Australian Government
Bureau of Meteorology

The Centre for Australian Weather and Climate Research
A partnership between CSIRO and the Bureau of Meteorology

CSIRO

Training the Ancilla Naïve Bayes Classification System to provide quality control for weather radar data

Susan J Rennie, Mark Curtis, Justin R Peter, Alan W Seed and Peter J Steinle

CAWCR Technical Report No. 073

June 2014

www.cawcr.gov.au

Training the Ancilla Naïve Bayes Classification System to provide quality control for weather radar data

Susan J Rennie, Mark Curtis, Justin R Peter, Alan W Seed and Peter J Steinle

Centre for Australian Weather and Climate Research, Australian Bureau of Meteorology

CAWCR Technical Report No. 73

June 2014

National Library of Australia Cataloguing-in-Publication Entry:)

Author: Susan J Rennie, Mark Curtis, Justin R Peter, Alan W Seed and Peter J Steinle

Title: Training the Ancilla Naïve Bayes classification system to provide quality control for weather radar data

ISBN: 9781486303250

Series: CAWCR Technical Report; No. 073.

Notes: Includes bibliographical references and index.

Subjects: Radar meteorology.

Radar cross sections.

Weather forecasting--Australia.

Dewey Number: 551.63530994

Enquiries should be addressed to:
Susan Rennie
Centre for Australian Weather and Climate Research:
Australian Bureau of Meteorology
GPO Box 1289, Melbourne
Victoria 3001, Australia

s.rennie@bom.gov.au

Copyright and Disclaimer

All images reproduced in grayscale. A colour version of CAWCR Technical Report No.073 is available online: http://www.cawcr.gov.au

Contents

List of Figures

List of Tables

1 INTRODUCTION

Weather radar is an important tool for monitoring and prediction of severe weather. Radar also provides broad spatial coverage for rainfall estimation that complements observations from the rain gauge network. However, the radar signal can be returned from a wide variety of targets, which gives erroneous rainfall estimations. An experienced human observer can usually filter out the extraneous radar echoes when interpreting radar scans, which is adequate for qualitative applications. This is facilitated by examining a time series of radar scans. However, for quantitative applications such as data assimilation or quantitative precipitation estimation, there needs to be an automated method of selecting the appropriate radar echoes that relies as little as possible on information external to an individual radar scan. This means that the method does not rely on the availability of e.g. previous scans, other observation types, or numerical weather prediction (NWP) output that may not be reliably available.

The principal applications for radar data include:

Quantitative precipitation estimation (QPE): the estimation of rainfall over an area.

Nowcasting: the advection of a precipitation field to provide up to 2–3 hour rainfall predictions;

Data assimilation: the use of reflectivity or precipitation estimates and Doppler radial wind velocity to improve the initialisation of NWP models.

Mostly these use only precipitation echo, although radial wind estimation can be achieved with some types of clear air echo. In all of these applications, the presence of non-precipitation echo can damage the result, for example producing gross errors in a forecast.

The types of echo typically detected by weather radar include:

- precipitation,
- aerofauna that fly at sufficient heights and in sufficient numbers (primarily insects, birds and bats),
- smoke from large bushfires,
- chaff (metallic particles) released from military aircraft,
- permanent ground echo
- anomalous propagation echo from the ground,
- anomalous propagation echo from the ocean,
- echo from side-lobes, primarily from the ocean surface,
- second-trip echo,
- echoes from individual mobile objects including aerials, ships, aeroplanes, including interfering signals.

These types of echo occur with varying frequency, and have different properties. Several of these fall under the heading of clear air echo, i.e. returns from the atmosphere when no precipitation is present. Aerofauna can have echo strength comparable to light precipitation, and an independent velocity. Insect air velocity may or may not be negligible (Achtemeier 1991; Geerts and Miao 2005), but birds and bats are strong fliers (Wilson *et al.* 1994; Koistinen 2000), and any observed 'wind' velocity will be a combination of the true wind velocity and the aerofauna air velocity. Insect echo is particularly prevalent in the summer months in Australia, during both night and day (Rennie 2013). Bird echoes are mostly seen at dawn or dusk (Rennie 2012). Smoke appears sporadically, but large fires such as the Black Saturday bushfires in Melbourne on 7 February 2009 can produce smoke at comparable signal strength and height as precipitation. Chaff is sporadically seen off the coast, for a short period, and may have a large initial velocity from when exiting the aircraft (Rennie 2012). Observation of many chaff examples suggests that chaff is only useful as a passive tracer after a long dispersal period (unpublished data).

Surface echo is another major type of non-precipitation echo. Ground echoes can cause serious contamination, particularly if there is no Doppler filtering of the zero radial velocity. Permanent echoes are most easily identified because of their consistent position. Anomalous propagation echo from the ground or sea is harder to detect, and may have an appearance similar to rain. Anomalous propagation (AP) occurs when the radar beam undergoes non-standard propagation through the atmosphere that results in it being bent back to the surface. Some Australian regions are more prone to AP because of the frequency of atmospheric conditions that permit ducting. A study of ducting conditions near Sydney is detailed in Peter *et al.* (2014).

Second-trip echo occurs when the returned signal arrives after the next pulse has been emitted, so the echo source appears much closer to the radar than it actually is. It is relatively rare and usually appears when there is very heavy rainfall far from the radar. Finally, echoes from individual targets are common, but generally of short duration and extent. These may be more easily removed with e.g. a speckle filter, than by attempting to classify them.

The ideal radar classification method would distinguish all types of echoes. Such a comprehensive echo classification has been demonstrated for dual-polarisation radars; various studies have shown that polarisation diversity is very useful for various types of classification frameworks, including fuzzy logic and Bayesian methods (Rico-Ramirez and Cluckie 2008 and references therein). There are many examples using dual polarisation to identify different hydrometeor types as well as echo types (e.g. Schuur *et al.* 2003; Dixon *et al.* 2005; Bachmann and Zrnić 2008; Koistinen *et al.* 2009). Dual-polarisation radars are some years away for the Australian radar network, so single polarisation methods must be used. There are examples of single-polarization classifiers, primarily focused on discriminating surface (usually AP) echoes from precipitation (e.g. Moszkowicz *et al.* 1994; Peter *et al.* 2013). Discrimination of precipitation and non-precipitation echoes using a neural network approach has been demonstrated with single-polarization radar (Lakshmanan *et al.* 2007). Fortunately, many applications of rainfall estimation and prediction only require distinguishing precipitation echo from other echoes. For Doppler radial velocity data assimilation, distinguishing echoes suitable for wind velocity estimation is the main requirement.

The objective of the radar classification method described in this report is to provide means to select only the desired echo type for the particular application. There is particular focus on Doppler radar usage because the assimilation of radial velocity into NWP models is a high priority. Precipitation is the preferred target for wind estimation because its independent velocity (non-wind driven) can be estimated as its fall velocity. However, some clear air echoes may be useful for wind estimation, where it has a negligible independent velocity (smoke and smaller insects). Clear air echo can provide a substantial source of wind observations during fine weather. Therefore detection of clear air echo suitable for wind estimation is a secondary goal to the discrimination of precipitation echo. However for data assimilation and NWP, defining the wind field between areas of rain is also important. This report documents the development of a Naïve Bayes classifier, expanding on the version developed by Peter *et al.* (2014), which aims to discriminate many types of echo. The feature fields available include reflectivity, textures and gradients, echo top height, and Doppler parameters. The naïve version means that the feature fields used by the classifier are assumed to be independent. This is not strictly the case, but reasonable results are expected to be possible (Friedman *et al.* 1997; Peter *et al.* 2014). The naïve version is much simpler to calculate and faster to implement, which is preferable in an operational context. The new classifier encompasses many echo types, uses Doppler information in the feature fields, and a wide range of probability distribution functions (pdfs) are available to describe the feature fields. Thus the most accurate representation of pdfs can be used to characterise all typical echo types.

2 THE ANCILLA CLASSIFIER

The classification algorithm forms part of the Ancilla software, which is developed within CAWCR to handle the radar data flow, from the radar site to the various outputs. The scope of Ancilla is much greater than will be discussed here. Ancilla provides several tools for classification.

- The creation of feature fields.
- The aggregation of values for the training dataset
- The definition of various pdfs
- The Bayes classifier
- Probability of detection map calculation

Additionally, there is a viewer which allows the user to manually classify radar volumes in order to create the training dataset. The training dataset is described in section 3.

2.1 Feature fields

The radar volumes include up to four fields transmitted from the radar. The first is reflectivity, which is available from all radars. This is reflectivity corrected for range-dependent noise, filtered to remove permanent echo (zero velocity filtering) and thresholded using the signal quality index (SQI). Secondly, Doppler radars may also return the radial velocity, which has been independently quality-controlled. Some Doppler radars additionally return spectrum width (the third moment) and uncorrected reflectivity. Information on the radar signal processing can be found in the radar documentation (Australian Bureau of Meteorology).

The dual Pulse Repetition Frequency (PRF) Doppler radars operate with alternating PRFs for each azimuth, with dual-PRF unfolding applied by operating across adjacent azimuths. There is a small difference in the value of the spectrum width that is dependent on the PRF. Therefore, the spectrum width is spatially averaged over adjacent azimuths before use.

From the first, second and third moments of the radar measurements, various feature fields can be derived. These are described in detail in Peter *et al.* (2014) The texture (Kessinger *et al.* 2004; Hubbert *et al.* 2009) and spin (Steiner and Smith 2002) are available in one or two dimensions. Texture (T) is defined as

$$T = \left[\sum_{j}^{N} \sum_{i}^{M} \left(X_{i,j} - X_{i-1,j} \right)^2 \right] \Big/ (N \times M) \qquad (2.1)$$

where X is a moment of the radar signal, typically reflectivity, and measures the difference between adjacent bins (radar observation of specific range and azimuth) within a kernel of $N \times M$ (in practice, we generally set $N = M$).

Spin is defined as a measure of the change in sign of the difference between adjacent bins within a kernel. The valid measurable spin fulfils the following conditions, for successive bins X_{i-1}, X_i and X_{i+1}.

$$sign\{X_i - X_{i-1}\} = -\, sign\{X_{i+1} - X_i\} \tag{2.2}$$

$$\frac{|X_i - X_{i-1}| + |X_{i+1} - X_i|}{2} > spin\ threshold \tag{2.3}$$

Where the spin threshold for reflectivity is e.g. 2 dB.

The vertical gradient of reflectivity is a measure of the difference between bins of the same range and azimuth, at adjacent elevations. This is applied to a smoothed field of reflectivity, after a Gaussian kernel filter has been applied, to compensate for the discontinuous values of reflectivity (see Section 3.2.2). It has no value at the highest elevation.

The echo top height is the height at which the vertical profile of reflectivity drops below some threshold. The vertical profile is determined as the value at a specific range and azimuth for all elevations. The height is the beam centre height assuming a standard $^4/_3$ earth approximation.

The feature fields considered for classification are as follows.

- Reflectivity (DBZH)
- Spectrum width (smoothed) (WAVG)
- 2D texture of reflectivity, with a kernel of 11 (TEX2D)
- 2D texture of radial velocity, with a kernel of 15 (VTEX2D)
- 2D spin of reflectivity with a kernel of 19, and threshold of 3 dB (SPIN2D)
- Vertical gradient (difference between levels) of smoothed reflectivity (VTDL)
- Echo top height with a threshold of 4 dB (ETH)
- Echo top height with a threshold of −5 dB (ETH2)

These were selected to be as independent a range of feature fields as possible. The texture and spin parameters are similar, but the difference in kernel size reduces correlations. Echo top height is correlated, but independence can be achieved by only using one threshold at any location, since for the higher threshold a value might not exist if the reflectivity never exceeds the threshold.

2.2 The probability distribution functions

Rather than use the Box-Cox transformation and make a Gaussian fit to the data, a range of pdfs are available to be empirically fit to the histogram of raw data. Many of these are standard pdfs. Full details are in Appendix B. All pdfs are continuous distributions. Their domain may be within all real numbers, or all positive numbers (some including zero). A transposition parameter is included to shift the pdf to a different region, e.g. to shift the lowest valid pdf value to just below zero if zero is not normally defined for the pdf. Since the fits are to be empirical, the coverage only needs to look reasonable.

The pdfs are summarised below.

- The uniform distribution, constant between two values.
- Linear distribution that increases from 0 to maximum between two values.
- Triangular, that increases linearly from 0 to the maximum, then decreases linearly.
- Trapezoidal distribution, that linearly increases to a plateau.
- Normal (Gaussian) distribution.
- Inverse normal distribution
- Log-normal distribution
- Skew-normal distribution
- Box-Cox normal distribution (not used)
- Truncated normal distribution
- Exponential distribution
- Gamma distribution
- Inverse gamma distribution
- Laplace distribution
- Laplace-normal distribution, a composite pdf with the Laplace distribution centred on 0 and combined with normal distributions centred at ± the mean. This can be set to exist in only the positive domain.
- Laplace-Laplace distribution, a composite of two co-located Laplace distributions.
- Laplace-skew-normal distribution, a composite of a Laplace distribution and a skew-normal distribution.
- Log-binormal distribution, a composite of two log-normal distributions.

2.3 Function of Ancilla to perform the classification

The relevant Ancilla functions are controlled by xml scripts which specify the functions and parameters. Details about these various functions and parameters are described in Sections 3 and 4. Here only the functionality of Ancilla is described.

Each feature field is generated by a specified function of radar moments and parameters. For example, echo top height is a function of reflectivity and a particular threshold. The texture and spin feature fields are created using the specified field, kernel size and if relevant, a threshold.

To run the classification algorithm, first the prior probabilities are specified. Climatology files for each radar are provided to be used for specialised prior probabilities. Some classes are dependent on location over land or sea, and the prior probability is a function of the distance from the coast. Permanent echoes have a prior probability that is a function of the probability of detection map. Otherwise prior probabilities are constant.

The pdfs are defined to be valid only over certain ranges, generally ignoring the extreme values of feature fields. Therefore any extreme values are set to Not a Number (NaN) so that the feature field is not used for that bin.

For each class and feature field, the type of pdf and its parameters are specified. These are used for the Bayesian classification algorithm.

The actual values of prior probabilities and pdf parameters will be discussed in Sections 3 and 4.

3 METHODS FOR TRAINING THE CLASSIFICATION SYSTEM

3.1 Manual classification of training data

The Naïve Bayes Classifier (NBC) requires a training dataset to create the pdfs for each feature field and each class. The training dataset was created by manually classifying a large number of volumes from a variety of radars. It is not feasible building up a sufficient database to capture the climatology of all classes due to the time-consuming and laborious nature of manual classification. However, examples were chosen from a range of cases, and the resultant pdfs were recognised to be smooth. The resultant pdfs were either sufficient so that new data had little impact, or that any irregularities in the pdf due to sparse sampling (problematic for echo top height for rare classes) were compensated for when fitting the pdf.

Twelve classes were selected for manual classification, which describe the types of echoes that are normally detected by the radars. The classes used are listed in Table 1 and detailed below. Examples of these are shown in Fig 1 through Fig.7. Identification of echo types by an expert was aided through the use of time sequences of images, combined with access to the various feature fields and the ability to compare scans for different elevations. Additionally, knowledge of the type of event e.g. weather forecasts or news reports of fires. In general a trained observer would have no difficulty identifying echo types provided that a suitable time series is available.

Precipitation
There were three precipitation classes: stratiform, convective and shallow convective. These were classified as separate classes in order to help represent the variability of a precipitation field. Convective precipitation is more spatially variable than stratiform, and shallow convective tends to have even smaller spatial variability as well as being height-limited. The discrimination was user subjective, and the NBC is not intended to be used to discriminate these types, as there are other methods that are sufficiently accurate and already in use (e.g. Steiner et al. 1995). Shallow convective is most similar to other classes and is therefore most difficult to classify.

Chaff
Chaff usually appears only over the ocean, off the central east coast (Sydney to Brisbane) and near Perth. However, chaff was also released inland north of Namoi. Generally this is related to the proximity to air force bases used for training exercises. Chaff classification is limited to the radars at which it has been seen. Chaff has been simultaneously detected by both C and S band radars, and appeared to be less easily detected by the C band radars. Radar calibration and sensitivity may be a factor here and this should be revisited periodically. At present, chaff at different radars is not discriminated by band.

