed by the water vapour channels. An alternative configuration of repeating the zenith view 5 times, interspersed with a view of the internal reference target would allow more rapid sampling where conditions are not favourable for tip curve calibrations or for elevation scanning to be exploited (i.e. when the atmosphere is not sufficiently horizontally homogeneous). This would allow more frequent cloud information to be provided in these situations where it is likely to be more variable. Switching between these modes could be achieved automatically based on the variance of recently observed brightness temperatures.

The high frequency channels (55-60 GHz) could then be filtered to reduce their radiometric noise by applying a rolling average, with a window of 300 s, prior to using their data in retrievals. Calibration data should also be averaged over this period prior to calculating the brightness temperatures to reduce the impact of this on radiometric noise (§3.5.1).

A period of 5 minutes should be set aside each day to allow scans of the 5 elevation angles at high spectral resolution to be performed, specifically to monitor cases of RFI. The brightness temperatures measured in this period should be analysed for self-consistency between adjacent frequencies and angles, in line with theoretical expectations, based on radiative transfer modelling. The scheduling of this period should be varied to allow detection of transmissions that are not constant in nature.

7.2 Integrating data from other instruments

The variational method allows different instruments to be combined and can provide a basis for the development of *Integrated Profiling Systems*, which have the potential to improve the profiles retrieved from the microwave radiometer alone. This has already been demonstrated to a limited extent in this thesis (§4.4). Although the inclusion of surface temperature and humidity sensors in the observation vector was trivial, as they provide a direct observation of two elements of the state vector, the infrared radiance needed a more complex forward model to describe its sensitivity to the humidity profile as well as the cloud-base temperature. For

each new observation type a forward model operator is required together with estimates of the observation error covariance. In future the 1D-VAR retrievals will be further extended to include observations from other instruments, described below. These are expected to improve the vertical resolution of the retrievals.

7.2.1 Ceilometer

Ceilometers (or *laser cloud base recorders*) are simple backscatter LIDAR (Light Detection and Ranging) systems that form part of the operational network of observations, commonly deployed at airports to provide measurements of the cloud-base height. They work by transmitting laser pulses near vertical and measuring the backscattered signal in a range of discrete gates. Although ceilometers typically operate in the near infrared (e.g. 905 nm), they are capable of penetrating thin cloud and able to measure the height of up to 3 distinct cloud layers. Ceilometers are typically configured to report cloud bases up to 8 km every 30 s. In fact, they are capable of providing much more information in the form of full backscatter profiles, which can indicate the height of inversions that trap aerosol layers near the surface. However, this is not currently exploited operationally due to difficulties interpreting their signal, which is a complex function of aerosol content, type and humidity. The signal can also be adversely affected by the sun when it is high in the sky. Lidars operating at infrared wavelengths are often simpler to interpret and could also be used in a similar way to ceilometers.

The height of the lowest cloud base could be used in the 1D-VAR retrieval in a similar way to the infrared brightness temperature. It could either be used as a cloud classification indicator (§4.8.1) and the first guess modified to ensure that it is cloud free below the cloud base, but saturated at that level – or by constructing a simple forward model for the cloud-base height and including this in the observation vector. The latter approach is likely to cause convergence problems of the type encountered with the infrared radiometer due to its non-Gaussian error characteristics (i.e. on/off). The ceilometer backscatter signal is generally more sensitive to cloud than the thermal infrared radiance and in many cases it will detect cloud layers that are transparent at microwave frequencies. For these reasons, it is unlikely that the ceilometer alone will have a large impact on the retrievals when an infrared radiometer is also available, although it remains attractive in principle, as this combination would provide both the height and temperature of the cloud base, and thus fix a point in the profile.

It may also be possible in future to use time-series of ceilometer observations to estimate profiles of fractional cloud occurrence in situations with broken cloud. This could be used to complement the time-series of retrievals from microwave radiometer observations with high time-resolution.

7.2.2 Cloud Radar

Cloud radars typically operate in the designated bands at 35 GHz and 94 GHz, where backscatter is dominated by liquid and ice drops, for which signal is proportional to their number density and the 6th moment of the drop size distribution [Westwater *et al.*, 2005]. This results in a much greater sensitivity to large cloud droplets, especially precipitation and insects, while clouds comprising of small droplets, such as small cumulus and fog may go undetected. It also makes the relationship between backscatter and liquid water content a nonlinear function of the drop size distribution. Cloud radars operate in window regions of the spectrum, where absorption is dominated by the water vapour continuum, which becomes significant at higher frequencies and must be accounted for in quantitative interpretation of the signal.

In their simplest form, observations from a cloud radar can be used to define cloud boundaries, which can then be combined with microwave radiometer data in a retrieval by methods similar to those suggested above for ceilometers. A more sophisticated approach is to define a relationship between the backscatter and the liquid water content (*LWC*) and use this as a forward model to include the backscatter profile in the observation vector. This approach was adopted in the *Integrated Profiling Technique* [Löhnert *et al.*, 2004] for pure liquid water cloud. Cloud radar and ceilometer data have been combined to classify cloud conditions in a profile as either clear, pure liquid water, pure ice or mixed phase cloud, or drizzling cloud [Hogan and O'Connor, 2006]. These classifications may be used to improve the accuracy of the relationship between backscatter and *LWC*.

7.2.3 Integrated Water Vapour column from GPS

Data from high precision GPS (Global Positioning System) receivers at fixed positions can be used to estimate the Integrated Water Vapour (IWV) by measuring the phase delay of signals transmitted from multiple GPS satellites [Bevis *et al.*, 1992]. This information is also available from microwave radiometer observations at two or more channels in the 20-40 GHz band and with comparable accuracy (§5.5), so the integration of the two systems does not seem to offer much potential for improving the retrievals. However, GPS sensors may allow a reduction in the number of channels needed for a microwave radiometer to retrieve temperature profiles accurately above 1 km (§4.9.6) or provide a calibration reference for water vapour channels.

7.2.4 Wind profiling radar signal

A wind profiler is a ground-based Doppler radar with 3 or more beams which measures signals backscattered from refractive index fluctuations associated with turbulence. The Doppler components of the beams from different directions are used operationally to estimate profiles of the wind vector, assuming the wind field to be homogeneous over their projected area. Typically wind profilers operate in the UHF band for profiling the boundary layer, while

VHF systems can be used to profile most of the troposphere. In addition to measuring the wind vector, the signal to noise ratio from wind profiler radars is sensitive to the magnitude of the gradients in the refractive index. These provide indications of the height of temperature inversions (for VHF systems in the upper troposphere) and hydrolapses (for UHF systems in the lower troposphere), which can be used to monitor the height of the top of the boundary layer. This information could improve profiles retrieved from microwave radiometer data.

The first guess profiles could be modified from the backgrounds available from NWP models so the height of any temperature inversions or hydrolapses are consistent with the profile of signal to noise ratio observed by a co-located wind profiler. However, attempts to modify the first guess in a similar way did not improve the retrieval of an inversion that was present at the wrong height in the NWP model background (§4.9.7). So it is unlikely that this technique alone would be more successful when applied to wind profiler data, without additional modifications to the assumed background error covariance. A slightly different approach is to use only the height of the wind profiler's peak signal strength by including an additional term in the cost function to constrain this to coincide with corresponding features in the state vector.

A more sophisticated approach would be to construct a reliable forward model to predict the wind profiler's signal to noise ratio from the NWP model fields, together with estimates of its error characteristics. This may be possible to model the signal from the elements used in the state vector of the retrievals described here – i.e. temperature and humidity. However, a more accurate model may need additional terms in the state vector to account for turbulence. A thorough review of this topic was performed by Gaffard and Nash [2006].

7.3 Applications

Remote-sensing observations are most useful for nowcasting applications when they provide extra information over and above what can be seen by cheaper in situ surface sensors. As many of the situations where observations can be helpful are driven by surface energy exchanges, surface sensors may provide the first indication of change – for example where fog lifts as a result of surface heating. However, the radiometer retrievals can be useful for monitoring to see if inversions weaken to allow the dissipation of a surface fog layer (§6.3).

7.3.1 Research applications

Cloud radiative properties are of prime importance in the detection of cloud-climate interactions, which currently dominate the uncertainty in the radiative forcing of future climate predictions [IPCC, 2001]. Calculation of the radiative properties of inhomogeneous cloud remains a central problem in cloud physics, which can be addressed by observational evidence [Wood and Taylor, 2001]. Microwave radiometers with at least two channels in the

region of 20-50 GHz can be used to retrieve column amounts of Integrated Liquid Water (ILW) due to the different absorption characteristics of the water vapour continuum and liquid water (§6.2). Price [2004] used observations of the statistical characteristics of ILW from a microwave radiometer to develop a parametrization for the albedo of inhomogeneous boundary layer cloud.