Insects
Insects are recognised by their typically low reflectivity, low altitude and relatively homogeneous distribution. Their detectability varies according to location and radar sensitivity, and likely other difference among the radars in the heterogeneous network that the Bureau operates. For example, Terrey Hills detects far more insects than Wollongong despite their

proximity. However, insect reflectivity can range up to 40 dB, and texture can be very smooth, or rough if convergence lines (e.g. around convective cells) cause aggregation.

Birds and bats – large aerofauna

Birds and bats are both likely to appear during dusk or dawn dispersal after roosting, and are not distinguishable by the radar. Hence the phenomenon is referred to as 'birds' throughout this document. Birds are not detected at all radars, but if present are usually seen at dawn or dusk in isolated patches, although Melbourne has a corridor of migration across Port Phillip Bay. Often they appear as a patch of velocity running counter to the wind, visible for a few scans only.

Smoke

Smoke is generally rare, though during the dry or summer seasons bushfires can occur relatively often. Smoke observations are available for many radars. However, large fires where the smoke extends above the convective boundary layer are very rare, so only one example (Black Saturday 2009) was used. The reflectivity and echo top height thus can vary as widely as precipitation.

Ground clutter

All radar locations suffer from permanent ground clutter echo to some degree, depending on the surrounding orography. The permanent echo has been characterised with probability of detection (POD) map. Much of the ground clutter is filtered by on-site processing by zero-isodop filtering, resulting in 'holes' where the POD is almost zero. However, there is a 'halo' surrounding these holes where the POD is relatively high, e.g. > 0.4 for some reflectivity threshold, e.g. ~0 dB. In some cases, the filtering is consistently poor and the POD is near 1. These halos are used for this class. AP ground clutter is more closely related to the occurrence of super-refraction and ducting, so is very common in northwest WA and sporadic in other locations. The onsite processing removes some of the AP ground clutter. The AP echo is expected to resemble the permanent echo in some respects, such as texture and vertical gradient.

Sea clutter

Sea clutter similarly depends on variability of atmospheric refraction. There are two type of sea clutter, which are characterised separately. Side-lobe sea clutter, which appears close to the radar and is more uniform in appearance, is common on Sydney radars, but has been observed at many coastal radars. It appears on the POD maps with probability ranging up to 0.6, and is not filtered since it does not have zero velocity. In contrast, AP sea clutter has a less uniform texture and so is more difficult to distinguish from precipitation. It is very common on the northwest of WA, but is also often seen off Sydney, as well as other coastal radars.

2nd trip echo

Second trip echo is observed rarely, but most often is seen at Terrey Hills, where the 2nd trip echo is used to increase the range of detection. Note that second trip echo in cases where the echo is from beyond the unambiguous range appears shallow and with near-zero velocity (unpublished data). Ultimately, this echo type was ignored during classification due to its rarity, to limit the number of classes the NBC works with.

Not included in the manual classification are examples of multi-body echo. For example, storm echo that is reflected off a building then to the radar, or delayed return from a hailstorm. Such cases were ignored as they are very rare, and the Bayes method is unlikely to recognise echo with low prior probabilities if its characteristics are indistinguishable. Singular targets such as aircraft are also ignored. With single polarisation radars, there are too few distinguishing characteristics for an exhaustive selection of classes.

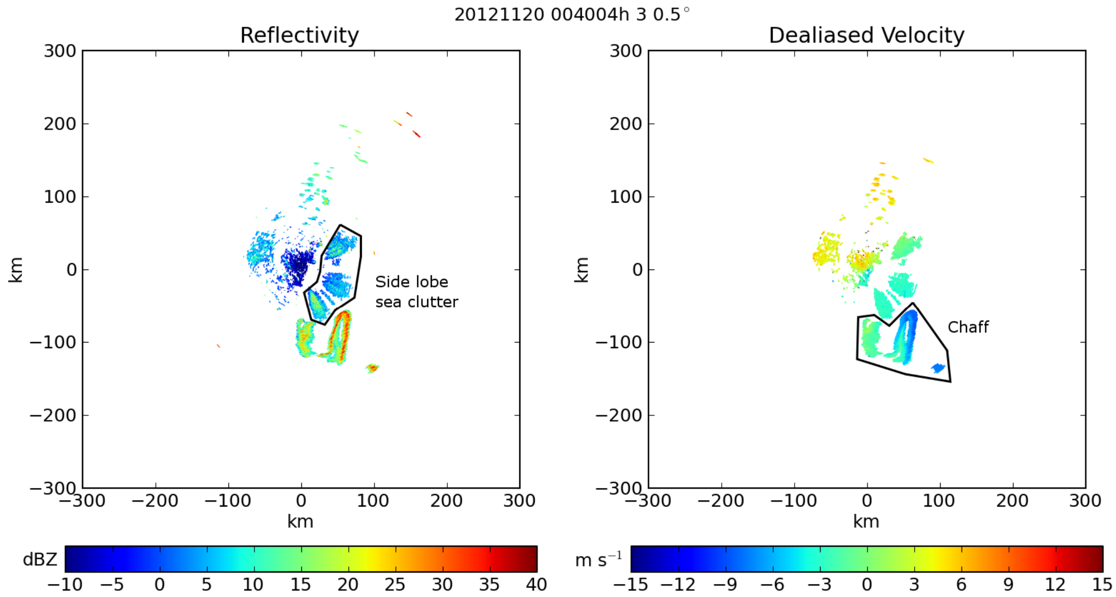

Fig. 1 Examples of side lobe sea clutter and chaff, in reflectivity and velocity, taken from 20 November 2012 at 0040 UTC from the Wollongong 0.5° elevation scan.

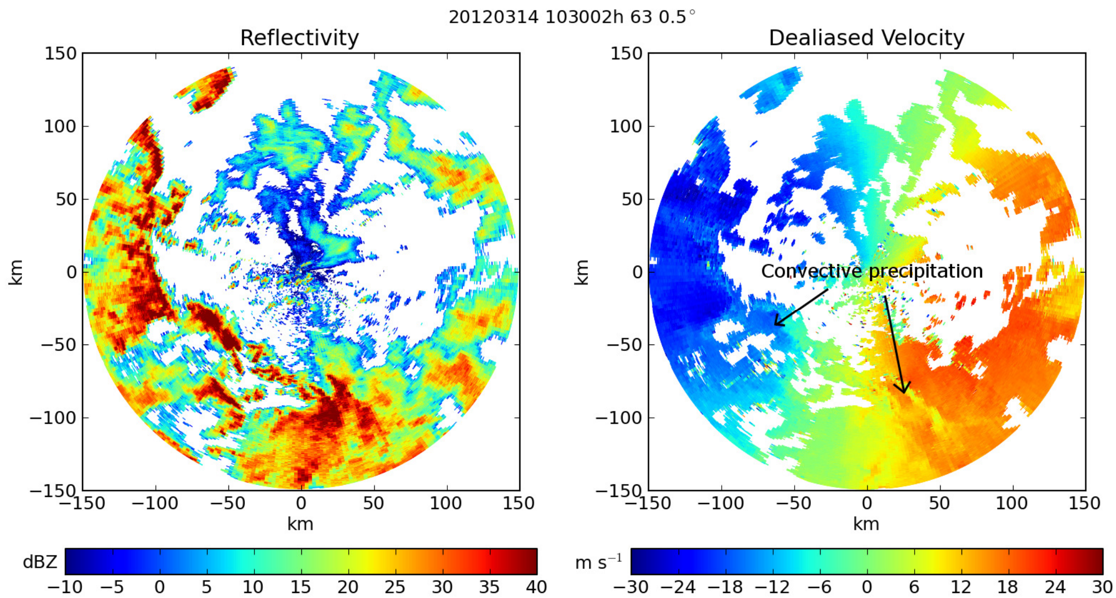

Fig. 2 Example of convective precipitation, in reflectivity and velocity, taken from 14 March 2012 at 1030 UTC from the Darwin (Berrimah) 0.5° elevation scan.

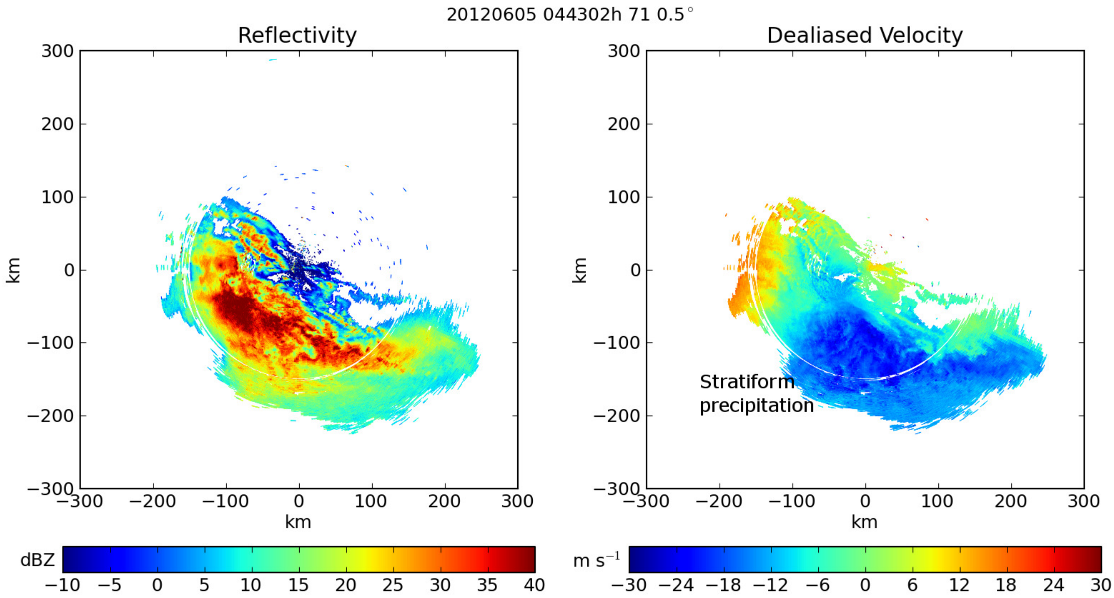

Fig. 3 Example of stratiform precipitation, in reflectivity and velocity, taken from 5 June 2012 at 0443 UTC from the Sydney (Terrey Hills) 0.5° elevation scan.

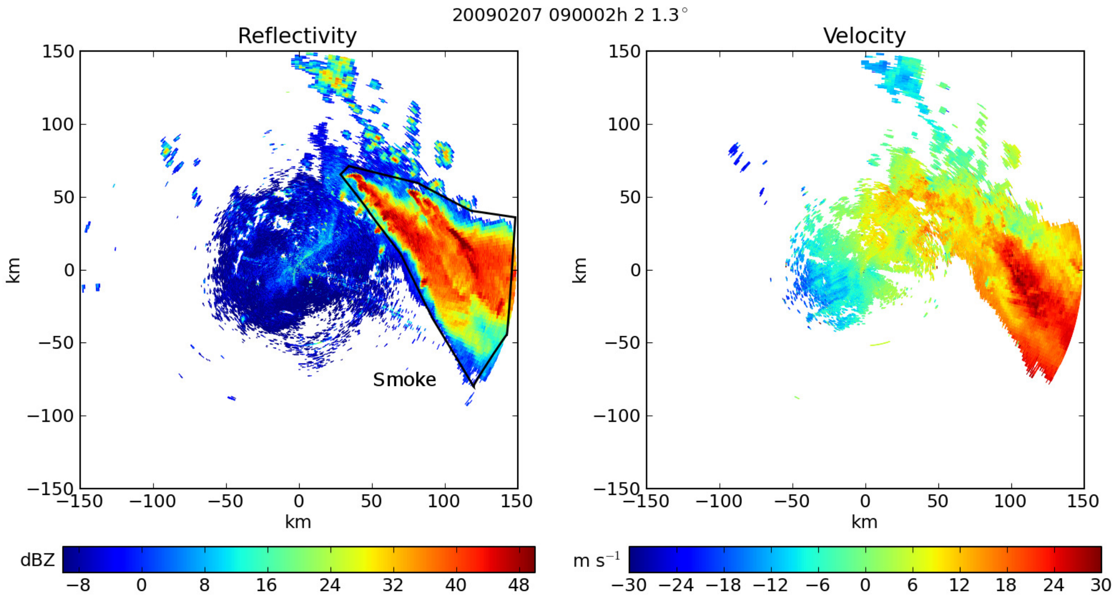

Fig. 4 Example of smoke from the Black Saturday bush fires, in reflectivity and velocity, taken from 7 February 2009 at 0900 UTC from the Melbourne 1.3° elevation scan.

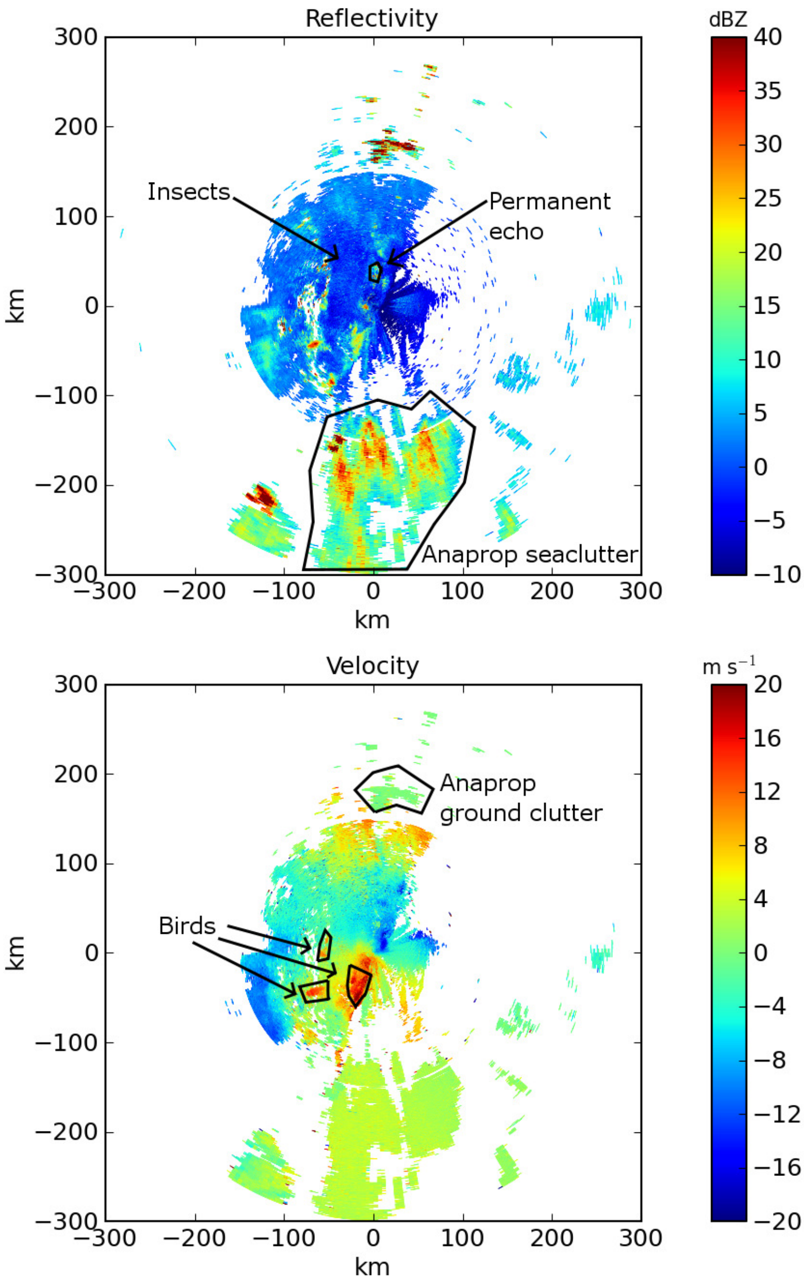

Fig. 5 Example of Insects, birds, permanent ground echo and AP ground clutter, in reflectivity and velocity, taken from 3 January 2012 at 0949 UTC from the Sydney (Terrey Hills) 0.5° elevation scan.

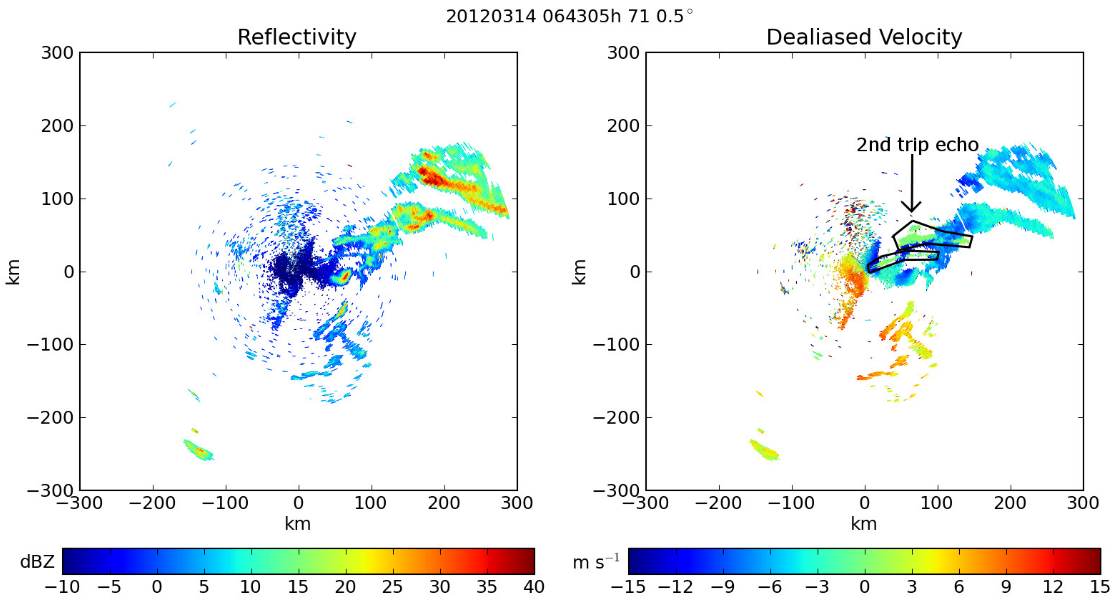

Fig. 6 Example of 2nd trip echo, in reflectivity and velocity, taken from 14 March 2012 at 0640 UTC from the Sydney (Terrey Hills) 0.5° elevation scan.

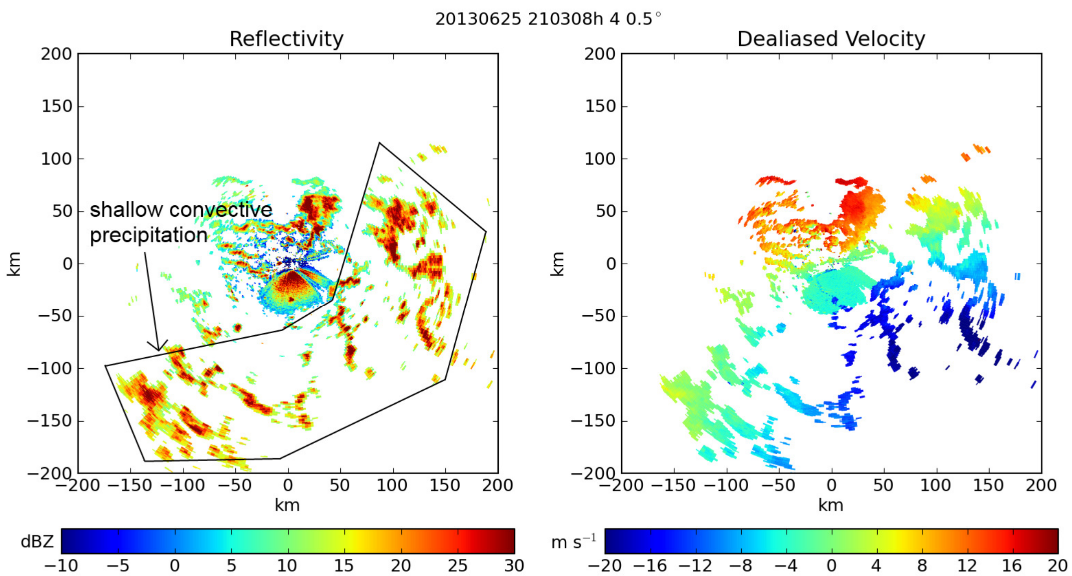

Fig. 7 Example of shallow convective precipitation echo, in reflectivity and velocity, taken from 25 June 2013 at 2103 UTC from the Newcastle (Lemon Tree) 0.5° elevation scan.

Table 1 List of the classes used for manual classification: Class number, class name, long and short class abbreviation as used for figure labels, etc.

Class number	Class	Class abbrev	Short abbrev
1	Convective precipitation	conv	con
2	Shallow convective showers	sh con	shc
3	Stratiform precipitation	strat	str
4	Insects	insects	ins
5	Smoke (bushfires)	smoke	smk
6	Chaff	chaff	chf
7	Birds	birds	brd
8	Permanent ground clutter	perm echo	pe
9	AP ground clutter	AP GC	gc
10	AP sea clutter	AP seaclutter	ap
11	Sidelobe sea clutter	SL seaclutter	sl
12	Second trip echo	2^{nd} trip	2tp

The number of examples from each radar and for each class is shown in Table 2. Table 3 summarises the number of pixels, volumes and days that contributed to the training dataset for each class. The radars' locations are shown in Fig. 8. More details on the training dataset are found in Appendix A.

Fig. 8 Map of the location of the radars used in the Strategic Radar Enhancement Project. The radar identification numbers mark the radar location, with the name of the radar location nearby.

Table 2 The number of volumes from each radar used to build the training data set for each class. Radars for which there was Doppler information are in shaded rows.

Radar	con	shc	str	ins	smk	chf	brd	pe	gc	ap	sl	2tp
2	2	0	3	1	3	0	4	7	2	0	0	0
3	11	1	7	4	6	18	6	2	0	6	20	0
4	0	1	0	0	0	3	0	0	0	0	0	0
8	1	0	0	1	0	0	0	2	0	0	0	0
14	0	0	1	0	0	0	0	0	0	0	0	0
16	0	0	0	0	0	0	0	1	7	6	7	0
19	1	0	0	0	0	0	0	0	0	0	0	0
28	0	0	0	0	0	2	0	0	0	0	0	0
29	0	0	0	0	0	0	0	1	2	2	5	0
49	1	0	3	8	4	0	1	3	0	0	0	0
50	1	0	0	0	0	2	0	0	0	0	0	0
52	0	0	1	0	0	0	0	0	0	0	0	0
54	0	1	0	2	6	5	0	6	2	2	8	0
63	2	0	0	0	0	0	0	0	0	0	1	0
64	0	0	0	1	0	0	0	2	2	1	4	0
66	0	0	0	1	0	3	0	0	1	0	0	0
68	0	2	0	0	0	0	0	0	0	0	0	0
69	2	0	0	9	0	0	0	0	0	0	0	0
70	0	2	0	3	0	3	0	3	0	0	0	0
71	3	1	8	19	6	8	8	17	5	3	18	5
72	4	0	0	6	0	0	1	3	0	0	0	0
75	0	0	0	6	7	0	2	0	0	0	0	0
76	0	1	2	0	0	0	0	5	0	0	1	0

Around 200 radar volumes were manually classified, of which ~160 were from Doppler radars. These were from a range of radars including C band and S band, different beam widths and range resolutions, and one with spectral width. Emphasis was given to Sydney radars since this is the first area for which Doppler wind assimilation will be tested. A range of classes was classified at each radar, depending on what types of echo are seen at that radar. Not all classes were classified at every radar; however since Wollongong was the only radar that provided spectral width at the time, all classes were classified there excluding 2nd trip (which was not found) and AP GC (not found but whose spectrum width should resemble perm echo). During development of the classification scheme, the training set was expanded if any dependencies were found, to try and improve the representativeness of the pdfs for each class. For example, if there are too few examples, the pdfs are highly dependent on those examples and the radars they came from.

In order to have a sufficiently large dataset, it is assumed that all radars will have similar properties regardless of band or beamwidth or range resolution. These assumptions are tested later. The files used for classification are summarized in Appendix A.

Data were classified using the following guidelines:

1. Only classify if identification is certain.

2. Try to avoid biased classification (e.g. only classifying high reflectivity values).

3. Avoid classifying areas where velocity was not correctly dealiased.

4. Minimise classification of regions where DBZH = −30. These areas are likely to skew the feature field statistics and will not be used, therefore need not be classified. In practice, data from these areas should be removed before use, because they don't contribute to the precipitation total and the velocity estimate from such a weak signal may be suspect.

Table 3 Summary of classification totals for each class. The number of pixels, radar volumes, radars and days (or events) that were classed for each classification type.

class	pixels	volumes	radars	days
conv	2853025	28	10	18
shconv	503032	9	7	9
strat	7638054	25	7	17
insects	8399289	61	12	29
smoke	645165	32	6	9
chaff	1326463	44	8	11
birds	49785	22	6	10
perm echo	145111	52	12	29
AP	373810	21	7	12
AP seaclutter	366761	20	6	9
SL seaclutter	735581	64	8	30
2^{nd} trip	19414	5	1	2

3.2 Aggregation and histogram creation

The training dataset was processed as follows. Using Ancilla, the feature fields described in Section 0 were created for each manually classified file. The values of the feature fields for each class were aggregated in a text file. The aggregated data were used to create a normalised histogram for each class and feature field. A range of pdfs were fit to the histogram, in order to find the most representative pdf, which is explained in detail in Section 3.3.

The radar data are stored in discrete intervals, and the feature fields may not be continuous. This can cause some artefacts to appear when creating histograms. Some processing was applied to the histograms to improve the fit of the pdfs, and remove artefacts. This prevented over-fitting of the continuous pdfs to artefacts in the histograms like the spikes described in Section 3.2.2.

3.2.1 Histogram creation

A normalised histogram is created using Python's NumPy histogram routine and specifying the number of bins. The bin centres are calculated and used as the *x*-values for the pdf.

Selecting the number of bins to produce a histogram that is smooth but does not lose information is important. The bin width is determined following Izenmann (1991), i.e.

$$W = 2\,(IQR)\,N^{-1/3} \tag{3.1}$$

where W is the bin width, IQR is the inter-quartile range and N is the number of data points. Thus the number of bins is the difference between the minimum and maximum values of the data divided by the bin width, and rounded up to the nearest integer.

In the case of some feature fields, using the recommended number of bins produced too much detail, or emphasised artefacts. Using a coarser histogram, a smoother result was achieved. For SPIN2D the number of bins is the maximum value of SPIN2D times 240. For echo top height, the histograms tended to have multiple peaks associated with the height of the scan elevations. The bin calculation above produced a too-detailed histogram for large sample sizes, so the number of bins was limited depending on sample size and variable to 250 (ETH > 40000 data) or 500 (ETH2, > 75000 data).

3.2.2 Histogram smoothing

Some feature fields, such as reflectivity and functions of reflectivity, have discrete values. Reflectivity, for example, has discrete values with varying, radar-dependent resolution. The typical values for reflectivity for one of the Doppler radars are [−30, −9, −7.5, −6, −4.5, −3, −1.5, 0, 1, 2, …, 10, 10.5, 11, …, 60, 60.5, …]. Most other radars have resolution half this, though a few have very coarse resolution (3–10 dB intervals). When creating a histogram, artefacts such as empty bins or spikes appear as a result of the irregular discrete values. This will reduce the accuracy of the pdf fit. The histogram bin size must be sufficiently small to capture the features of the histogram, and to ensure that the pdf follows the shape of the histogram without permitting features smaller than the bin width.

To alleviate the artefact problem, some feature fields are dithered to form a continuous distribution. The dithering model treats the discrete values as rounded values, because the histogram bins should be centred on the discrete values. It assumes that 'raw' non-rounded values were evenly distributed about the rounded values, and so half the raw values were rounded up from the midway point between discrete values, and half were rounded down. Therefore to reverse the rounding, a random number scaled to the range between adjacent midway points is added. For example, for reflectivity and a rounded value of −6, these midway points would be −6.75 and −5.25, so a random value in the range of ±0.75 is added.

The discrete, 'rounded' values are taken to be the unique values in the dataset. However, the rounded values may vary for different radars, but once the data are collated, there is no indication which radar they came from. Using unique rounded values from the whole dataset would result in poor dithering of values from radars with coarse resolution. Therefore the data are subdivided before dithering. A subsample of e.g. a thousand contiguous data values will likely come from the same radar, and the dithering process will work correctly. Dithering was applied to DBZH and WAVG.

The dithering algorithm is as follows:

1. Take a subsample of data

2. Determine the sorted unique values of the discrete variable.

3. Find the difference between them, and calculate the range over which a random value (the pad) should be added (within half the difference on each side).

4. For instances of each value, add either a random amount over the lower pad or the upper pad. For the lower limit of the feature field, only add an upper pad.

The result is not perfect, but gives a smoother histogram. For reflectivity, artefacts occur where the resolution changes, around 0 and 10 dB.

The histograms for data before and after dithering are shown in Fig. 9.

Finally, to remove any remaining artefacts in the histograms, which normally manifest as spikes, a pdf (from a selection of pdfs previously found to be representative of the feature field) is fit, and any histogram values that differ greatly from the pdf value at the same bin are replaced by the pdf value. For example, this is useful for DBZH where the resolution changes cause a histogram artefact (Fig.9). It was also necessary for SPIN2D histograms.

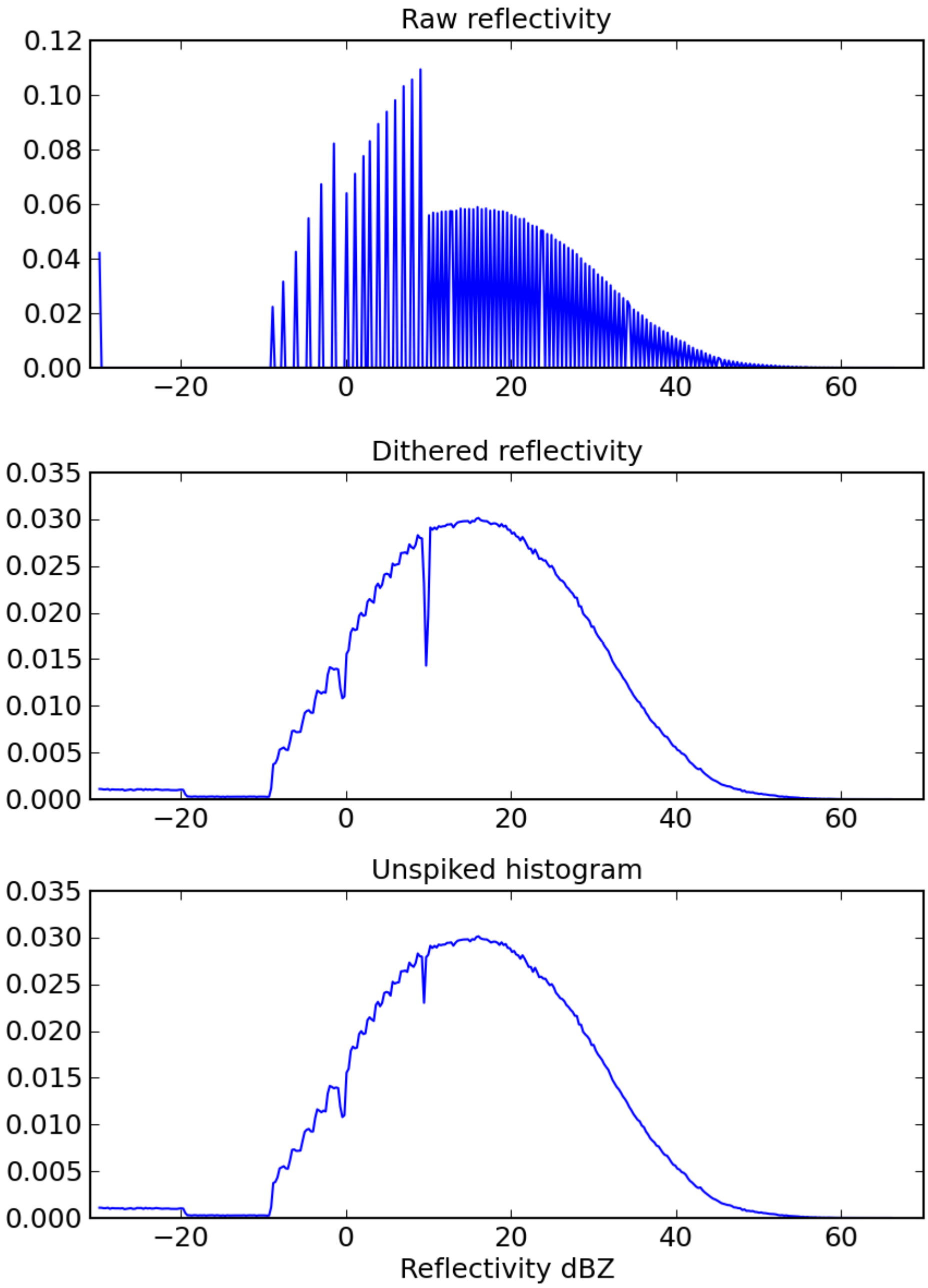

Fig. 9 Example of dithering of reflectivity. The top panel is the normalised histogram of original reflectivity values, the middle panel is the histogram of dithered reflectivity values, and the bottom panel is the histogram with spike removed.