Another application for which radiometer observations have proved useful is in the re-analysis of case studies, where their information can help make up a more complete picture of the atmospheric state and aid understanding of the processes involved. These applications do not require real-time retrievals and may benefit from more intensive processing and more careful retrospective calibration. An example of this is the identification of changes in the profile's static stability during convection (§6.1.2).

7.3.2 Demonstrating impact on NWP

Observations from ground-based microwave radiometers have the potential to address recognised deficiencies in mesoscale NWP models. These include the height and strength of boundary layer capping inversions [Browning *et al.*, 2006], humidity profiles [Nash *et al.*, 2006] and integrated cloud liquid water [Price, 2004].

When considering assimilation of new observation types into NWP models, it is desirable to demonstrate that they will have a positive impact on the accuracy of the forecasts. Typically this requires parallel runs of NWP models over an extended period (>1 month) with and without assimilation of data from real instruments. In turn, this requires reliable instruments to be operated and substantial investment in the development of code to quality control, process and assimilate their observations. However, this appears to be a “chicken and egg” problem to justify the investment in the instruments and the code needed to assimilate their observations.

In the present version of variational assimilation used at the Met Office, the background error covariance is a property of large ensembles of model forecasts, i.e. they represent the global error characteristics of the background field and are assumed to be constant between assimilation cycles. Although the assimilation of an observation reduces the analysis error, this grows again during the forecast step to the level of the background error by the next assimilation cycle. In this case, it is difficult to demonstrate an impact from the assimilation of high temporal resolution observations that are statistically consistent with the model background. For an observation's increment to be maintained in the model, it must be converted into a set of analysis increments that are dynamically consistent [Lorenc and Hammon, 1988]. In a situation where a new single observation is assimilated what needs to be considered is how the global error characteristics change in response to that new observation type (when weighed against the existing large observation network driving the

analysis). Locally, the analysis may be improved significantly from a single observation. However, this must be maintained in balance with the surrounding model state. This problem is common to the assimilation of data from all new observing networks, particularly ground-based systems, which are typically deployed piecewise. *Observing System Simulation Experiments* partially address this problem, but still require a full assimilation system and realistic simulation of all observation types, with associated development costs.

In the absence of full assimilation trials, 1D-VAR retrievals from a prototype instrument can be used to demonstrate the accuracy of profiles retrieved from its data. The performance of the 1D-VAR retrievals has been compared with the theoretical accuracy from the error analysis (§4.6.1) and can be compared to the stated user requirements (§1.1.5). Furthermore the residuals of the observations with respect to the forward modelled background have been used to identify instrument, forward model and/or background biases and confirm their random errors are consistent with the assumptions (§4.6.3). To demonstrate an impact on the analysis, offline 1D-VAR retrievals have been set up, based on NWP model background fields. Experiments have been conducted to assess whether the observations can help add absent features to a background profile, suppress erroneous features and correct vertically displaced features. These experiments included an independent data source as truth – radiosondes that have not been assimilated. This allowed the demonstration of which model parameters are improved by the use of data from the radiometer. These results are summarised in §4.9.

While it is unlikely that data assimilation will be conducted at the full resolution of the model, future improvements in the model's resolution are expected to reduce the magnitude of the representativeness errors, which dominate the error budget of some channels sensitive to humidity and cloud.

7.3.3 Other retrieval techniques

Both the 1D-VAR and neural network described above are physical retrievals – they use a forward model based on physical principles to translate from state space to observation space. It is also possible to develop retrievals based on a purely empirical regression between the observations and state space. This technique was demonstrated by Güldner and Spänkuch [2001] using a data set of 18 months' observations from a Radiometrics TP/WVP-3000 and co-located radiosondes from Lindenberg, Germany. Although their technique retrieved profiles with lower variance against radiosondes than the neural network, it requires extensive data sets and is limited in application to the location of the training data. Cimini *et al.* [2006b] showed this technique produced similar results to the neural network when applied to a different data set.

The problem of introducing or moving a feature in the background profile is common to retrievals from other systems. This needs further investigation to fully exploit observations from *Integrated Profiling Systems*. It may be possible to modify the correlations in the background error covariance, \mathbf{B} , dynamically to enforce a discontinuity – for example, at the top of the boundary layer. Although this approach is attractive in a 1D framework, it becomes more problematic when applied to retrievals in higher dimensions because of the need to maintain balance between model fields. This balance is partly achieved through the correlations in \mathbf{B} , so if used in 3D-VAR and 4D-VAR it would be difficult to recompute dynamically. Future developments in data assimilation are needed to allow full use of this and many other observation types. For example, the retrieval problem can be re-defined as minimising the position (height, temperature) of significant points in the profile [Berkels *et al.*, 2004]. This technique has the potential to address the problem of variational retrievals' inability to correct features misplaced in the background profile.

7.3.4 Filling-in observing networks

One of the key potential applications of a ground-based microwave radiometer is to form part of a network of upper air observing systems, complementing radiosondes and aircraft profiles, either co-located or spatially distributed. In this concept, data from the radiometer is used to *fill-in* gaps in the observing network – either spatially or temporarily.

Löhnert *et al.* [2007] presented an analysis of the benefit of observations from an integrated profiling system (including a microwave radiometer) for NWP models using an optimal estimation framework, referred to as the *Integrated Profiling Technique*. This was based on calculating the analysis error resulting from the retrieval of temperature, humidity and cloud profiles from an integrated profiling system, using synthetic backgrounds, simulating the *a priori* information that would be available from the nearest or the latest radiosonde profile. Where radiosondes have been launched recently from nearby sites, the background error covariance, \mathbf{B} , is tightened up, while it is relaxed as time elapses since the launch, or as we move further from the launch site. Although they used NWP model fields to simulate a radiosonde network with variable density, their technique is independent of the NWP model and allows the benefit of different observing networks to be quantified. They express this in terms of *profit*, which they define as the average reduction in r.m.s. error of the temperature and humidity profile in the lowest 4 km. Like the *Degrees of Freedom for Signal* used in this study (§7.1), this is dependent on the assumptions inherent in the \mathbf{B} matrix.

Löhnert *et al.* [2007] concluded that the accuracy of the temperature profile degrades linearly with time/distance from the radiosonde that provides the background up to ~12 hr/400 km. However, for humidity profiles, they found that radiosondes more than the equivalent of 6 hr/200 km from the observation did not provide a more accurate background than a climatological average. This illustrates the greater atmospheric variability of humidity than

temperature. Their results can be used to specify the network density required to meet a given user requirement for the accuracy of temperature or humidity profiles. For example, a typical *breakthrough* user requirements for the accuracy of temperature and humidity profiles of 1 K and 10% (§1.1.5) would require a network densities of <400 km and <100 km, respectively. This can be compared to the equivalent spacing of radiosonde observations in the UK of 250 km. However, to achieve the *maximum* user requirements of 0.5 K would require an equivalent network density of 150 km for temperature, whereas the stated *maximum* user requirement of 5% for humidity is not achievable at any density of these observations.

It is difficult to extend this analysis to include profiles from commercial aircraft on the upper air network as their distribution is highly irregular in space and time. However, each participating airport could be crudely considered as an additional radiosonde station for temperature.

7.3.5 Real-time Retrieval Visualisation

In addition to the assimilation of their data into NWP models, there may be an ongoing requirement to visualise microwave radiometer observations in near real-time to monitor the performance of the hardware and to provide guidance to forecasters. For this application it would be beneficial to continue the development of neural network retrievals, as these can be implemented in real-time and independently of NWP to allow radiometer observations to be visualised in state space, which is meaningful to the forecaster. Proprietary software now allows real-time processing and display of radiometer data and retrieved profiles as time height cross-sections and vertical profiles for comparison with radiosondes. This should be further developed to increase the flexibility for the user to zoom in on areas of interest and select more retrieved parameters, such as static stability (§6.1.2) and thickness (§6.4). The user requirement for nowcasting systems states the optimum delay of <5 minutes between observation time and the availability of their data (§1.1.5). This does not present a problem if retrievals are implemented as neural networks, but it does limit the time window available to average the observations or the profiles retrieved from them.

7.4 Exploiting High Time-Resolution Observations

The geometry of ground-based observations typically means their information is concentrated in the planetary boundary layer. This is particularly beneficial to NWP as it complements the information available from aircraft and satellites over land, the application of which near the surface is limited by variable emissivity and surface temperature in the case of microwave sounders and extinction by cloud for infrared sounders. For this reason, ground-based observations are expected to have most impact on very short-range NWP.