3.2.3 VTDL Spike removal

VTDL has a large number of zero values from where the reflectivity is identical across two scans. Many of these occur where the reflectivity is −30; most otherwise occur at low reflectivity, probably due to the coarser reflectivity resolution. This means that regions where VTDL = 0 may not be clearly classified using this feature. The size of the zero spike in a histogram of VTDL depends on how much low reflectivity was classified, and is not assumed to reflect the climatology. The problem is partly alleviated by using the smoothed reflectivity to calculate the gradient; any remaining artefact should be removed.

The preferred method of fitting a histogram to VTDL would be to ignore zero values. The pdf should then be a result of the distribution of non-zero differences. This creates a histogram resembling a skew normal or Laplace distribution, with mode typically <0. The problem values appear in the range between ±0.5 (i.e. the minimum difference in reflectivity) divided by the relevant difference in beam heights. An example is shown in Fig.10. Logically, one could interpolate across the zero values. However, this may substantially affect the normalization of the histogram. Therefore the normalisation is done after the interpolation.

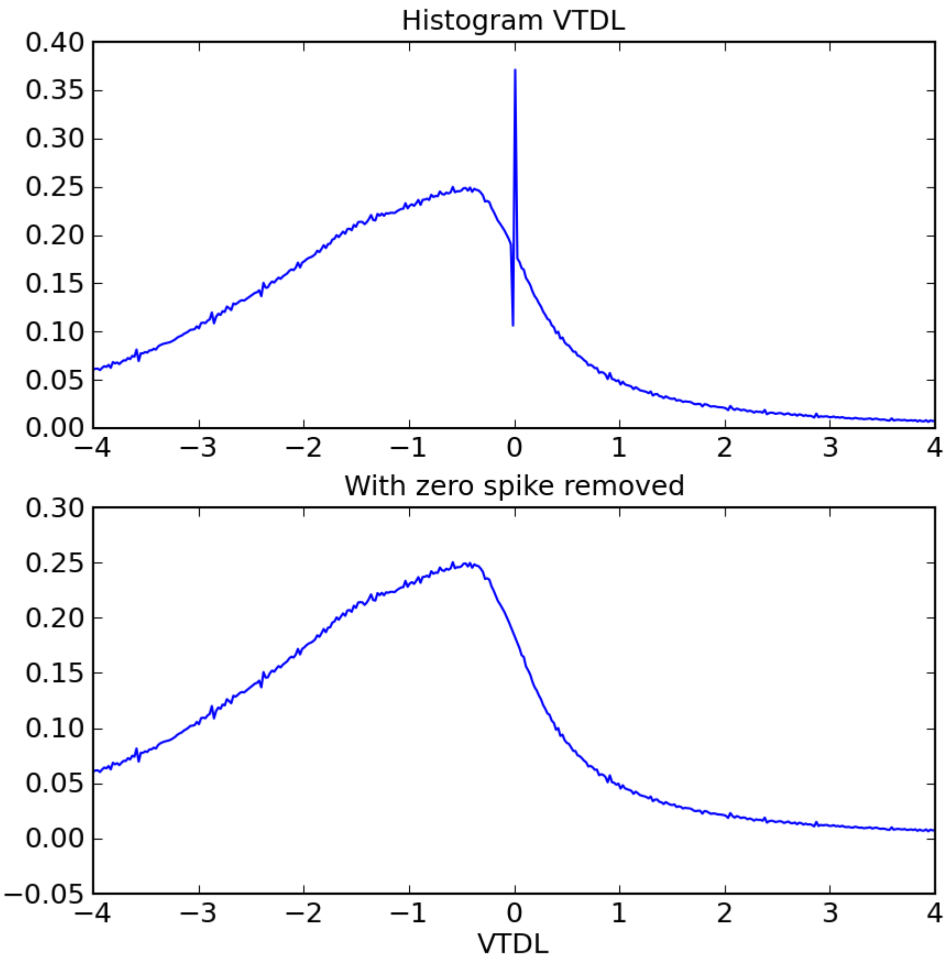

Fig. 10 Example of spike in the histogram at 0 VTDL, where the vertical reflectivity gradient is zero (top panel), and the histogram after the removal of the spike and interpolation across (bottom panel).

3.3 Fitting the pdfs

The pdf functions detailed in Section 2.2 and Table 4 were fit using Python. The SciPy statistics package supplied many pdfs and the rest were manually coded, e.g. the composite pdfs. The SciPy optimization package was used to perform a least-squares fit (using FITPACK) to find the parameters describing the best fit pdf.

The initial guesses were determined as per Table 4. A good first guess is important because the optimization package finds local minima only. Note that SciPy sometimes requires the pdf descriptors in a different format to the Ancilla program, and the SciPy descriptors are specified here. Generally SciPy parameters are normalised. Conversions between the two are trivial.

Table 4 Value of first guess for pdf parameters as defined in earlier section. These are usually a function of the data.

Pdf type	Guess parameters
normal	mean=mean, stddev=standard deviation (s.d.)
uniform	min=1^{st} percentile (pctl) and max=99^{th} pctl
linear	max, zero =1^{st} pctl and 99^{th} pctl, ordered by whether the mode is greater than the 50^{th} pctl value. This changes the slope
triangular	mid=mode, min=1^{st} pctl, max=99^{th} pctl (normalised for SciPy function)
trapezoidal	min=1^{st} pctl, loc1=20^{th} pctl, loc2=80^{th} pctl, max=99^{th} pctl
log normal	mean=exponent of mean of log zeroed data, stddev=s.d. (of log of data for SciPy), location =data minimum
skew normal	location=mean, scale=s.d. and shape=skew of data
truncated normal	min=data min or 1^{st} pctl if less than mean , max=99^{th} pctl (for SciPy transformed to the standard normal as (a-mean)/s.d.), mean=mean, stddev=s.d. of best normal fit to data
exponential	lambda=1
gamma	location = 0, shape=1.5, scale=mode
inverse gamma	location = 0, shape=1, scale=mode
Laplace	location=mode, scale=s.d./2
Laplace normal	scale=s.d./2, mean=mean, stddev=s.d./4.,mix=0.5
Laplace Laplace	location=mode, scale1= $(s.d.)^{½}$, scale2= $(s.d.)^{½}$/2,mix=0.5
Laplace skew normal	location=mode, scale= $(s.d.)^{½}$,mean=mean, stddev=s.d., shape=skew, mix=0.5
inverse normal	location = value where count is 1% of max counts (b0), mean=mean-b0, shape =$(mean−b0)^3$/variance
log binormal	mean1=mode(1^{st} half data), scale1=s.d./2, mean2=mode (2^{nd} data), scale2=s.d./2, location=data min, mix=0.8 (all of log, zeroed data)

The pdfs that were fit for each parameter are described in Table 5.

Table 5 The range of values or domains of the feature fields, and the distributions that are valid for that domain. The distributions likely to produce a reasonable fit are in bold

Parameter	Value range	Valid distributions
ETH ETH2	≥0	Uniform, linear, **triangular**, **inverse normal, log normal**, box cox, **truncated normal**, exponential, **gamma, inverse gamma, Laplace normal, log binormal, trapezoid**
SPIN2D TEX2D VTEX2D	≥0	Uniform, linear, triangular, **inverse normal, log normal**, **skew normal, truncated normal**, exponential, **gamma, inverse gamma, Laplace normal, log binormal,** trapezoid
DBZH	(−30, ∞)	Uniform, linear, triangular, normal, **inverse normal, skew normal**, **truncated normal**, Laplace, Laplace normal, Laplace Laplace, **Laplace skew normal, trapezoid**
WAVG	≥0	Uniform, linear, triangular, normal, **log normal, gamma, inverse gamma, Laplace normal, inverse normal, log binormal**
VTDL	(−∞,∞)	Uniform, linear, triangular, normal, **skew normal**, L**aplace**, **Laplace normal, Laplace Laplace, Laplace skew normal**

Only the most likely pdfs were fitted for the feature fields. For each histogram, the pdfs were fit, and the parameters and r.m.s. residuals were calculated. The best two (with smallest residuals) were noted. Then the histogram and best fits were plotted for closer inspection. The best fits were also plotted at higher resolution across a larger range to ensure a realistic representation across all feature field values. The best fit was selected to capture the characteristics of the pdf. In some cases manual modifications were made, e.g. to the truncated normal distribution, to ensure that the distribution covered the range of possible values for the feature field. Often the best fit was the pdf with the most degrees of freedom, i.e. the composite distributions like Laplace skew normal. If a simpler distribution provided almost as good a fit, that pdf was used instead, to improve computational efficiency.

The appropriate parameters for each selected fit were written out to be input to the Ancilla classification routine. Because the tails of the distributions less desirable for use in classification, extreme values of some feature fields were converted to NaN so that they would not be used by the classifier. These included DBZH<−25, ETH > 15 km, VTEX2D > 10, VTDL > 60, VTDL < −60, WAVG > 20. This masking of extreme values was conducted after the feature field values were calculated. Appendix B shows the raw histograms and the final pdfs for all the feature fields. Details of the feature fields are found in Section 4.

4 TUNING THE CLASSIFICATION SYSTEM

4.1 Choice of feature fields

For best results, it was resolved to use as many feature fields as possible. This means reflectivity and reflectivity-derived fields for all radars, and velocity and spectrum width fields for Doppler radars. The fitted pdfs and the overlaps between each are shown for each feature field, to demonstrate the usefulness. The overlap is the area shared by two pdfs, where a value of 100% signifies total overlap. Smaller overlaps mean more separation between the classes and so better use for classification should result.

Reflectivity

Reflectivity shows good separation (Fig. 11). Although rationally there is not an exact limit on the reflectivity that can be observed from most class types, there is nevertheless a climatological difference in the typical reflectivity values seen.

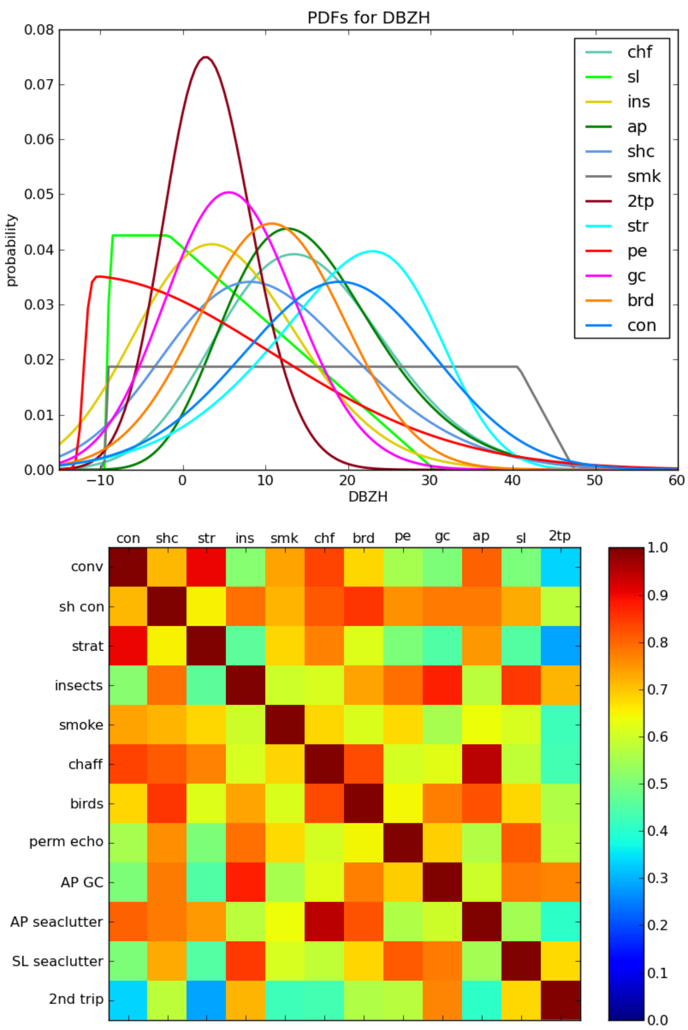

Fig. 11 Final pdfs (left panel) for reflectivity (DBZH) for all classes and overlaps (right panel) indicated by colour, where dark red is 100% overlap and dark blue is 0% overlap.

Texture and spin

Texture of reflectivity and velocity are expected to be independent. In general, the results were not found to be particularly sensitive to the kernel size, apart from near either extreme of kernel size. Very small kernels would not contain enough information to be useful, and very large kernels would average too much as well as be more likely to encompass multiple echo types. Kernel sizes of 11, 15 and 19 were tested and were not found to be very different. The pdfs and overlaps for SPIN2D for the three kernel sizes are shown in Fig. 12, Fig. 13 and Fig.14. The equivalent plots for TEX2Dare shown in Fig. 15, Fig.16 and Fig. 17. For reflectivity, a smaller texture kernel and larger spin kernel were considered to optimise separability and independence of features. The size of the kernel had a greater influence on SPIN2D than TEX2D, though SPIN2D overall showed less separability between classes. The texture of velocity used a moderate kernel of 15×15, with pdfs and over laps in Fig. 18.

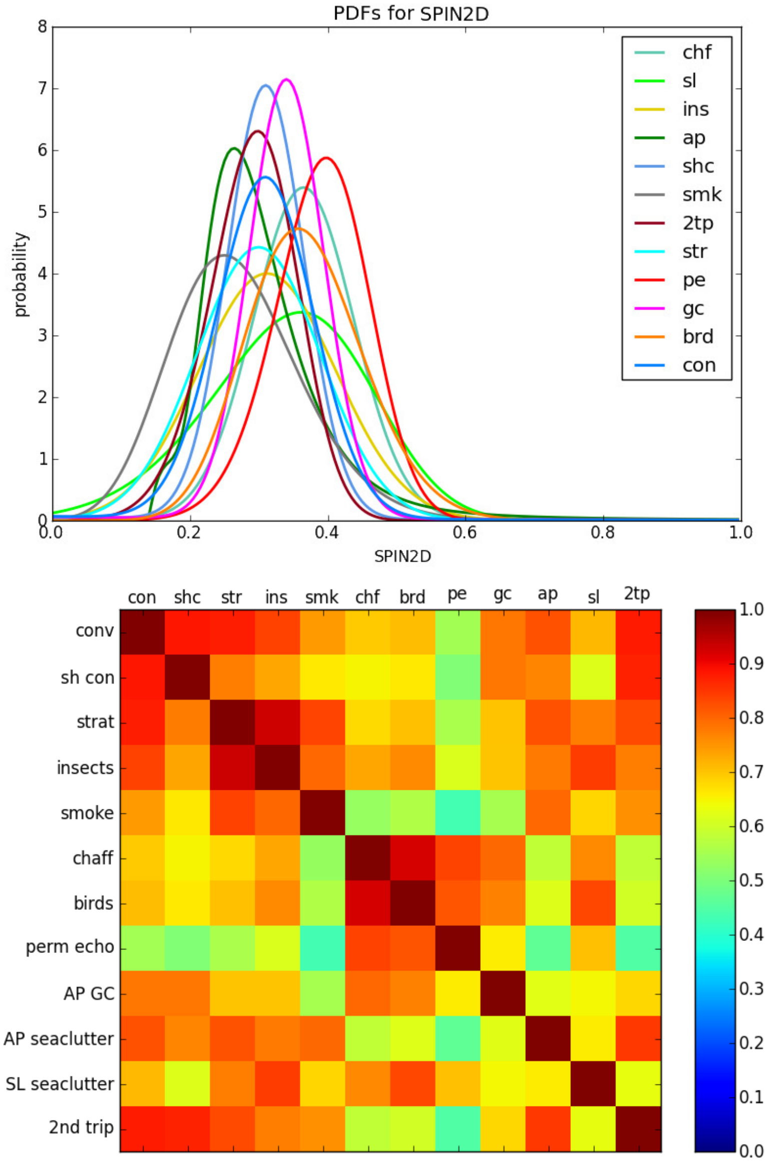

Fig. 12 Pdfs and overlap grid for spin of reflectivity for a kernel size of 11×11. The overlap percentages range from dark blue (0%) to dark red (100 %) where blue is desirable.