One of the greatest potential benefits of ground-based observations is their high temporal resolution. However, to fully exploit the high time-resolution of data available from these instruments requires a 4-dimensional assimilation method, such as 4D-VAR. These methods have the potential to correct gross errors in the timing of (for example) fronts passing the observation location, and this has knock-on effects on much larger scales. As well as direct assimilation of observations in 4D-VAR, it is also possible to use it to assimilate profiles retrieved by 1D-VAR [Marécal and Mahfouf, 2003]. 4D-VAR is computationally expensive. So, although it is now operational in global models at the Met Office and other centres, it is not expected to be implemented in operational convective-scale NWP models for several years – and it is in these high-resolution models that ground-based observations are expected to have most impact. Meanwhile, other methods of exploiting the observations' high time-resolution suggested in the following paragraphs could be investigated.

One of the basic assumptions about the potential benefit of observations from ground-based instruments is that their high temporal resolution can be used as a proxy for horizontal resolution in the direction of atmospheric advection. In 3D-VAR, their data is assimilated at nominal observation times with relatively short windows every few hours. In practice, most operational 3D-VAR schemes use the First Guess at Appropriate Time (FGAT) method, comparing the observation with the background at the observation time.

Lopez *et al.* [2006] demonstrated a novel assimilation method to exploit the high time-resolution observations is 2D-VAR using a single column NWP model. They set up this unconventional system to investigate the problems of assimilating observations from a co-located microwave radiometer, cloud radar, GPS and rain gauge. This assimilation aims to find initial temperature and humidity profiles that give the best fit between *a priori* information from the model and the observations over a 12 h time window at a given geographical location. They assimilated radiometer and GPS observations at 30 minute time-resolution, although the cloud radar data needed to be averaged in time to achieve convergence. The obvious limitations of this method are that it only produces increments of temperature and humidity co-located with observations and has no feedback from the 3D dynamics (unlike 4D-VAR) and that it cannot modify the forcing – e.g. by including convergence in 2D.

A relatively simple scheme has been demonstrated in this thesis that uses the variability seen by the observations to calculate the representativeness errors dynamically (§4.9.5). Although this showed little reduction in the variance of the retrieved profiles compared to co-located radiosondes, the errors in this comparison were limited by those of the radiosondes and the background, which was assumed to be fixed. Similar limitations were found when averaging retrieved profiles over a short period prior to validating them against radiosondes (§5.3.2). However, this method is expected to produce an average profile which is closer to the average of the NWP model's grid box.

It is also possible to use observations at high time-resolution to estimate the sub-gridscale variability within each model box and relate this to other quantities in the model. For example, the variance could be calculated as well as the mean of a series of profiles retrieved by 1D-VAR, or these could be used to estimate the cloud fractional occurrence within a model grid-box. The variance could be empirically correlated to other model parameters as indicators of convective activity. However, while the background error covariance remains constant, it will be difficult for this method to improve the accuracy of 1D-VAR retrievals. Alternatively, cloud fraction could be added as a control variable in retrievals using a time-series of observations.

It is also possible to use profiles retrieved from a time-series of observations, taking the analysis of one 1D-VAR retrieval as the background for the next and take its analysis error covariance as the background error covariance for the next. This makes the reasonable assumption that the observation errors are independent between successive retrievals. However, it also assumes that the errors introduced at different stages in the successive retrievals are independent, which is more dangerous as it could result in the reinforcement of retrievals that converge on erroneous minima.

7.5 Summary of Optimisations for Operational Use

The information content analysis showed instantaneous radiometer observations at zenith (including surface temperature and humidity sensors) provide ~3 degrees of freedom (*DFS*) for temperature and ~2 for humidity in clear conditions and ~1 more *DFS* for humidity under cloud, partially thanks to the use of infrared radiance. *DFS* for temperature could be increased by averaging the noise-limited high frequency channels before using them in retrievals or by adding more elevation angles with these channels. The benefits of these techniques were traded-off to recommend the optimum observing strategy for the existing radiometer configuration, which could approach ~5 *DFS* for temperature and ~2 for humidity in clear conditions, the latter increasing to ~3 under cloud.

This led to the specifications of a microwave radiometer suitable for use as part of an *Integrated Profiling Station*, where its data is integrated with that from a co-located ceilometer, wind profiler and cloud radar. Methods were discussed to exploit the high time-resolution of the observations from these instruments in NWP models.

Chapter 8 Conclusions

This closing chapter summarises the conclusions of the preceding material of this thesis and examines the prospects for the future use of data from ground-based microwave radiometers.

8.1 Hardware

Observations from ground-based microwave radiometers can be used to retrieve profiles of temperature and humidity and the Integrated Liquid Water path (ILW). Down-welling radiances measured in a range of frequencies near absorption features or at different elevation angles are sensitive to thermal emission from different levels in the atmosphere. This study concentrated on the use of a commercial instrument – the Radiometrics TP/WVP-3000. This has 7 channels in the oxygen band from 51-59 GHz for temperature profiling and 5 channels between the water vapour line at 22.235 GHz and the window region at 30 GHz to provide information on humidity and liquid water cloud. It can measure at 5 elevation angles in a 60 s cycle. This system has participated in several trials and experiments at different locations over 4 years (§1.5) and proved very reliable. During this period extensive liaison with the manufacturers has resulted in a greater understanding of the system and a number of improvements – both hardware, to improve its performance in rain and resilience to corrosion, and software, to speed up the observing cycle and improve the calibrations (§3.2).

The stability of the system was of great concern because the massive amplification needed to boost the thermal signals to detectable levels can result in gain fluctuations, which cause the accuracy to deteriorate between calibrations. Analysis of the radiometer's data from the trials has demonstrated it was sufficiently stable that these did not dominate the overall observation errors (§3.5.4). The observations have been validated against forward-modelled radiosondes in clear skies (§3.4.2), which showed the observations can be well modelled in the channels used for temperature profiling. These comparisons highlighted a bias in the calibration of the radiometer's high frequency channels, which has since been corrected, confirmed a dry bias in the radiosondes, and raised concerns about the accuracy of the absorption model near 51-53 GHz, which warrants further investigation. Budgets of the radiometric noise have also been developed (§3.5), showing, while it is adequate to allow retrievals of temperature and humidity profiles, it could be further improved by averaging the calibration coefficients and/or the observations over a few minutes. However, it was later shown that atmospheric variability was more important than instrument noise for those channels sensitive to humidity and cloud.

Although it has not been experienced during these trials, Radio Frequency Interference (RFI) poses the greatest threat to the future viability of microwave profiling due to greater pressures on the use of the electromagnetic spectrum. It is recommended that some of the following techniques are implemented in future radiometer designs to mitigate this risk: frequency, angular and polarisation diversity or statistical analysis of the signal (§3.6.1).

8.2 Retrieval Techniques

The retrieval of temperature and humidity profiles is *ill-posed*, as an infinite number of atmospheric states fit a given set of observations. To resolve this ambiguity requires additional *a priori* background information – either in the form of a short-range forecast from an Numerical Weather Prediction (NWP) model, or as a statistical climatology of profiles. Chapter 4 and Chapter 5 have shown these approaches can be implemented by different retrieval techniques: a one-dimensional variational (1D-VAR) retrieval and a neural network.

The 1D-VAR retrieval has been developed to allow radiometer observations to be combined with a background from an NWP model in an optimal way, which accounts for their error characteristics. This has been used to retrieve profiles of temperature, humidity and cloud using a novel total water control variable (§4.2.2). The retrievals required development of a forward model to translate state space to observation space (§2.1) and the specification of the error covariances of the NWP background (§4.3.2) and the observations, which include contributions from radiometric noise, forward model and representativeness errors (§4.4.1).

A series of experiments was conducted to investigate different aspects of the 1D-VAR method. The greater nonlinearity of the observations in cloudy conditions caused slow convergence using the Gauss-Newton minimisation method. This was improved by adopting the Levenberg-Marquardt method (§4.9.1.3). Other methods of improving the retrievals' convergence rate in cloudy situations were investigated by making the forward model more linear between the first guess and the solution (§4.8). Variational retrievals also have the advantage of providing an estimate of the error on the retrieved profile. Error analysis (§4.6.1) has shown the microwave radiometer system is expected to improve the NWP background up to 4 km, retrieving temperature profiles with <1 K uncertainty and 2.8 Degrees of Freedom for Signal (DFS) and humidity with <40% uncertainty and 2.2 DFS. In the lowest 1 km, the accuracies improve to <0.5 K and <20%. Retrievals from real observations still improved the NWP backgrounds' temperature in the lowest 3 km, but had limited impact for humidity.

Neural networks offer an alternative approach to the inverse problem by combining radiometer observations with a statistical climatology of radiosonde profiles. The neural network itself is a nonlinear statistical regression between a large ensemble of observation vectors and state vectors and was found to generate biases if applied outside the range of its training data set (§5.3.4). For this reason, it is recommended that the range of the training data set is extended when training neural networks in future. However, this extension of the training data set represents a trade-off between performance in extreme conditions and the average conditions, which may be slightly degraded as a result of less representative background data. However, it may be possible to improve the representativeness of the background by training a new neural network to include time and date as additional inputs (§5.7.2).

The 1D-VAR retrievals have been shown to be advantageous over methods taking their background from statistical climatology, such as the neural network (§5.7.1), primarily because of the more accurate background information available from NWP models. However, variational methods have the disadvantages of being computationally expensive and requiring more investment to develop. Despite their problems, retrievals from the neural network are fast and, being independent of NWP, can be used to validate their forecasts.

Retrievals using elevation scans have also been developed (§5.4). These were found to give small improvements in the retrieved temperature profile, but their application to humidity profiling was limited by antenna sidelobes and atmospheric variability.

8.3 Benefits and Limitations

Microwave radiometer observations can be used to retrieve profiles of temperature and humidity as well as the Integrated Liquid Water path (ILW) or be assimilated directly into NWP models. This analysis has shown the retrieved profiles have only moderate vertical resolution in the boundary layer, which further degrades at higher levels (§4.6.2) to fall short of the *minimum* user requirement above 3 km. However, ground-based systems have the benefit of high time-resolution (~1 minute). The profiles can be used to calculate geopotential thickness or static stability and can represent the basic features of the convective boundary layer's development (§6.1). They can also monitor the evolution of near-surface inversions, although they can only resolve their general structure, not detailed features (§6.2).

The performance of the temperature profile retrievals approach the *breakthrough* user requirements for regional NWP (§1.1.5). However, the humidity profiles fall short of the stated *minimum* requirements due to the large spatial and temporal variability of water vapour combined with its poor representation in NWP background fields.

In addition to its poor vertical resolution, there are other limitations to the use of microwave radiometer observations. Together with the formulation of the variational method, this makes it very difficult for the retrievals to modify the height of features in the profile if they are incorrect in the background (§4.9.7). This problem is in common with other observation types. Radiometers are also sensitive to scattering by falling rain drops and their calibration is degraded by emission from liquid water on the instrument's window. However, the effect of this was much reduced by the introduction of a hydrophobic coating and improved blower mechanism to the point where retrievals were not significantly biased in light rain and drizzle (§5.5.2).

8.4 Information Content of Observations

Following the analysis of observation errors it became clear that the accuracy of the temperature profiles could be improved if the radiometric noise of the channels near the centre of the 60 GHz oxygen band could be reduced – e.g. by averaging successive samples. The information content of the observations was quantified in the variational retrieval framework using the concept of Degrees of Freedom for Signal, *DFS* (§7.1). This allowed different observing strategies to be traded-off (§7.1.5). For example, averaging more frequent zenith observations was found to give less benefit than performing less frequent elevation scans. The additional channels in the water vapour band were found to provide little more information than using a single channel here to provide information on the Integrated Water Vapour (IWV), equivalent to that available from cheaper GPS (Global Positioning System) sensors. However, this result depends on the IWV – more humidity information is available from the radiometer in warm, humid conditions.

8.5 Towards Operational Use

This study has highlighted a number of issues that need to be addressed if observations from ground-based microwave radiometers are to be used for operational applications. In addition to the recommendations mentioned in preceding sections of the conclusions, the paragraphs below list the logical order of activities required as part of a follow-on project:

Computational time can present a bottleneck in operational systems using variational retrievals from observations made at high time-resolution, and also in their development cycle. This should be improved by developing a *fast model* to calculate the observation vector and its Jacobians from the state vector. It may be possible to modify the geometry of the RTTOV fast model (§2.5.4) used for satellite applications to suit ground-based instruments.

Quality control is of key importance to operational applications. This should include flagging dubious observations, based on their consistency with previous or background values. Tools need to be developed to allow routine monitoring of *Observations-minus-Background* residuals as part of the assimilation cycle (§4.6.3). The application of neural network retrievals should also include checks that the observations are not outside the range of the training data set, as this can induce gross errors (§5.3.4).

To further investigate the performance of microwave radiometers for operational use and the benefit of their data to NWP it is recommended that radiometers are operated at different sites as part of *Integrated Profiling Stations* for extended periods (>1 year). This would allow full characterisation of the instruments' biases with respect to an NWP model, validation of the forward model and error characteristics. Data sets from such *testbed* deployments can be

used to identify the impact of assimilating the observations on the NWP analysis and, ultimately, its forecasts.

These data sets can also be used to develop methods of integrating observations from the microwave radiometer with those from other instruments. The variational method provides a suitable framework to perform this integration. This has been demonstrated to a limited extent by the inclusion of infrared radiances in the observation vector (§4.4) to provide information on the cloud-base temperature. This method should be extended to include observations from ceilometers, cloud radar and wind profiling radar (§7.2). It is expected that this will help improve the vertical resolution of the retrievals.

The results of the trade-off between observing channels and view angles (§7.1.5) led to a recommended observing sequence for the Radiometrics TP/WVP-3000 (§7.1.6). Further study is needed to confirm the suggestion in this thesis that the channels in the radiometer's low frequency band (22-30 GHz) could be replaced with a high precision GPS receiver (§7.1.4.1). This would provide the information on IWV required for temperature profiling, while simplifying the radiometer hardware. However, it may result in a reduction of humidity information in summer conditions, especially when cloudy.

During the past decade the price of a microwave radiometer has reduced by a factor of ~2 and technological developments have improved their accuracy and stability. Yet at ~£100k each they are still too expensive to deploy in a dense observing network. However, future developments in Monolithic Microwave Integrated Circuit (MMIC) technology are expected to allow the integration of key components of a 50-60 GHz radiometer into a single package (§3.7). It is recommended that these developments are supported, as they would allow mass production of radiometers suitable for profiling temperature at much reduced costs in the longer term.

Further research is also needed to exploit fully what is potentially the greatest benefit of ground-based microwave radiometers – their high time-resolution. A first step has been investigated in this thesis by calculating representativeness errors dynamically, based on a time series of observations (§4.9.5). Further developments (§7.4) may be achieved through different approaches, including using the analysis as the first guess in successive retrievals, retrieving sub-gridscale variability, 4-Dimensional Variational assimilation (4D-VAR) of 1D-VAR retrievals and, ultimately, 4D-VAR of the observations directly into the new generation of convective-scale NWP models. Together these offer the possibility of much improved short-range weather forecasts.

Acronyms and Abbreviations

Acronym	Meaning	Page
CMB	Cosmic Microwave Background	20
CSIP	Convective Storm Initiation Project	14
DFS	Degrees of Freedom for Signal	146
GPS	Global Positioning System	157
HATPRO	Humidity And Temperature PROfiling radiometer from Radiometer Physics	47
IF	Intermediate Frequency	46
ILW	Integrated Liquid Water content [g/m^2]	9
IR	Infrared	47
IWV	Integrated Water Vapour content [kg/m^2]	9
LIDAR	Light Detection And Ranging	156
LN2	Liquid Nitrogen	53
MP3000	Microwave Profiler 3000 model from Radiometrics, aka TP/WVP-3000	45
MPM	Millimetre-wave Propagation Model [Liebe and Layton, 1987]	23
NWP	Numerical Weather Prediction	1
O-B	Observation minus Background residuals	41
pdf	Probability Density Function	67
RFI	Radio Frequency Interference	62
RH	Relative Humidity	34
r.m.s.	Root mean square	12
Ros98	Rosenkranz [1998] Millimetre-wave absorption model, successor to MPM	24
RTE	Radiative Transfer Equation	19
RTM	Radiative Transfer Model	18
SD	Standard Deviation	31
TP/WVP- 3000	Microwave Profiler 3000 model from Radiometrics, aka MP3000	45
TSF	Time Structure Function	60
TUC	Temperature hUmidity and Cloud experiment, Payerne, Switzerland, Winter 2003/04	14
UM	Unified Model (Met Office Numerical Weather Prediction system)	1
VAR	Variational Assimilation/Retrieval	11

Symbols

Symbol	Meaning	Page
A	Analysis Error Covariance Matrix	81
B	Background Error Covariance Matrix	71
$H(\mathbf{x})$	Forward Model, mapping state space into observation space	79
H	Matrix of Jacobian vectors	67
I	The Identity matrix (diagonal terms = 1, off-diagonal terms=0)	81
ΔT_{\min}	Minimum detectable change in brightness temperature (radiometric resolution) [K]	57
$P(\mathbf{x})$	Probability Density Function (pdf) of \mathbf{x}	67
q	Specific humidity [kg/kg]	68
q_t	Total water content [kg/kg]	68
R	Observation Error Covariance Matrix	73
T_b	Brightness Temperature [K] – a measure of microwave radiance	19
T_{ir}	Infrared Brightness Temperature [K] – measured by IR radiometer	42
T_{MR}	Mean Radiative temperature of the atmosphere [K]	51
T_{CMB}	Effective T_b of the Cosmic Microwave Background [K]	20
T_{ND}	Noise Diode Brightness Temperature [K]	50
\mathbf{x}	State vector, representing profile of atmospheric temperature and logarithm of total water ($\ln q_t$)	68
\mathbf{x}^a	State vector of analysis (retrieval)	85
\mathbf{x}^b	State vector of background (prior)	70
\mathbf{y}^o	Observation vector, representing a set of radiances in different channels, surface temperature and humidity and infrared T_b	73
ν	Frequency [Hz]	
λ	Wavelength [m]	
τ	Opacity [Np]	
θ	Zenith Angle ($\theta=0^\circ$ at zenith)	

The standard notation suggested by Ide *et al.* [1997] has been adopted in this thesis.