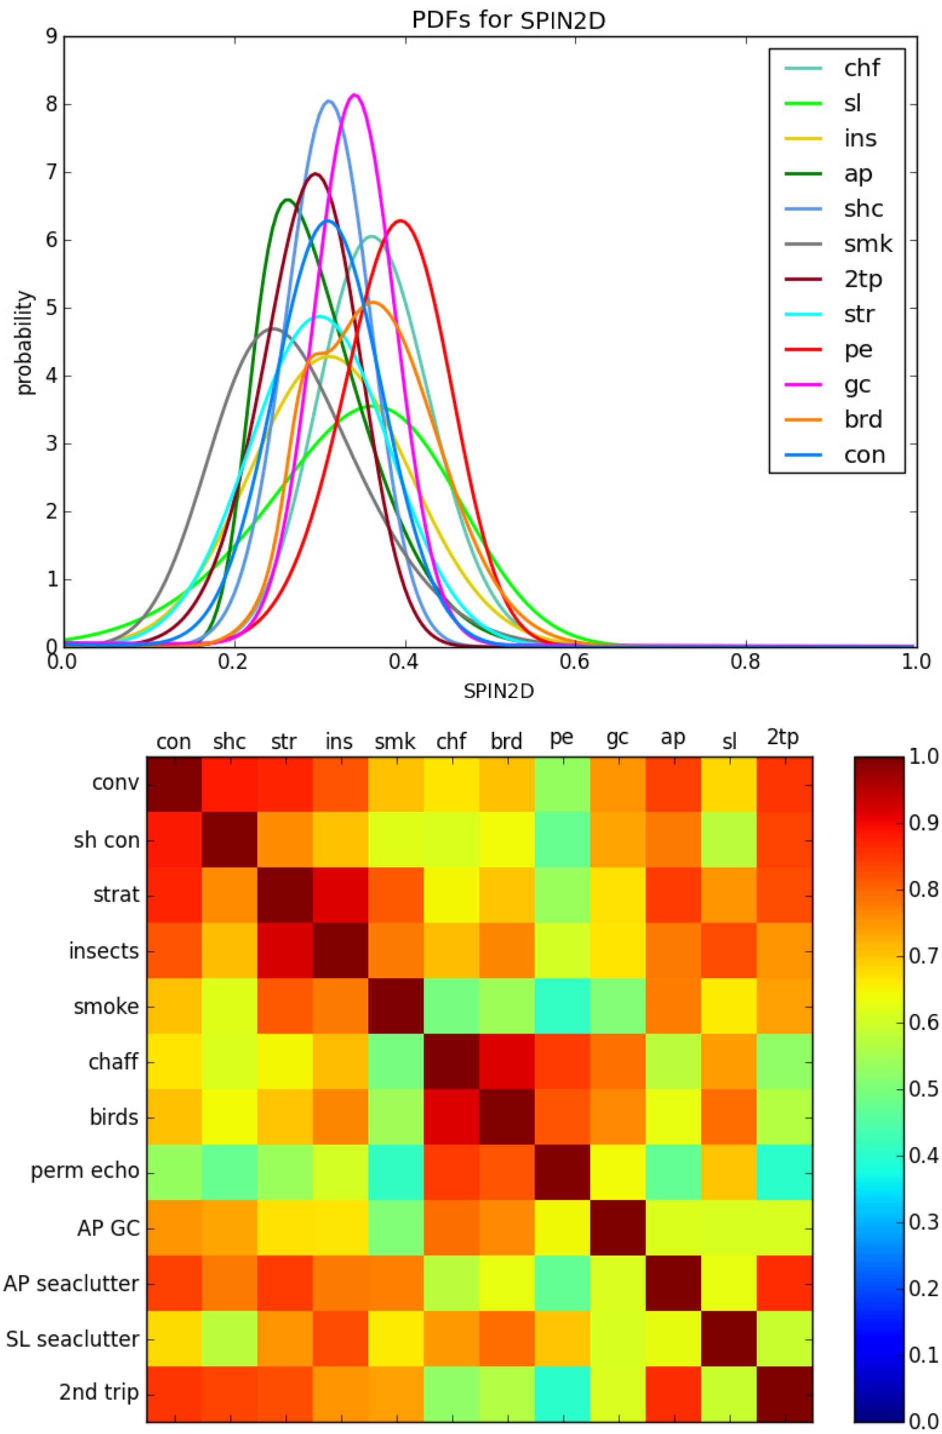

Fig. 13 As per Fig. 12 but for a kernel size of 15×15.

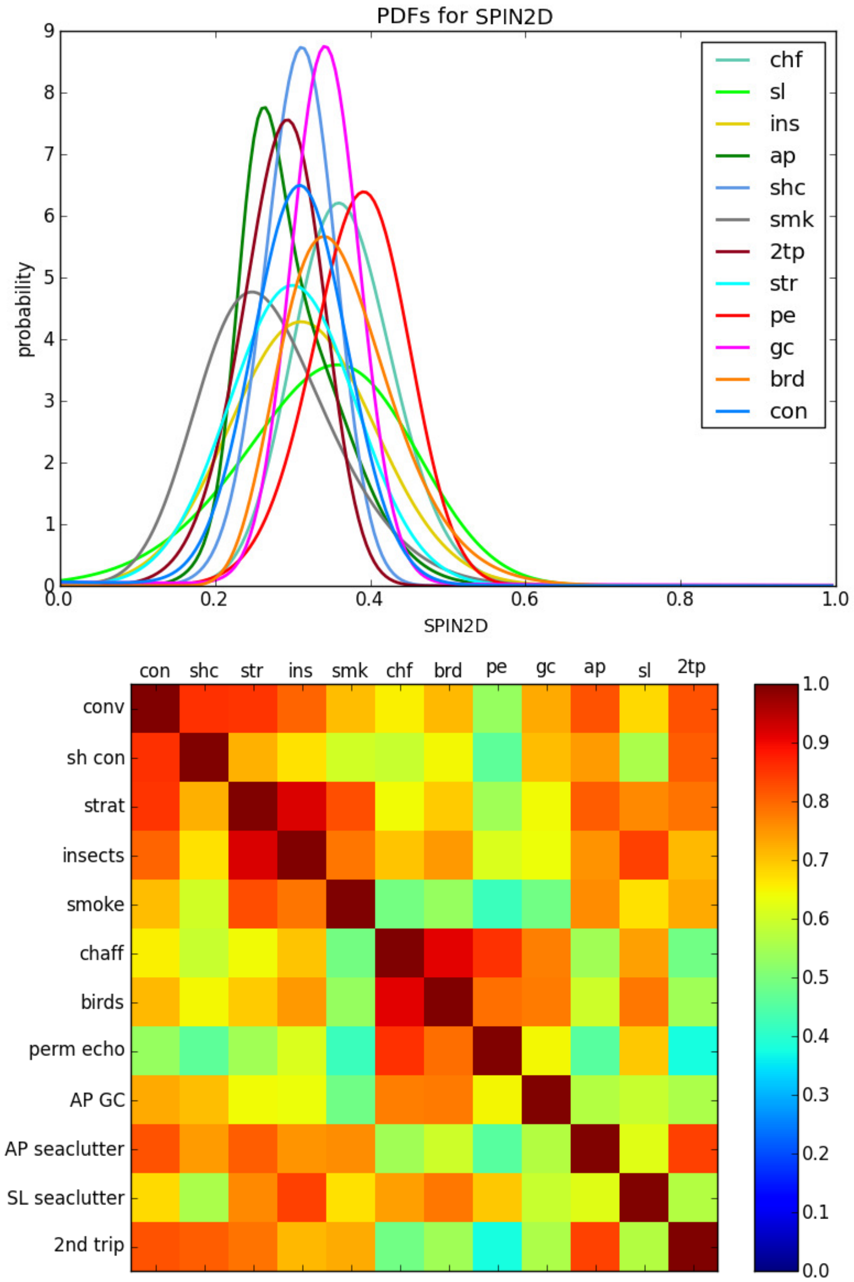

Fig. 14 As per Fig.13 but for a kernel size of 19×19.

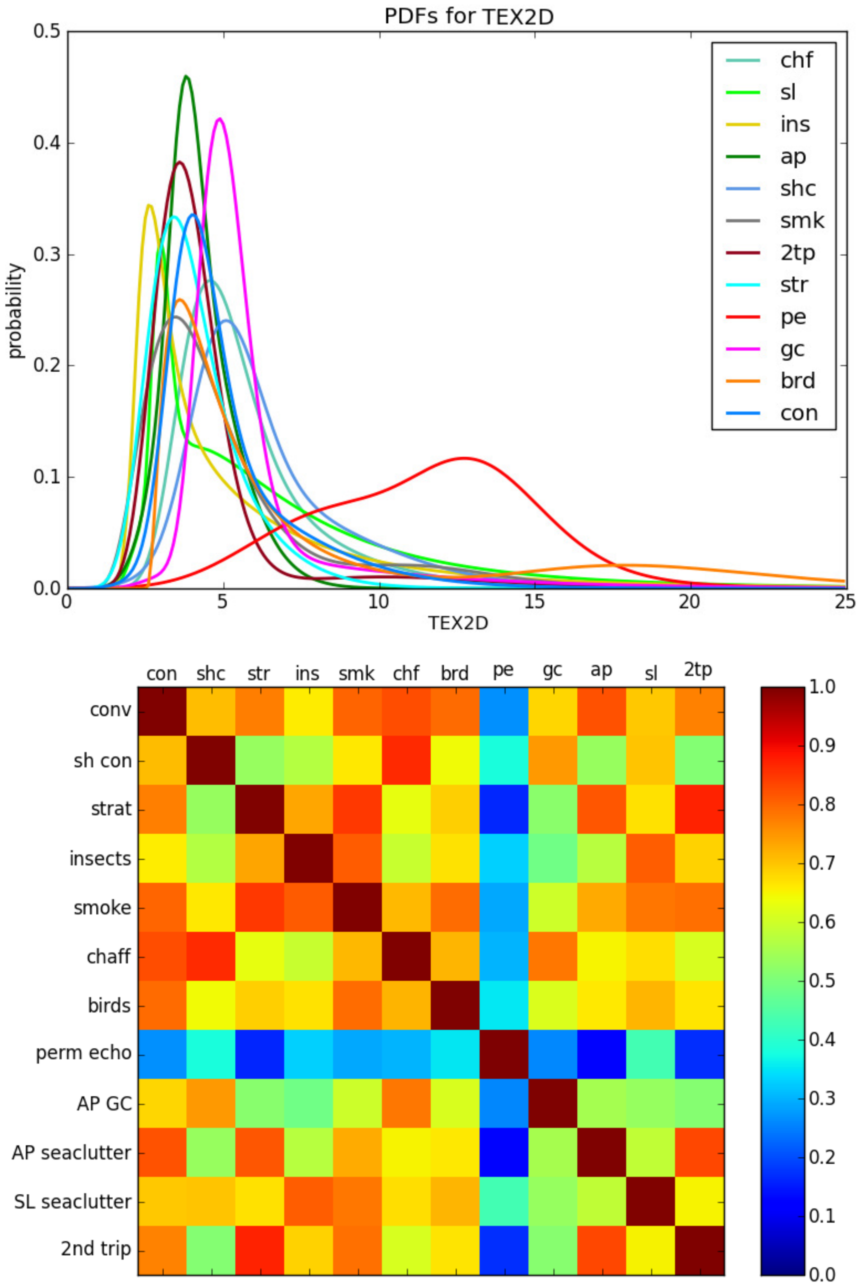

Fig. 15 Pdfs and overlap percentages for reflectivity texture with a kernel size of 11×11. The overlap percentages range from dark blue (0%) to dark red (100 %) where blue is desirable.

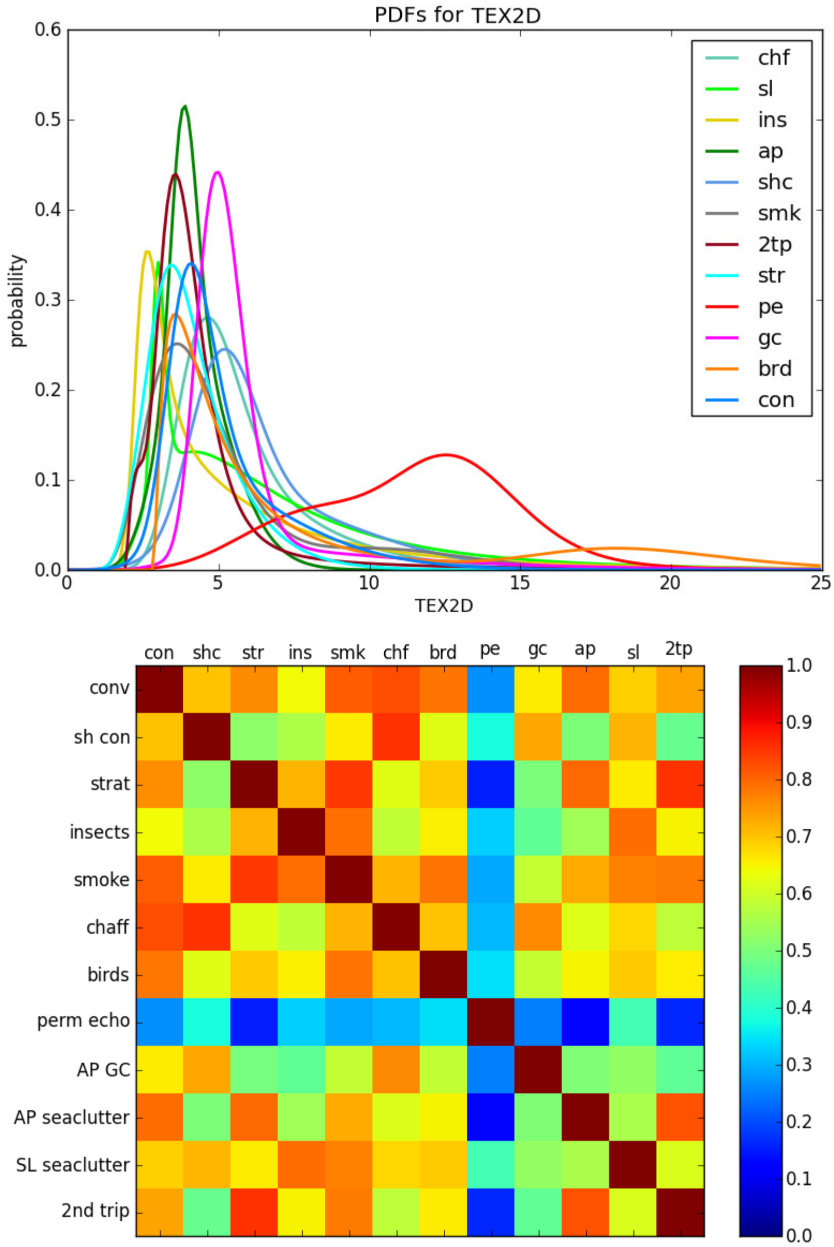

Fig. 16 As for Fig. 15 but for a kernel size of 15×15.