References

- Agnew, J. and T.Hewison, 2004:** "Comparison of water vapour profiles measured using a Raman lidar and a microwave radiometer at Chilbolton Observatory, UK", Proceedings of 22nd ILRC (International Laser Radar Conference), July 12-16 2004, Matera, Italy, Also available from http://tim.hewison.org/Agnew_Hewison.doc.
- Allan, D.W., 1996:** "Statistics of Atomic Frequency Standard", Proc IEEE, 54, No. 2, 221-231.
- Berkels, B., U.Clarenz, S.Crewell, U.Löhnert, M.Rumpf and C.Simmer, 2004:** "A physical temperature profiling method using gradient flows", In "Microwave Radiometry and Remote Sensing Applications", Feb. 2004, Rome, Italy, 2004.
- Bevis, M.,S. Businger, T.A.Herring, C.Rocken, R.A.Anthes and R.H.Ware, 1992:** "GPS meteorology: Remote sensing of atmospheric water vapor using the global positioning system", J. Geophys. Res., 97, 15787-15801.
- Bloemink, H.I., A.C.A.P.van Lammeren, A.J.Feijt, and S.C.H.M.Jongen, 1999:** "Active-passive sensor synergy for cloud observation; IR cloud properties and cloud liquid water", in Remote Sensing of Cloud Parameters, Procs. Symposium, 21-22 Oct. 1999, IRCTR, Delft University of Technol., ISBN 90-804551-6-4, The Netherlands.
- Böttger, H., P.Menzel and J.Pailleux, 2004:** "Proceedings of the Third WMO Workshop on the Impact of Various Observing Systems on Numerical Weather Prediction", Alpbach, Austria, 9-12 March 2004.
- Boukabara, S.A., S.A.Clough, J-L Moncet, A.F.Krupnov, M.Tretyakov and V.V.Parshin, 2004:** "Uncertainties in the Temperature Dependence of the Line Coupling Parameters of the Microwave Oxygen Band: Impact Study", IEEE Trans. Geosci. Rem. Sensing, Vol.43, No.5.
- Browning, K., A. Blyth, P. Clark, U. Corsmeier, C. Morcrette, J. Agnew, D. Bamber, C.Barthlott, L.Bennett, K.Beswick, M.Bitter, K. Bozier, B.Brooks, C.Collier, C.Cook, F.Davies, B.Deny, M.Engelhardt, T.Feuerle, R.Forbes, C.Gaffard, M.Gray, R.Hankers, T.Hewison, R.Huckle, N.Kalthoff, S.Khodayar, M.Kohler, S.Kraut, M.Kunz, D.Ladd, J.Lenfant, J.Marsham, J.McGregor, J.Nicol, E.Norton, D.Parker, F.Perry, M.Ramatschi, H.Ricketts, N.Roberts, A.Russell, H.Schulz, E.Slack, G.Vaughan, J.Waight, R.Watson, A.Webb, A.Wieser, K.Zink, 2006:** "The Convective Storm Initiation Project", Submitted to Bulletin Am. Meteorol. Soc.
- Buehler, S.A., M.Kuvatov, V.O.John, U.Leitner and H.Dier, 2004:** "Comparison of Microwave Satellite Humidity Data and Radiosonde Profiles: A Case Study", J. Geophys. Res., 109, D13103, [doi:10.1029/2004JD004605](https://doi.org/10.1029/2004JD004605).
- Cadeddu, M.P., G.E. Peckham and C.Gaffard, 2002:** "The vertical resolution of a ground-based microwave radiometer analyzed through a multiresolution wavelet technique", IEEE Trans. on Geoscience and Remote Sensing, Vol.40, No.3, pp. 531-540.
- Cadeddu, M.P., J.C.Liljegren and A.Pazmany, 2006:** "Measurements and Retrievals from a New 183-GHz Water Vapor Radiometer in the Arctic", Proceedings of 9th Specialist

Meeting on Microwave Radiometry and Remote Sensing Applications (MicroRad'06),
Puerto Rico, 28 Feb to 3 Mar 2006.

- Chandrasekhar, S., 1960:** "Radiative Transfer", Dover, New York.
- Chylek, P. and V.Ramaswamy, 1982:** "Simple Approximation for Infrared Emissivity of Water Clouds", J. Atmos. Sci.: Vol. 39, No. 1, pp. 171-177.
- Churnside, J.H., T.A.Stermitz and J.A.Schroeder, 1994:** "Temperature Profiling with Neural Network Inversion of Microwave Radiometer Data", J. Atmos. Ocean. Technol., Vol.11, No.1, pp.105-109.
- Cimini, D., E.R.Westwater, Y.Han and S.Keihm, 2003a:** "Accuracy of Ground-based microwave radiometer and balloon-borne measurements during the WVIOP2000 field experiment", IEEE Trans. Geoscience Remote Sensing, Vol.31(11), pp.2605-2615.
- Cimini, D., E.R.Westwater, and Y.Han, 2003b:** "Theoretical Analysis of the Frequency Allocation of the Hinge Points Around 22.235 GHz", Proceedings of the Thirteenth Atmospheric Radiation Measurement (ARM) Science Team Meeting, ARM-CONF-2003, April 2003, Broomfield, Colorado.
- Cimini, D., T.J.Hewison and L.Martin, 2006a:** "Comparison of brightness temperatures observed from ground-based microwave radiometers during TUC", Meteorologische Zeitschrift, Vol.15, No.1, pp.19-25, [doi:10.1127/0941-2948/2006/0096](https://doi.org/10.1127/0941-2948/2006/0096).
- Cimini, D., T.J.Hewison, L.Martin, J.Gueldner, C.Gaffard and F.Marzano, 2006b:** "Temperature and humidity profile retrievals from ground-based microwave radiometers during TUC", Meteorologische Zeitschrift, Vol.15, No.1, pp.45-56, [doi:10.1127/0941-2948/2006/0099](https://doi.org/10.1127/0941-2948/2006/0099).
- Collard, A., 1998:** "Notes on IASI Performance", Forecasting Research Technical Report No.256, Available from National Meteorological Library, Met Office, UK.
- Collard, A., 2004:** "NWPSAF Met Office 1D-Var User Manual", NWPSAF-MO-UD-006, Available from National Meteorological Library, Met Office, UK.
- Crane, R.K., 1971:** "Propagation phenomena affecting satellite communications systems operating in the centimetre and millimetre wavelength bands", Proc. IEEE, 59, pp.173-188, as reported in Ulaby *et al.* [1981].
- Crewell, S., H.Czekala, U.Löhnert, C.Simmer, T.Rose and R.Zimmermann, 2001:** "Microwave Radiometer for Cloud Carthography: A 22-channel ground-based microwave radiometer for atmospheric research", Radio Sci., Vol.36, No.4, pp.621-638.
- Crewell, S., M.Drusch, E.van Meijgaard, and A.van Lammeren, 2002:** "Cloud observations and modeling within the European BALTEX Cloud Liquid Water Network", Boreal Env. Res. 7: 235-245.
- Crewell, S. and U.Löhnert, 2003:** "Accuracy of cloud liquid water path from ground-based microwave radiometry: 2. Sensor accuracy and synergy", Radio Sci., Vol.38, No.3, 8042.
- Czekala, H., S.Crewell, C.Simmer and A.Thiele, 2001:** "Discrimination of cloud and rain liquid water path by groundbased polarized microwave radiometry", Geophys. Res. Letters, Vol.28, No.2, pp.267-270.

- Debye, P.J.W. ,1929:** "Polar molecules", New York : Chemical Catalog Co., as reported in Ulaby *et al.* [1991].
- Deblonde, G. and S.J.English, 2003:** "One-Dimensional Variational Retrievals From SSMIS Simulated Observations", *Journal of Applied Meteorology* 42: 1406-1420.
- Dharssi, I., A.C. Lorenc and N.B. Ingleby, 1992:** "Treatment of gross errors using maximum probability theory", *Q.J.R. Meteorol. Soc.*, 118, 1017-1036.
- Ellison, W.J., S.J.English , K.Lamkaouchi, A.Balana, E.Obligis, G.Deblonde, T.J.Hewison, P.Bauer, G.Kelly and L.Eymard, 2003:** "A comparison of new permittivity data for sea water with AMSU, SSM/I, TRMM Microwave Imager and airborne radiometers observations", *J. Geophys. Res.*, 108 (D21), 4663, [doi:10.1029/2002JD003213](https://doi.org/10.1029/2002JD003213).
- English, S.J., 1995:** "Airborne radiometric observations of cloud liquid-water emission at 89 and 157GHz: Application to retrieval of liquid-water path", *Q.J.R. Meteorol. Soc.*, vol.121, pp.1501-1524.
- Eyre, J.R., 1987:** "On Systematic-Errors in Satellite Sounding Products and their Climatological Mean-Values", *Q.J.R. Meteorol. Soc.*, Vol.113, pp.279-292.
- Eyre, J.R., 1990:** "The information content of data from satellite sounding systems: A simulation study", *Q.J.R. Meteorol. Soc.*, Vol.116, No.492, pp.401-434.
- Gaffard, C. and T.J.Hewison, 2003:** "Radiometrics MP3000 Microwave Radiometer Trial Report", Observations/Development Technical Report TR26, Met Office, National Meteorological Library, Exeter, UK. Also available from <http://tim.hewison.org/TR26.pdf>.
- Gaffard, C. and J.Nash, 2006:** "Investigation of the Relationship Between Windprofiler Signal to Noise and Temperature, Humidity, and Cloud Profiles" in COST720 Final Report, Eds: D.Englebart, J.Nash and W.Monna, European Science Foundation, In press.
- Gaussiat, N., R.J.Hogan and A.J.Illingworth, 2006:** "Accurate liquid water path retrievals from low-cost microwave radiometers using additional information from lidar and operational forecast models", Submitted to *J. Atmosph. Oceanic Technol.*
- Giering, R., 1999:** "Tangent linear and Adjoint Model Compiler", Users manual 1.4, <http://www.autodiff.com/tamc> (accessed 19/11/04).
- Golitsyn, G.S., E.N.Kadygrov and I.N.Kuznetsova, 2002:** "Microwave Remote Sensing Investigation of the Atmospheric Boundary Layer Thermal Regime Above an Urban Heat Island", 12th ARM Science Team Meeting Proceedings, St. Petersburg, Florida, 8-12 April.
- Güldner, J. and D.Spänkuch, 2001:** "Remote Sensing of the Thermodynamic State of the Atmospheric Boundary Layer by Ground-Based Microwave Radiometry", *J Atmos. Ocean. Technol.*, Vol.18, pp.925-933.
- Güldner, J. and J.-P.Leps, 2005:** "Analysis of the CLIWA-NET intensive operation period data as part of the monitoring activities at the German Meteorological Service site Lindenberg", *Atmos. Res.*, Vol.75, pp.151-166.
- Han, Y. and E.R.Westwater, 2000:** "Analysis And Improvement Of Tipping Calibration For Ground-Based Microwave Radiometers", *IEEE-TGARS*, Vol.38, No.3, pp.1260-1276.

- Hewison, T.J., 1999:** "Land surface emissivity at millimetre wavelengths," MSc Dissertation, Department of Meteorology, University of Reading, UK. Also available from http://tim.hewison.org/MSc_dissertation.pdf.
- Hewison, T.J. and A.J.McGrath 2001:** "Performance assessment of Liquid Nitrogen Calibration Target supplied by Fred Solheim (Radiometrics) at 89, 157 and 183 GHz", MRF Technical Note 39. Available from National Meteorological Library, Met Office, UK and http://tim.hewison.org/Radiometrics_LN2.doc.
- Hewison, T.J., 2003:** "Microwave Radiometer Comparison at Cardington", Spring 2003, Obs. Development Technical Report TR30, Met Office, National Meteorological Library, Exeter, UK. Also available from <http://tim.hewison.org/TR30.doc>.
- Hewison, T.J. and C.Gaffard, 2003:** "Radiometrics MP3000 Microwave Radiometer Performance Assessment", Obs. Development Technical Report TR29, Met Office, National Meteorological Library, Exeter, UK. Also available from <http://tim.hewison.org/TR29.pdf>.
- Hewison, T.J., C.Gaffard and J.Nash, 2003:** "Validation of Microwave Radiometer Measurements in Clear Air", Proc. ISTP2003, Leipzig, Germany, September 2003, Also available from http://tim.hewison.org/ISTP_Hewison_Abstract.pdf.
- Hewison, T., C.Gaffard, J.Nash, D.Ruffieux, R.Nater, H.Berger, M.Perroud and B.Andrade, 2004:** "Monitoring Inversions From Ground-Based Remote Sensing Instruments During Temperature, Humidity, And Cloud Profiling Campaign (TUC)", Proceedings of microRad'04, Rome, 24-27 February 2004. Also available from http://tim.hewison.org/microRad04_Hewison.pdf.
- Hewison, T.J., 2006:** "Aircraft Validation of Clear Air Absorption Models at Millimeter Wavelengths (89-183 GHz)", J. Geophys. Res., 111, D14303, [doi:10.1029/2005JD006719](https://doi.org/10.1029/2005JD006719).
- Hewison, T.J., D.Cimini, L.Martin, C.Gaffard and J.Nash, 2006:** "Validating clear air absorption model using ground-based microwave radiometers and vice-versa", Meteorologische Zeitschrift, Vol.15, No.1, pp.27-36, [doi:10.1127/0941-2948/2006/0097](https://doi.org/10.1127/0941-2948/2006/0097).
- Hewison, T.J. and C.Gaffard, 2006:** "Report on Met Office Participation in Convective Storm Initiation Project (CSIP) at Linkenholt", Obs. Dev. Technical Report, Met Office, National Meteorological Library, Exeter, UK. Also available from http://tim.hewison.org/CSIP_linkenholt_data.doc.
- Hogan, R. J., N. Gaussiat and A. J. Illingworth, 2004:** "Stratocumulus liquid water content from dual-wavelength radar", J. Atmos. Oceanic Technol., 22, 1207-1218.
- Hogan, R. J. and E.J. O'Connor, 2006:** "Facilitating cloud radar and lidar algorithms: the Cloudnet Instrument Synergy/Target Categorization product", CloudNet product documentation, <http://www.cloud-net.org/data/products/categorize.html>, accessed 10/7/06.
- Houghton, J. T., F. W. Taylor and C. D. Rodgers, 1984:** "Remote Sounding of Atmospheres", Cambridge University Press.
- Ide, K. P.Courtier, M.Ghil and A.C.Lorenc, 1997:** "Unified Notation for Data Assimilation: Operational, Sequential and Variational", J. Meteor. Soc. Japan, Vol.75(1B), pp.181-189.

- Ingleby, B.N., 2001:** "The statistical structure of forecast errors and its representation in The Met. Office Global 3-D Variational Data Assimilation Scheme", Q. J. Roy. Meteorol. Soc., Vol.127, No. 571, pp.209-231.
- IPCC, 2001:** "Climate Change 2001: The Scientific Basis", Intergovernmental Panel on Climate Change, Cambridge University Press.
- Iturbide-Sanchez, F., S.C.Reising and R.W.Jackson, 2004:** "Design and Implementation of a Miniaturized Water Vapor Profiling Radiometer", Proc. IEEE Internat. Geoscience and Remote Sensing Symposium, Anchorage, Alaska, September (2004).
- Janssen, M.A., 1993:** "An introduction to the passive microwave remote sensing of atmospheres " in "Atmospheric remote sensing by microwave radiometry", Wiley Series in Remote Sensing, M. Janssen, Ed. Wiley, 1993.
- Jones, D.C., 1995:** "Validation of scattering microwave radiative transfer models using an aircraft radiometer and ground-based radar", PhD Thesis, Meteorology Dept., University of Reading.
- Kadygrov, E.N. and D. R. Pick, 1998:** "The potential for temperature retrieval from an angular-scanning single-channel microwave radiometer and some comparisons with in situ observations", Meteorological Applications, Vol.5, pp.393-404.
- Kitchen, M., 1989:** "Representativeness errors for radiosonde observations", Q.J.R. Meteorol.Soc, Vol.115, pp.673-700.
- Kleespies, T.J. and D.Crosby, 2001:**" A Method for Correlated Noise Modeling in Satellite Radiance Simulation", AMS preprint volume for the 11th conference on Satellite Meteorology and Oceanography, P5.26, 15-18 October 2001, Madison, Wisconsin.
- Levenberg, K, 1944:** "A method for the solution of certain nonlinear problems in least squares", Quart. Appl. Math, **2**, 164.
- Liebe, H.J and D.H.Layton, 1987:** "Millimetre-Wave Properties Of The Atmosphere: Laboratory Studies And Propagation Modeling", NTIA-Report 87-224, Natl. Telecommun. And Inf. Admin, Boulder, CO.
- Liebe, H.J., 1989:** "MPM - An Atmospheric Millimetre Wave Propagation Model, Int. J. Infrared and Millimetre Waves", Vol.10(6), pp.631-650.
- Liebe, H.J., G.A.Hufford and M.G.Cotton, 1993:** "Propagation modeling of moist air and suspended water/ice particles at frequencies below 1000GHz", AGARD 52nd Specialists' Meeting of the Electromagnetic Wave Propagation Panel, Ch3.
- Liljegren, J.C., S.A.Boukabara, K.Cady-Pereira and S.A.Clough, 2005:** "The Effect of the Half-Width of the 22-GHz Water Vapor Line on Retrievals of Temperature and Water Vapor Profiles with a Twelve-Channel Microwave Radiometer", IEEE Trans. Geosc. Rem. Sens., Vol.43, No.5, pp.1102-1108.
- Limb, M. and G.Spellman, 2001:** "Evaluating domestic tourists' attitudes to British Weather - a qualitative approach", Proceedings of the First International Workshop on Climate, Tourism and Recreation. Edited by A. Matzarakis and C. R. de Freitas, Report of a Workshop Held at Porto Carras, Neos Marmaras, Halkidiki, Greece, 5-10 October 2001.