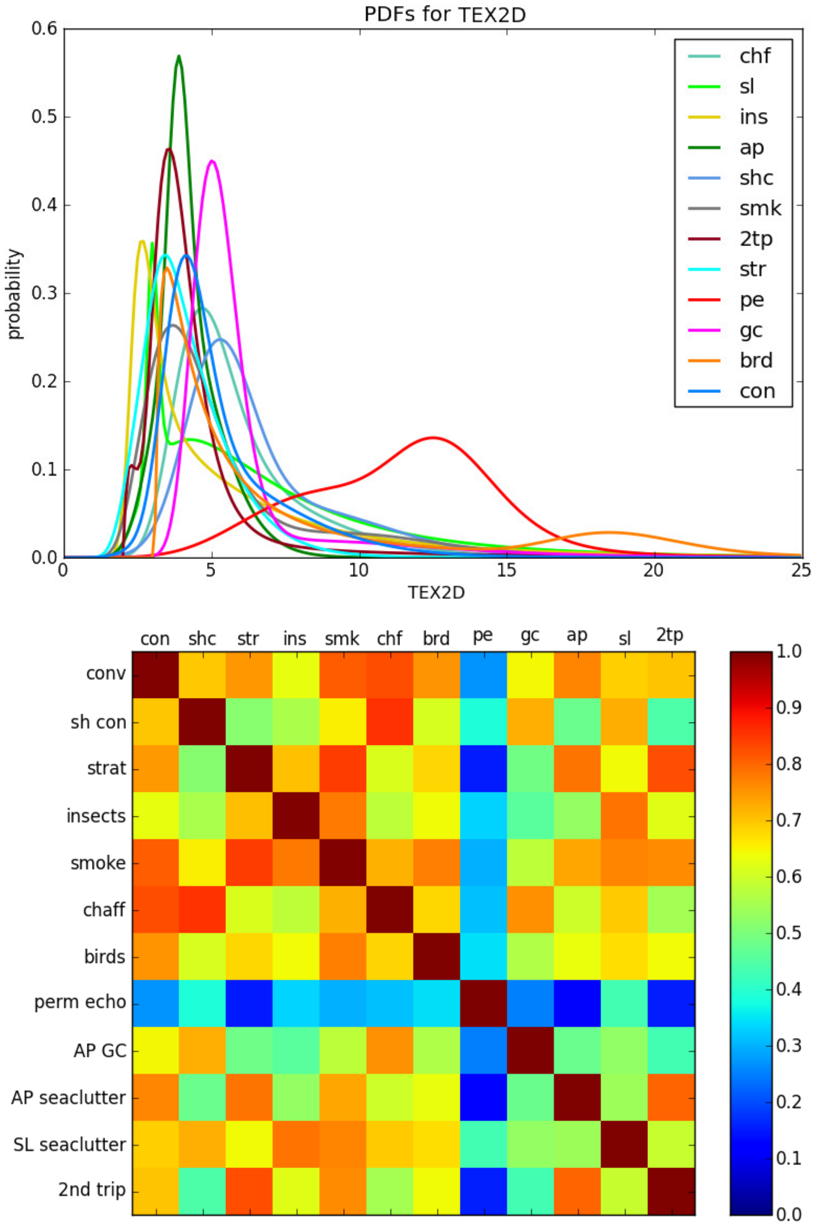

Fig. 17 As for Fig. 15 but for a kernel size of 19×19.

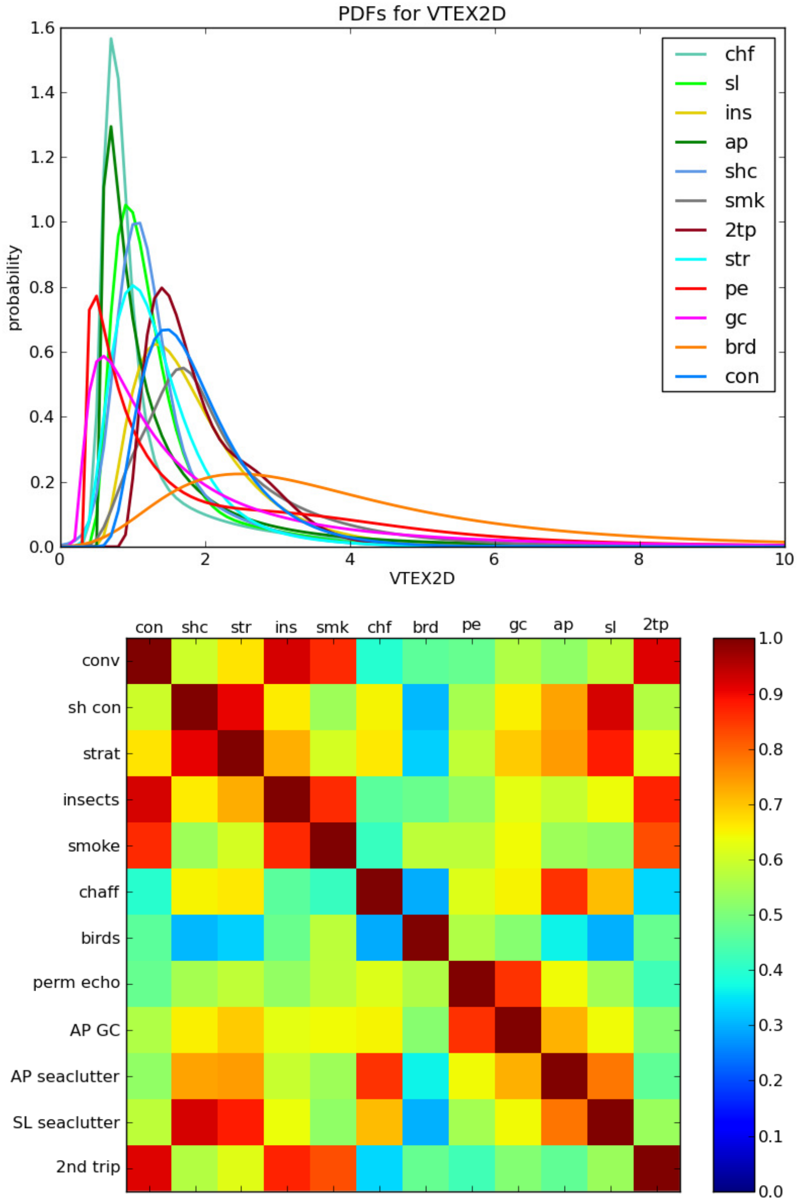

Fig. 18 The pdfs and overlaps for VTEX2D.

Echo top height

Echo top height provides valid values only where the reflectivity in the vertical column exceeds the threshold at some height. For precipitation, a higher threshold (e.g. 5–10 dB) may be more sensible, i.e. near the lower limit of measureable precipitation. Farther from the radar, the minimum reflectivity (noise threshold limit) is greater than 0 dBZ, and the distance between elevations is greater, so a low threshold is not effective. However, near the radar, there are regions where clear air echo does not exceed 0 dB, and so echo top height would not be available unless a lower threshold was used. Using the lower threshold gives almost the same result but a broader coverage. However, slightly better classification results (tested against the training dataset; see Section 5) were achieved by combining the thresholds. Echo top height is a quite effective field for classification (Fig. 19 and Fig. 20).

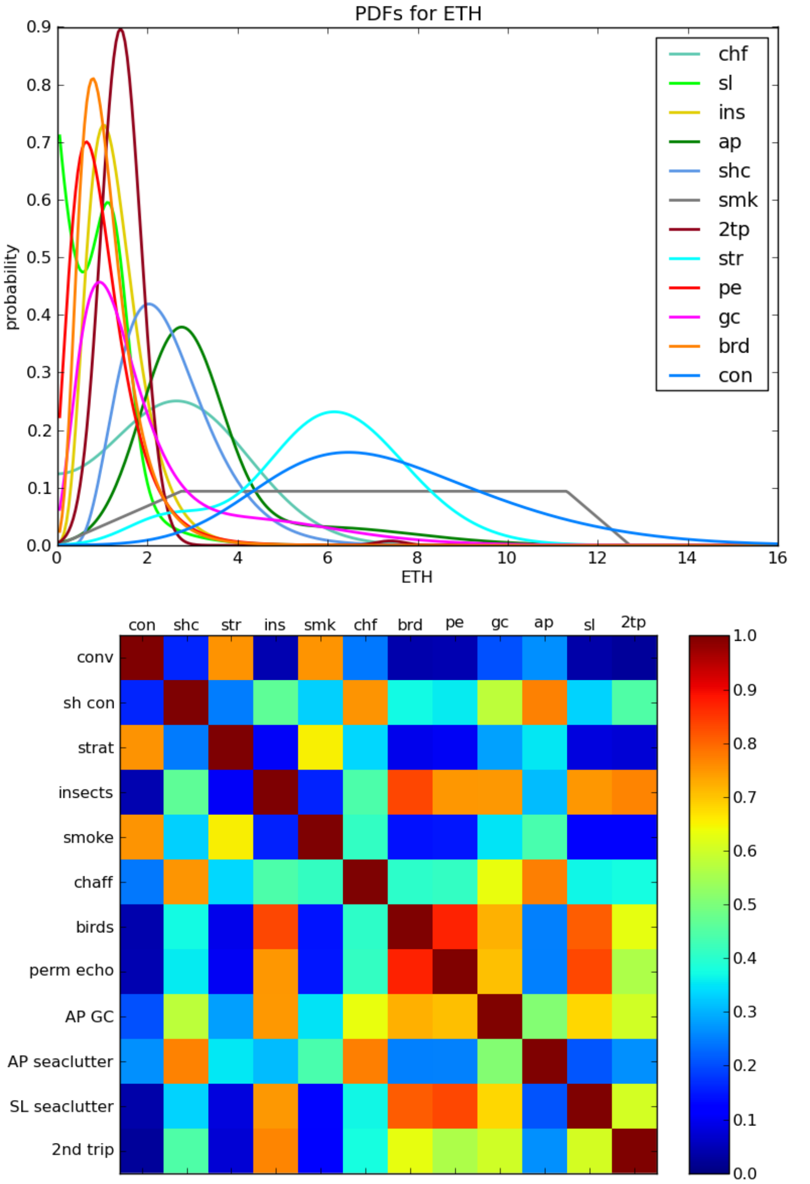

Fig. 19 PDFs and overlaps for echo top height with a threshold of 4 dB (ETH).

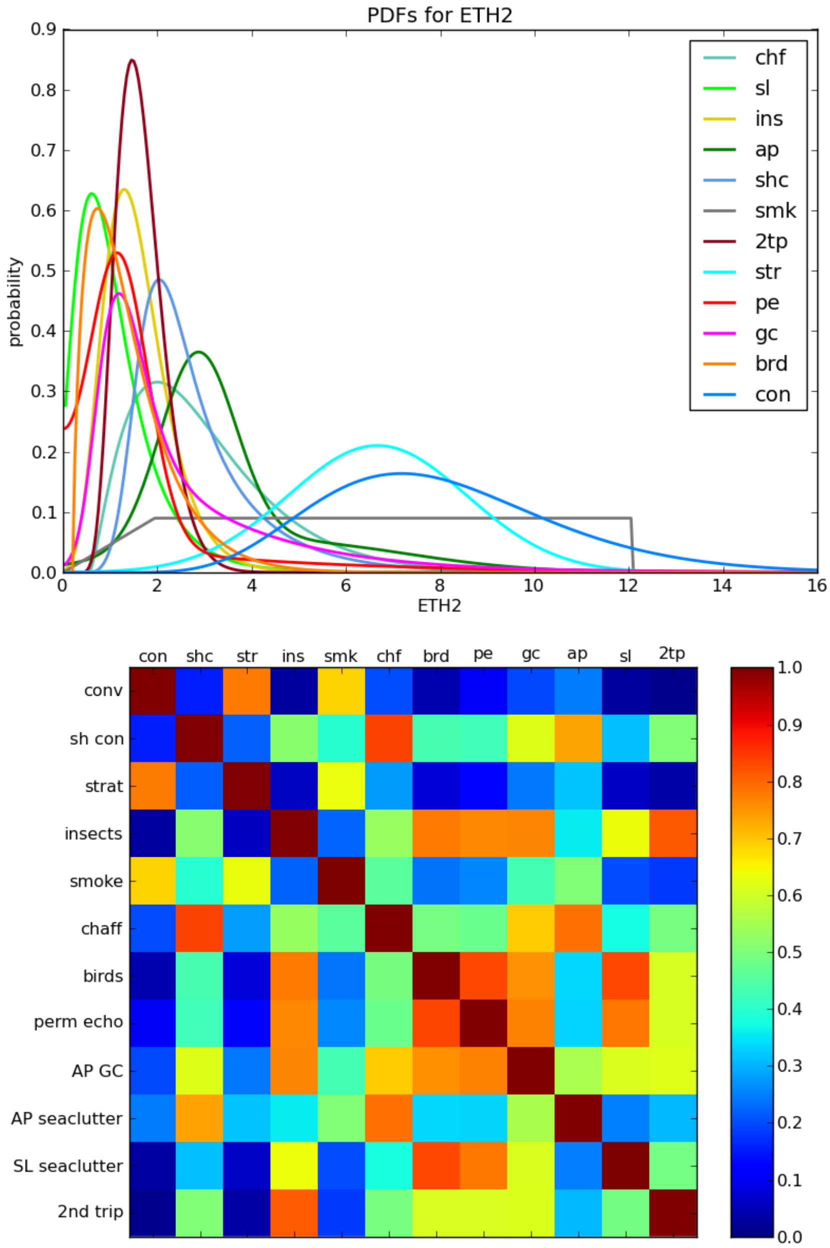

Fig. 20 Pdfs and overlaps for echo top height for a threshold –5 dB (ETH2).

Vertical gradient

The vertical gradient of reflectivity does not offer a large amount of discrimination as most differences are around the same values (Fig. 21). However, some classes have a tendency toward a higher negative gradient, such as surface echo, and vertical gradient helps to locate these. Precipitation, in contrast, should have a vertical gradient closer to zero.

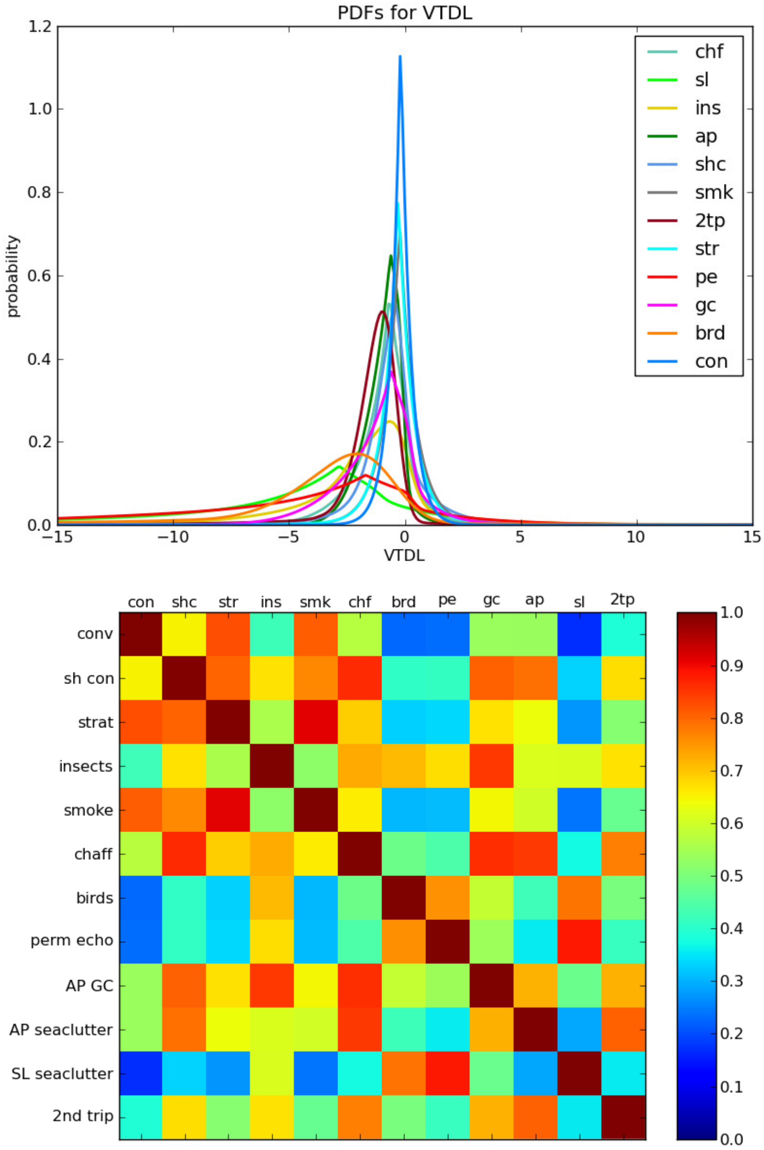

Fig. 21 Final pdfs for vertical gradient of reflectivity, and overlaps.

Spectrum width

The spectrum width had limited availability, coming from only one radar initially, but was found to be useful as an additional field. It is becoming more widely available. The point where spectrum width becomes less useful is at low values. Stratiform rain and dispersed chaff (that has been in the atmosphere for a while) both have very low spectrum width, and their pdfs spike at the lowest WAVG value (Fig. 22). This encourages any echo with extremely low spectrum width to be identified as stratiform precipitation, which will most adversely affect the identification of chaff.