- Löhnert, U. and S.Crewell, 2003:** "Accuracy of cloud liquid water path from ground-based microwave radiometry: 1. Dependency on cloud model statistics", *Radio Sci.*, Vol.38, No.4, 8041.
- Löhnert, U., S.Crewell and C.Simmer, 2004:** "An Integrated Approach toward Retrieving Physically Consistent Profiles of Temperature, Humidity, and Cloud Liquid Water", *J. Appl. Meteor.*, Vol.43, pp.1295-1307.
- Löhnert, U., E.Van Meijgaard, H.Klein Baltink, S.Groß and R.Boers, 2007:** "Accuracy assessment of an integrated profiling technique for operationally deriving profiles of temperature, humidity and cloud liquid water", *J. Geophys. Res.*, Vol.112, No.D4, D04205.
- Lopez, P., A. Benedetti, P. Bauer, and M. Janiskov'a, 2006:** "Experimental 2D-VAR Assimilation of ARM cloud and precipitation observations", *Quart. J. Roy. Meteor. Soc.*, 617, pp.1325-1347(23).
- Lorenc, A.C., 1986:** "Analysis methods for numerical weather prediction", *Quart. J. Roy. Meteor. Soc.*, Vol.112, pp.1177-1194.
- Lorenc, A.C. and O.Hammon, 1988:** "Objective quality control of observations using Bayesian methods. Theory, and a practical implementation", *Q.J.R. Meteorol. Soc.*, Vol.114, No.480, pp.515-543.
- Marchand, R., T.Ackermann, E.R.Westwater, S.A.Clough, K.Cady-Pereira, J.C.Liljegren, 2003:** "An assessment of microwave absorption models and retrievals of cloud liquid water using clear-sky data," *J. Geophys. Res.*, Vol.108, D24, doi:10.1029/2003JD003843.
- Marécal, V. and J-F.Mahfouf, 2003:** "Experiments on 4D-Var assimilation of rainfall data using an incremental formulation", *Q.J.R. Meteorol. Soc.*, Vol.129, No.594, pp.3137-3160.
- Marquardt, D.W., 1963:** "An algorithm for least-squares estimation of nonlinear parameters", *SIAM J. Appl. Math.*, **11**, 164.
- Martin, L., C.Mätzler, T.J.Hewison, D.Ruffieux, 2006:** "Intercomparison of integrated water vapour measurements during the TUC campaign", *Meteorologische Zeitschrift*, Vol.15, No.1, pp.57-64, [doi:10.1127/0941-2948/2006/0098](https://doi.org/10.1127/0941-2948/2006/0098).
- Mattioli, V., E.Westwater, S.Gutman and V.Morris, 2005:** "Forward Model Studies of Water Vapor using Scanning Microwave Radiometers, Global Positioning System, and Radiosondes during the Cloudiness Inter-Comparison Experiment", *IEEE Trans. Geosci. Rem. Sens.*, Vol.43, No.5, pp.1012-1021.
- Mätzler, C., E.R.Westwater, D.Cimini, S.Crewell, T.J.Hewison, J.Güldner and F.S.Marzano, 2006:** "Microwave Radiometers", in COST720 Final Report, Eds: D.Englebart, J.Nash and W.Monna, European Science Foundation, In press.
- McClatchey, R.A., R.W. Fenn, J.E.Selby, F.E.Volz and J.S.Garing, 1972:** "Optical properties of the atmosphere", 3rd. edn. AFCRL-72-0497, *Environ. Res. Papers* No. 411.
- McGrath, A.J. and T.J.Hewison, 2001:** "Measuring the accuracy of a Microwave Airborne Radiometer (MARSS)", *J. Atmos. Ocean. Technol.*, Vol.18, No.12, pp.2003-2012.

- Meissner, T. and F.J.Wentz, 2004:** "The Complex Dielectric Constant of Pure and Sea Water From Microwave Satellite Observations", IEEE Trans. Geosc. Rem. Sens., Vol. 42, No. 9, pp.1836-1849.
- Misra, S., C.Ruf and R.De Roo, 2006:** "Agile Digital Detector for RFI Mitigation", Proceedings of 9th Specialist Meeting on Microwave Radiometry and Remote Sensing Applications (MicroRad'06), Puerto Rico, 28 Feb to 3 Mar 2006.
- Mlawer, E. J., S.A.Clough and D.C.Tobin, 2003:** "The MT_CKD water vapour continuum: A revised perspective including collision induced effects", Atmospheric Science from Space using Fourier Transform Spectrometry (ASSFTS) Workshop, Bad Wildbad, Germany, 8-10 October 2003.
- Monna, W., J.Nash, D.Englebart, 2006:** "Introduction to Final Report of COST Action 720 – Integrated Ground-Based Remote Sensing Stations for Atmospheric Profiling", in COST720 Final Report, European Science Foundation, In Press.
- Nash, J., J. B. Elms and T. J. Oakley, 1995:** "Relative humidity performance observed in recent international radiosonde comparisons", Ninth Symp. on Meteorological Observations and Instrumentation, Charlotte, NC, Amer. Meteor. Soc., 43–48.
- Nash, J. C.Gaffard, T.Hewison, E.Norton, J.Agnew, E.Walker, J.Jones, M.Smees, M.Ramatschi and G.Dick, 2006:** "Humidity Evaluation using GPS, Radiometer, Wind Profiler and UV Lidar During CSIP Associated with Thunderstorms", Proceedings of COST720 Final Symposium, MeteoFrance, Toulouse, May 2006, Available from http://tim.hewison.org/COST720_proc_humidity.doc.
- Newton, I., 1669:** "Methodus fluxionum et serierum infinitarum", as cited by E.W. Weisstein, "Newton's Method." From MathWorld – A Wolfram Web Resource, <http://mathworld.wolfram.com/NewtonsMethod.html>, (accessed 7/9/06).
- Ofcom, 2004:** "UK Frequency Allocation Table 2004", <http://www.ofcom.org.uk/radiocomms/isu/ukfat/>, (accessed 21/6/06).
- Parrish, D. and J.C.Derber, 1992:** "The National Meteorological Center's spectral statistical interpolation analysis system", Monthly Weather Review, 120, pp1747-1763.
- Peckham, G.E. and M.Grippa, 2000:** "Improved retrieval of tropospheric temperatures from remote measurements of thermal radiation using the adiabatic lapse rate as a constraint", Q. J. Royal Meteorol. Soc., Vol.126, No.563, pp.749-760.
- Petersen, R.A. and W.R.Moninger, 2006:** "Assessing two different aircraft based sensing systems", 10th Symposium on Integrated Observing and Assimilation Systems for Atmosphere, Oceans and Land Surface, Jan 2006, Atlanta, GA, Amer. Meteor. Soc.
- Phalippou, L., 1996:** "Variational retrieval of humidity profile, wind speed and cloud liquid-water path with SSM/I: Potential for numerical weather prediction", Q. J. Royal Meteorol. Soc., Vol.122, No.530, pp.327-355.
- Price, J.D., 2004:** "Liquid water path and boundary layer cloud albedo bias calculations made at MRU Cardington, and comments on their parametrization", Cardington Technical Note No.63, Available from National Meteorological Library, Met Office, UK.