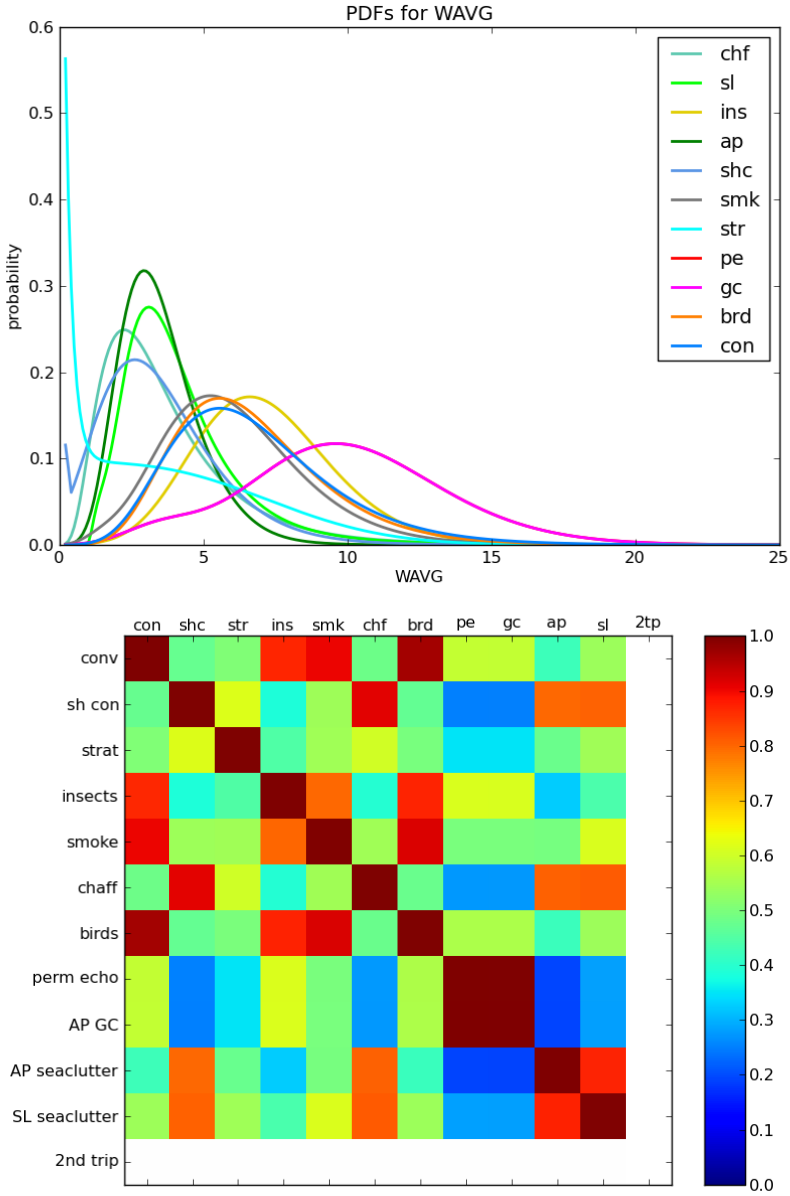

Fig. 22 Pdfs and overlaps for spectrum width.

4.2 Comparison of different range resolutions

The histograms were calculated for 250 m and 500 m range resolution, aggregating values from all radars with these resolutions. The proportion of each class's training data that belonged to these range resolutions varied. For the classes with ample sample size, there was not a great difference in the mean and spread of values for the feature fields (i.e. the precipitation classes). There were many differences between the pdfs for the two resolutions, but often the sample size for either the 250 m or 500 m radars was small for one of the classes, so a good comparison was not available. So if there is a significant difference in feature field values like texture or spin that relates to range resolution, it has not been demonstrated definitively.

Differences also do not appear to be systematic, i.e. the difference between pdfs for a feature field does not vary in a consistent manner, such as one having a lower mean than the other. However, for large sample sizes the histograms tend to converge, which suggests that differences are not significant ultimately. On the other hand, there are likely differences between different radars that are a combination of radar resolution, beamwidth, sensitivity, possibly band (wavelength), and geography that affect the retrieved echoes' characteristics. The sensitivity is a particular issue and is one reason why several different schema were adopted for different radar types, and permitted some customisation of pdfs.

4.3 Prior probabilities

The prior probability used in the NBC should reflect the likelihood of any class appearing in a scan. The first consideration is the relative occurrences of the different classes and how that information can be constrained.

Precipitation is generally the most common echo type. Among the three precipitation classes, there may be some difference between the classes' relative frequency of occurrence. Convection may be more common in tropical regions. In temperate regions, stratiform precipitation and shallow convection are relatively more common.

Insect echo is widely present day and night in many parts of Australia. It is generally confined to land, as insects actively avoid travelling over the ocean. Insects are thermally dependent, so the colder Tasmanian climate limits insect numbers. The insect echo also decreases in winter in temperate regions. Tropical regions do not seem to have as much insect echo as temperate regions. Insect echo is also undetected by the less sensitive radars, and the radar type and sensitivity probably biases the apparent insect density—unavoidable with such a heterogeneous radar network.

Bird or bat echo is typically present in concentrated areas for a limited time. Most instances are at sunrise or sunset over land. The examples are predominantly dusk or dawn dispersals from roosting areas, though some instances of migration corridors across the Port Phillip Bay near Melbourne are observed. Large scale migrations are not generally observed in Australia.

Smoke is present very sporadically during the hot/dry seasons, which differ for tropical and extra-tropical Australia. Smoke can be carried offshore although fires are generally constrained to onshore where the fuel is.

Chaff is present sporadically in somewhat limited regions, i.e. near air force bases and away from populated area—usually offshore. In the last few years chaff has been seen off Perth, offshore of the eastern coast between Wollongong and Brisbane, and inland north of Namoi. However, chaff carried by the wind can drift across coastlines. It takes several hours for a chaff trail to disperse and fall to earth.

Ground clutter is only detected over land, or along the coastline in mixed land/sea pixels. Permanent or normal-propagation ground clutter can be located with a probability of detection map. AP ground clutter occurs in most radar locations with varying degrees of frequency.

Sea clutter is analogously located over the ocean or along the coastline. The side-lobe sea clutter appears on probability of detection maps. AP sea clutter is common off northwest WA and Sydney although it can appear at any coastal radar.

There are several factors listed above that can be used to limit the number of classes that the classifier must distinguish at any location. Spatial factors include the land/sea location, the probability of detection map, and the radar location. Temporal factors are seasonality and time of day. Spatial variations have been built into the system. However, temporal conditions are much harder to implement because decisions of how the prior probabilities should change with time are more complicated. For example what classes are sensitive to sunrise and sunset, or vary with day/night, and how does the variation manifest. Developing and tuning this was not considered a priority with the resources available, although in future this option may be implemented. Note also that a different classification schema for each radar complete with its own range of prior probabilities and pdfs would be difficult to construct and maintain. Alternatively, grouping radars by, for example, whether or not chaff has ever been observed, or by radar type or location, is possible as it creates a limited range of schemata.

Probability maps were created for each radar, using a period of at least three months where possible. (On occasion this was not possible if the radar was new or if the scan parameters had changed.) The number of reflectivity gates falling within a range of thresholds was calculated. The POD for any threshold could then be determined by aggregating the count for all thresholds above it and dividing by the number of scans used for the POD. This made it easier to select a threshold to capture the permanent surface echo but avoid the omnipresent clear air echo.

The distance from the coastline was calculated for all sample cells in the standard volume from each radar, with overland values positive and over-sea values negative. This makes it easy to select land and sea bins and allow a buffer between them. For example, insect echo could extend to 8 km from shore. Ground clutter echo could extend to 1 km from shore, to account for coastal bin locations.

The values for the prior probabilities were next selected. These can be tuned, but were initially chosen to give a reasonable classification result with sensible values considering the climatology.

Precipitation is one of the most common echo types, so has the highest probability, especially if the precipitation classes are combined. Insects also have a high probability. The values shown below (Table 6) were selected as a first guess. The probabilities in the probability of detection map (a value between zero and one) may be used for the prior probability. Note that the probabilities act as weights in the classification algorithm, and do not need to sum to one.

Table 6 The prior probabilities for each class.

Class	Prior Probability	
Conv precip	0.4	
Shallow Conv	0.25	
Stratiform	0.5	
Insects	0.4 over land	0 over sea
Chaff	0.05	Not at all radars
Smoke	0.1	
Birds/bats	0.2 over land	0 over sea
Ground clutter	POD over land	0 over sea
SL sea clutter	1.2 * POD over sea	0 over land
AP GC	0.1 over land	0 over sea
AP sea clutter	0.1 over sea	0 over land
2^{nd} trip	0	

These probabilities may not reflect the true climatology. In the NBC, a very low probability for a class that is not easily distinguishable from other classes would mean that that class is never likely to be selected; therefore chaff and smoke have higher probabilities than climatology. Echoes from birds and bats are assigned a high probability, although they appear rarely except at certain times of day. These echoes are fairly easily distinguished from precipitation, and so any echoes that could be mistaken for birds are not desirable for assimilation. By permitting the bird class at all times, some outlying echoes may be excluded via this classification, particularly if the velocity texture is high. We elected to ignore second trip echo, so its probability is zero.

5 RESULTS USING THE CLASSIFIER

The schema described in Section 4 was applied to the training dataset, and implemented in the operational Ancilla version. The success of the schema can be evaluated in two ways. Firstly, by comparing the Ancilla automated classification with the manual classification in the training dataset. Secondly, by monitoring the outcome of the schema on the radar data in real time. An ideal alternative would be to create an independent manually classed dataset. However, this would be very time-consuming and the resources are not available for this.

Some examples of how the classifier performed for various 'typical' examples of a range of classes are shown below. These examples are not in the training dataset. Each example contains panels for the reflectivity and classification for a scan, the reflectivity identified as precipitation, and the radial velocity of echoes identified as clear air or precipitation. The latter two are the fields likely to be used for most applications. Figure 23 shows an example of clear air echo, Fig. 24 shows an example of shallow precipitation, Fig. 25 shows an example of sea clutter and Fig. 26 shows classification of chaff. The identification of clear air echo (Fig. 23), side-lobe sea clutter and shallow precipitation (Fig. 24) are all reasonable. The sea clutter is in part classified as precipitation, and in part as non-precipitation (clutter or chaff). An incorrect classification of sea clutter as chaff is considered acceptable for our applications. Note that echo with low reflectivity (e.g. <–10 dB) may be removed as well for Doppler applications, and for precipitation estimates reflectivity <15 dB may be excluded.

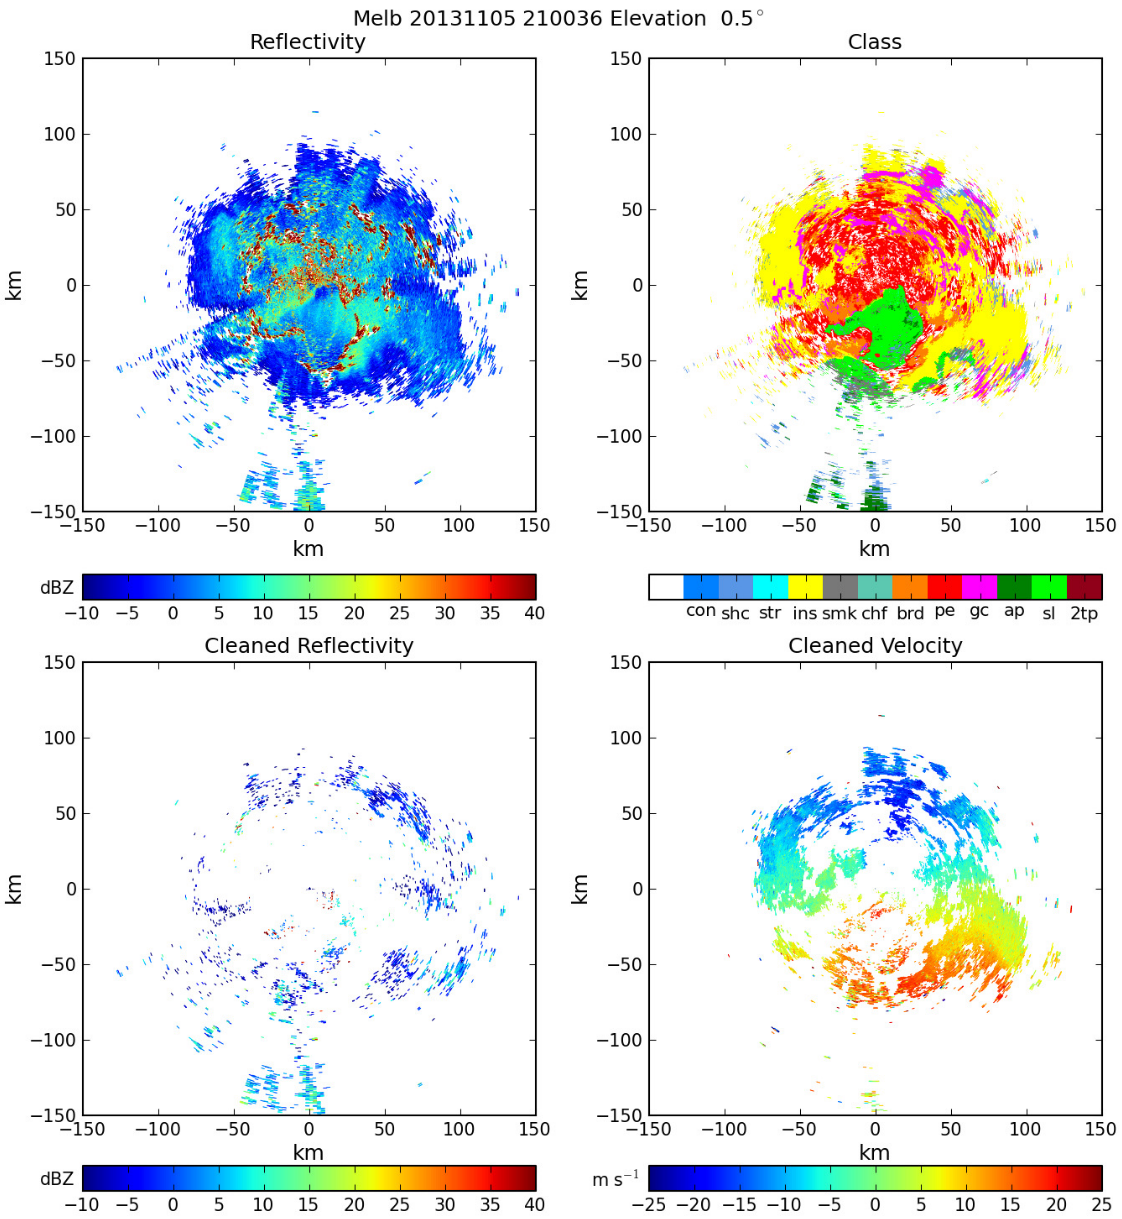

Fig. 23 Classification of clear air at Melbourne on 5 November 2103 at 2100 UTC. The top left panel is the reflectivity, top right is the classification. Bottom left is the reflectivity classified as precipitation. Bottom right panel contains the radial velocity classified as precipitation or clear air. The ground and sea clutter are mostly accurately classified, where the ground clutter is either permanent echo, AP ground clutter or classified as birds. Some insect echo has been identified also.