- Press, W.H., S.Teukolksy, W.T.Vetterling and B.Flannery, 1995:** "Numerical recipes: the art of scientific computing", Second edition, Cambridge University Press.
- Purser, R.J. and H.-L. Huang, 1993:** "Estimating Effective Data Density in a Satellite Retrieval or an Objective Analysis", J. Applied Meteorology: Vol. 32, No. 6, pp. 1092-1107.
- Rabier, F., N.Fourrje, D.Chafai and P.Prunet, 2002:** "Channel selection methods for Infrared Atmospheric Sounding Interferometer radiances", Q.J.R.Meteorol. Soc, 128, pp.1011-1027.
- Radiometrics, 2001:** "TP/WVP-3000 Temperature and Water Vapour Profiling Radiometer" Operators handbook, Radiometrics Corporation, Boulder, CO, USA
<http://radiometrics.com>.
- Rayer, P.J. 2001:** "Microwave transmittance models for RTTOV", EUMETSAT NWP SAF Deliverable 8.1, available from <http://www.metoffice.gov.uk/>.
- Rodgers, C.D., 2000:** "Inverse Methods for Atmospheric Sounding: Theory and Practice", World Scientific Publishing Co. Ltd.
- Rose, T., S.Crewell, U.Löhnert, C.Simmer, 2005:** "A network suitable microwave radiometer for operational monitoring of the cloudy atmosphere", Atmosph. Res., 75, 183-200.
- Rosenkranz, P.W., 1993:** "Absorption of Microwaves by Atmospheric Gases", in "Atmospheric remote sensing by microwave radiometry", Wiley Series in Remote Sensing, M. Janssen, Ed. Wiley.
- Rosenkranz, P.W., 1998:** "Water Vapour Microwave Continuum Absorption: A Comparison Of Measurements And Models", Radio Science, Vol.33, No.4, pp.919-928.
- Rowling, J. K., 1999:** "Harry Potter and the Chamber of Secrets", Bloomsbury, London.
- Ruffieux, D., J.Nash, P.Jeannet, J.L.Agnew: 2006:** "The COST 720 temperature, humidity, and cloud profiling campaign: TUC", Meteorologische Zeitschrift, Vol.15, No.1, pp.5-10.
- Sand, W., W.Cooper, M.Politovich and D.Veal, 1984:** "Icing conditions encountered by a research aircraft", J. Climate. and Appl. Meteor., 28, 856-868.
- Saunders, R.W., M.Matricardi and P.Brunel, 1999:** "A fast radiative transfer model for assimilation of satellite radiance observations - RTTOV5", ECMWF, Technical Memorandum No.282.
- Slonaker, R.L., B.E.Schwartz, W.J.Emery, 1996:** "Occurrence of Nonsurface Superadiabatic Lapse Rates within RAOB Data", Weather and Forecasting: Vol. 11, No. 3, pp. 350-359.
- Smith, W.L., W.F.Feltz, R.O.Knuteson, H.E.Revercomb, H.M.Woolf and H.B.Howell, 1999:** "The Retrieval of Planetary Boundary Layer Structure Using Ground-Based Infrared Spectral Radiance Measurements", J. Atmos. Oceanic Technol., Vol. 16, pp. 323-333.
- Smout, R. J.Nash, T.Hewison and M.Smees, 2005:** "Results of the RS92 Acceptance Test performed by the Met Office (UK)", WMO Technical Conference on Meteorological and Environmental Instruments and Methods of Observation (TECO-2005), Bucharest, Romania, 4-7 May 2005: [http://www.wmo.int/web/www/IMOP/publications/IOM-82-TECO_2005/Papers/2\(11\)_UK_6_Smout.pdf](http://www.wmo.int/web/www/IMOP/publications/IOM-82-TECO_2005/Papers/2(11)_UK_6_Smout.pdf).

- Solheim, F., J. Godwin and R.Ware, 1996:** "Microwave Radiometer for Passively and Remotely Measuring Atmospheric Temperature, Water Vapour, and Cloud Liquid Water Profiles", *Final Contract Report DAAL01-96-2009, White Sands Missile Range*. Available from <http://radiometrics.com/eigenvalue.pdf>.
- Solheim, F., J.Godwin, E.Westwater, Y.Han, S.Keihm, K.Marsh and R.Ware, 1998:** "Radiometric Profiling of Temperature, Water Vapor, and Cloud Liquid Water using Various Inversion Methods", *Radio Science* Vol.33, pp.393-404.
- Solheim, F.S., J.Vivekanandan, R.H.Ware, C.Rocken, 1999:** "Propagation delays induced in GPS signals by dry air, water vapour, hydrometeors, and other particulates", *J. Geophys. Res.*, Vol.104, No.D8, pp.9663-9670.
- Stringer, S., 2006:** "User Requirements for Upper Air Observations", *Observation Supply, Met Office Internal Report*, Available from Met Office, FitzRoy Road, Exeter.
- Thiebaut, H.J., 1973:** "Maximally Stable Estimation of Meteorological Parameters at Grid Points", *J. Atmosph. Sci.*, Vol.30, pp.1710-1714.
- Treuhart, R.N. and G.E.Lanyi, 1987:** "The effect of the dynamic wet troposphere on radio interferometric measurements", *Radio Science*, 22 (2): pp.251-265.
- Troitsky, A.V., K.P.Gaikovich, V.D.Gromov, E.N.Kadygrov and A.S.Kosov, 1993:** "Thermal sounding of the atmospheric boundary layer in the oxygen band center at 60 GHz", *IEEE Trans. Geosci. Remote Sens.*, Vol.31, No.1, pp.116-120.
- Ulaby, F.T., R.K.Moore and A.K.Fung, 1981:** "Microwave Remote Sensing - Active and Passive, Volume I: Microwave Remote Sensing Fundamentals and Radiometry", Reading, MA: Addison-Wesley.
- van Baelen, J., J-P.Aubagnac and A.Dabas, 2005:** "Comparison of Near-Real Time Estimates of Integrated Water Vapor Derived with GPS, Radiosondes and Microwave Radiometer", *J. Atmosph. Oceanic Techn.*, Vol.22, pp.201-210.
- van Meijgaard, E. and S.Crewell, 2004:** "Comparison of model predicted liquid water path with ground-based measurements during CLIWA-NET", *Atmos. Res.*, 75(3), pp.201-226.
- van Vleck, J.H. and V.F.Weisskopf, 1945:** "On the Shape of Collision-Broadened Lines", *Rev. Mod. Phys.*, 17, pp.227-236.
- Vasić, V., D.G.Feist, S.Müller and N.Kämpfer, 2005:** "An Airborne Radiometer for Stratospheric Water Vapor Measurements at 183 GHz", *IEEE Trans. Geosci. Remote Sensing*, Vol.43, No.7.
- Ware, R., F.Vandenberghe, J.Vivekanandan, and E.R.Westwater, 2004:** "Ground-Based Radiometric Profiling during Dynamic Weather Conditions", *Unpublished manuscript* available from <http://www.radiometrics.com/pubs.htm>.
- Westwater, E.R., Y.Han, M.D.Shupe, S.Y.Matrossov, 2001:** "Analysis of integrated cloud liquid and precipitable water vapor retrievals from microwave radiometers during the Surface Heat Budget of the Arctic Ocean project", *J. Geophys. Res.*, Vol.106, No.D23, p.32,019.

- Westwater, E.R., B.Stankov, D.Cimini, Y.Han, J.A.Shaw, B.M.Lesht and C.N.Long, 2003:** "Radiosonde Humidity Soundings and Microwave Radiometers during Nauru99", J. Atmos. Ocean. Technol., Vol.20, No.7, pp.953-971.
- Westwater, E.R., S.Crewell and C.Mätzler, 2005:** "Surface-based Microwave and Millimeter wave Radiometrics Remote Sensing of the Troposphere: a Tutorial", IEEE Geoscience and Remote Sensing Society Newsletter, March 2005.
- Wilks, D. S. , 1995:** "Statistical Methods in the Atmospheric Sciences", Academic Press.
- Wood, R. and J.P.Taylor, 2001:** "Liquid water path variability in unbroken marine stratocumulus cloud", Q. J. Roy. Meteorol. Soc., Vol.127, pp.2635-2662.
- Wu, R. and J.A.Weinman, 1984:** "Microwave radiances from precipitating clouds containing aspherical ice, combined phase and liquid hydrometeors", J. Geophys. Res., 89, 7170-8.
- Xie, S. and M. Zhang, 2000:** "Impact of the Convection Triggering Function on the Single-Column Model Simulations", J. Geophys. Res., Vol.105, D11, pp.14,983-14,996.
- Zell, A., N.Mache, R.Hubner, G.Mamier, M.Vogt, K.Herrmann, M.Schmalzl, T.Sommer, A.Hatzigeorgiou, S.Doring and D.Posselt, 1993:** "SNNS : Stuttgart Neural Network Simulator," Tech. Rep. 3/93, Institute for Parallel and Distributed High Performance Systems, University of Stuttgart, Germany.