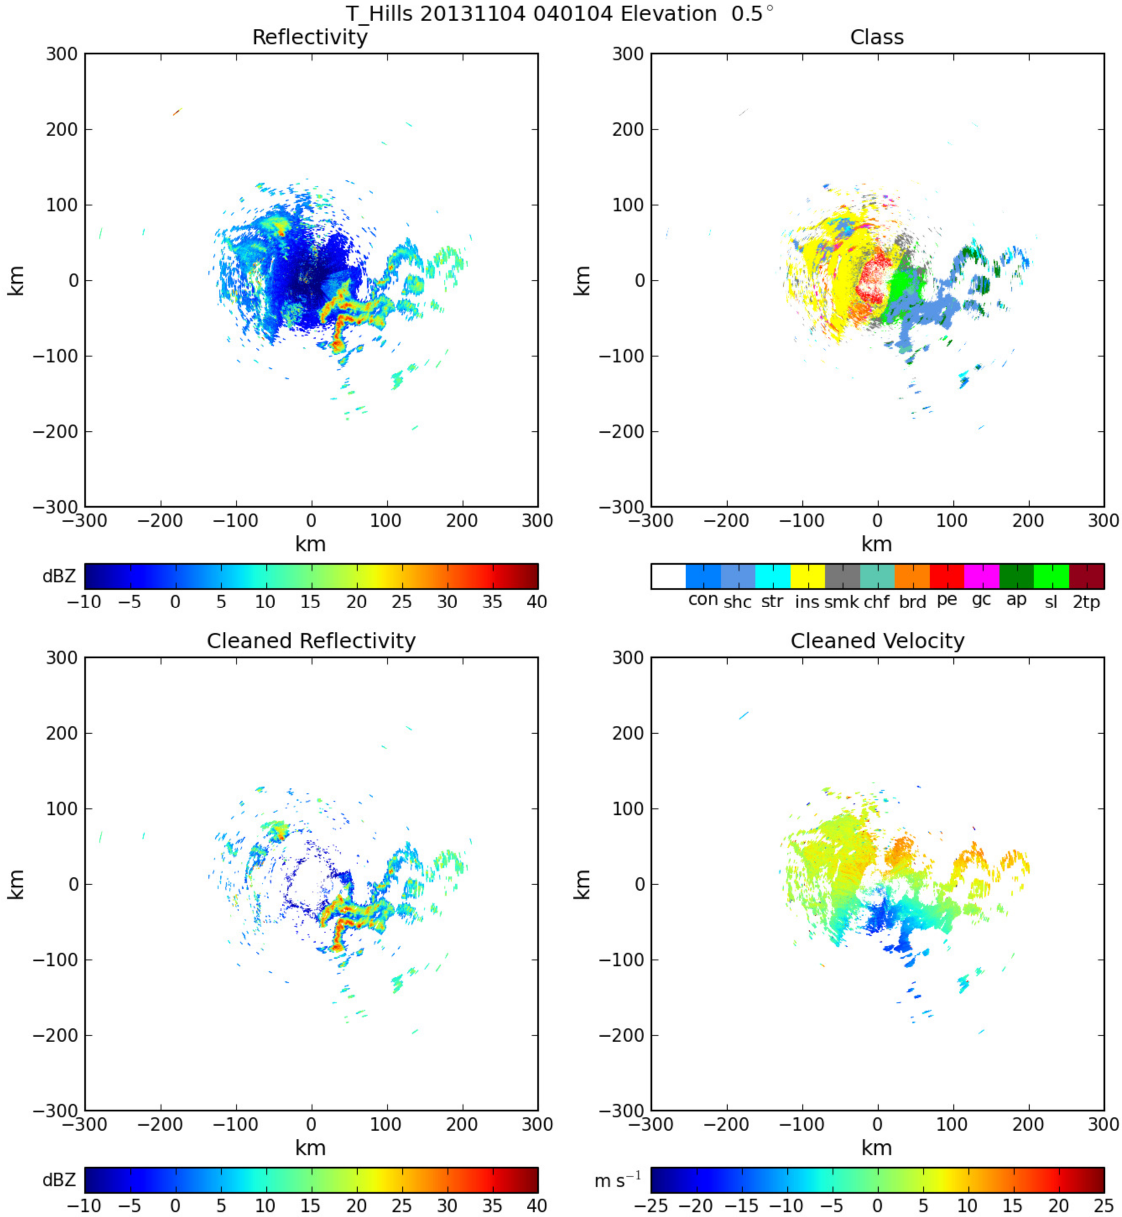

Fig. 24 Classification of shallow convection embedded in side-lobe sea clutter, from Sydney (Terrey Hills) on 4 November 2013 at 0401 UTC. Panels as per Fig. 23. Most of the precipitation is identified correctly; a small part is classified as AP sea clutter.

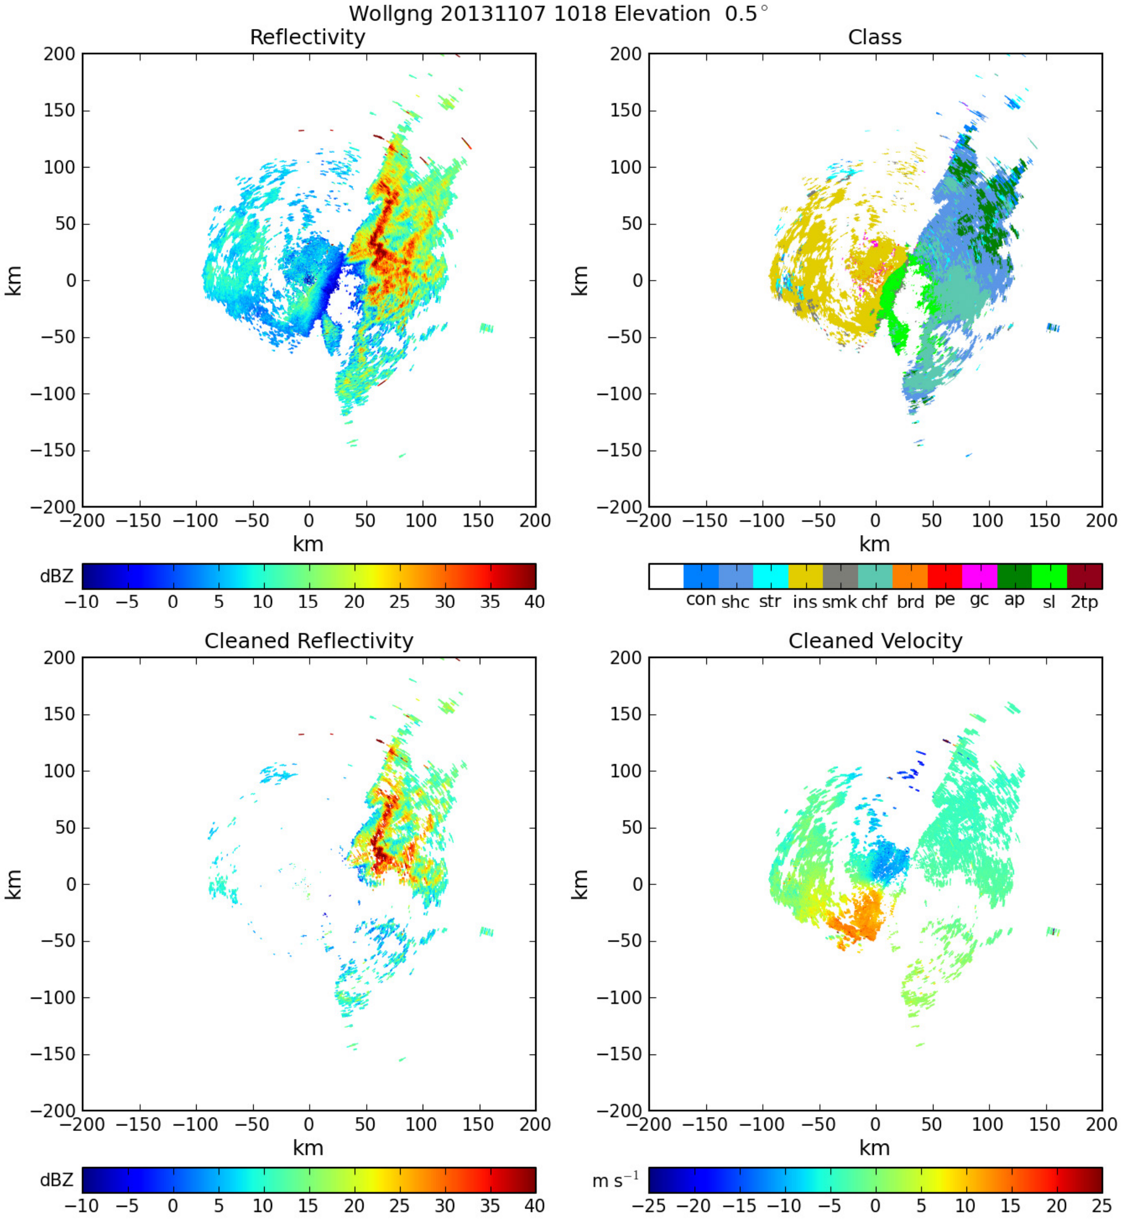

Fig. 25 Classification of AP sea clutter at Wollongong on 7 November 2013 at 1018 UTC. Panels as per Fig. 23. Much of the AP is classified as one of the unwanted clutter types. Some is classified as precipitation, primarily shallow convection. Some insect echo is also classified correctly. The resultant reflectivity could be described as reduced rather than removed AP.

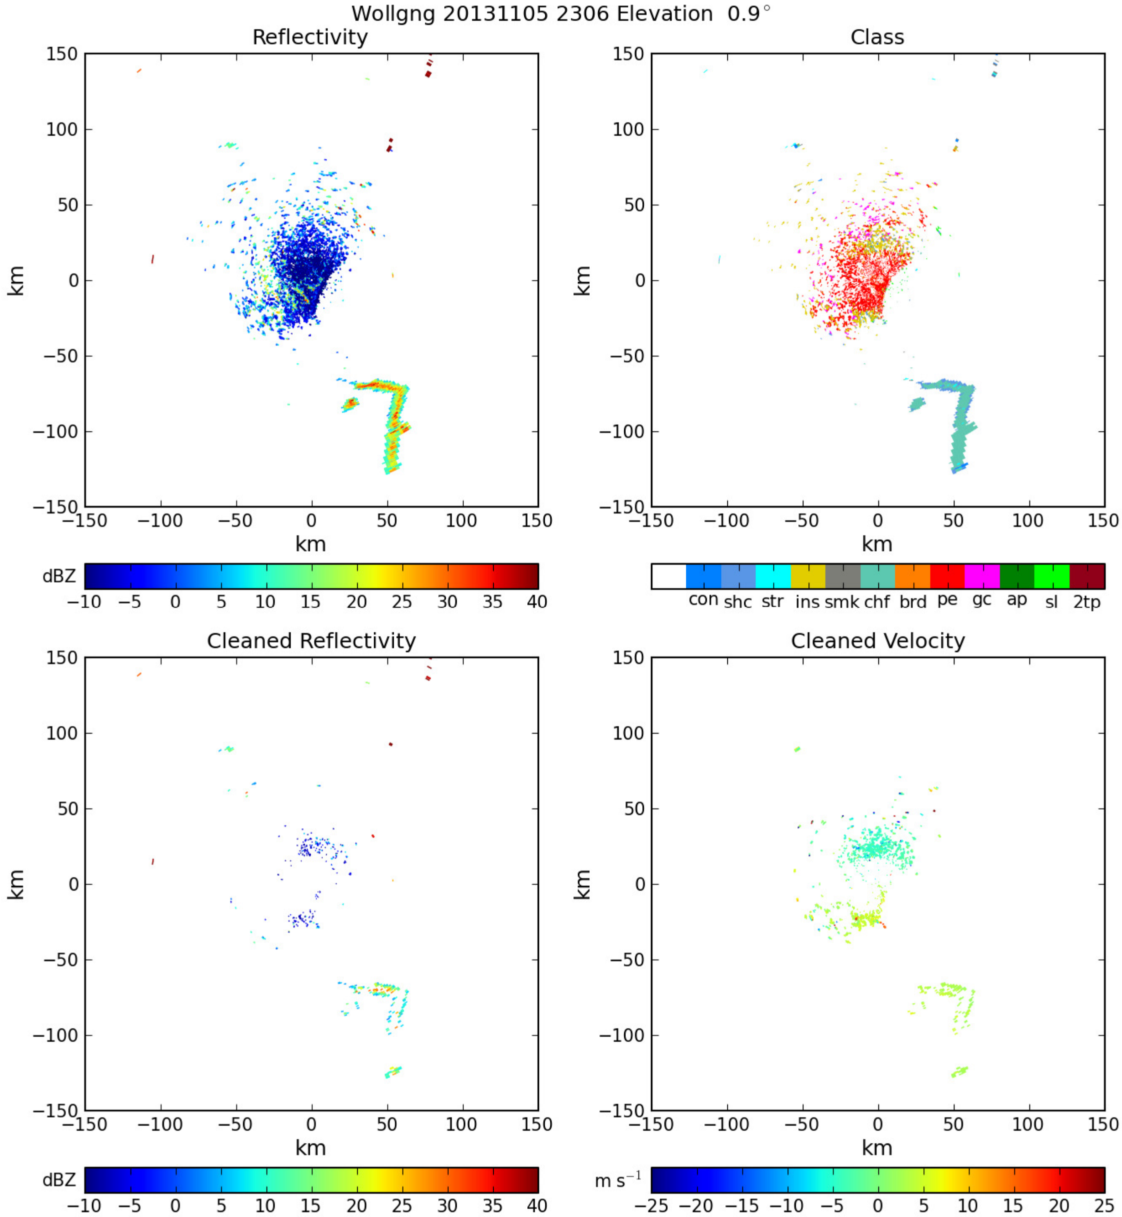

Fig. 26 Classification of chaff at Wollongong on 5 November 2013 at 2306 UTC. Panels as per Fig. 23. Some of the chaff is classified correctly. Most is classified as precipitation.

The standard version of the classification schema was applied to the training dataset. A correct classification is when the automated and manual classification values are identical. The overall success is found by counting the number of each automated class value for each manual class value. These are converted to percentages, where the values for each manual class sum to 100%. Results are shown in contingency tables, where the automated classification is in columns and the manual classification in rows. This format is used throughout this report. The classification results for the training dataset are shown in Table 7. Values are given as percentages, and the sum total pixels are shown in the far right column. The diagonal represents a correct classification.

Table 7 Results of the classification applied to the training dataset. The manually classed classes are in rows and the classifier output classes are in columns. All values are percentages and the rows add up to 100%. The total number of classed pixels is shown in the Counts column. The diagonal of correct classification is shown in bold.

man\aut	con	shc	str	ins	smk	chf	brd	pe	gc	ap	sl	2tp	Counts
con	**45.2**	**7.7**	**40.9**	1.5	2.2	0.3	0.1	0.2	1.6	0.3	0.1	0.0	2813232
shc	**3.4**	**66.9**	**6.0**	10.3	0.6	3.1	1.1	1.9	1.5	3.2	2.1	0.0	420606
str	**24.7**	**8.7**	**59.8**	1.2	1.1	2.4	0.2	0.4	0.4	0.8	0.2	0.0	7544563
ins	0.9	6.2	1.6	**73.6**	1.5	0.1	6.5	5.9	3.0	0.1	0.7	0.0	8166983
smk	42.5	15.6	8.3	14.2	**11.7**	0.0	2.1	3.3	2.2	0.0	0.0	0.0	631508
chf	6.0	19.8	29.1	6.9	0.5	**26.5**	1.6	1.3	0.6	6.0	1.6	0.0	1319968
brd	0.0	3.5	0.9	27.7	0.8	0.0	**51.4**	7.5	5.4	0.0	2.8	0.0	48201
pe	1.8	6.3	0.4	10.2	0.8	0.0	8.6	**68.6**	2.6	0.0	0.7	0.0	132677
gc	2.9	28.5	6.0	19.8	1.6	1.4	7.5	6.2	**26.0**	0.0	0.0	0.0	367101
ap	6.1	17.2	17.7	0.0	0.7	11.7	0.0	0.0	0.0	**41.7**	5.0	0.0	365141
sl	0.6	7.1	0.8	0.2	3.4	1.8	0.0	0.2	0.0	1.4	**84.4**	0.0	694275
2tp	1.0	18.3	0.9	57.4	1.2	1.7	3.7	0.6	1.2	3.3	10.8	0.0	19176

Precipitation types are not exactly discriminated, and these are expected to be combined in any case. However, most precipitation is correctly identified as such. Combining the precipitation classes results in a correct classification rate of 92.7%. Similarly the clear air (insects and smoke) and clutter classes have been combined, as shown in Table 8. This is more representative of how the classification information would be used.

Table 8 Collapsed classification results. Precipitation represents all three precipitation classes. Clear air includes insects and smoke. Clutter contains the remaining classes. The manual classes are in rows and the results of the classifier are in columns.

	precip.	clear air	clutter	Total pixels
precip.	**92.7%**	3.0%	4.3%	10778401
clear air	12.8%	**71.5%**	15.7%	8798491
clutter	36.9%	8.2%	**54.8%**	2946539

Some generalisations can be made about the performance of the classifier, which are borne out by the examples shown. Most precipitation is correctly identified, but sometimes spots are mistaken for clutter or clear air. These may be areas where the texture of reflectivity or velocity is unusually high. Many classes are easy to mistake for precipitation, i.e. AP, smoke, chaff. These have always been difficult to separate from precipitation from single polarisation radar because of the similar characteristics: reflectivity, height, scale. Consequently these classes form the majority of incorrect 'false alarms', i.e. non-precipitation classed as precipitation. AP clutter is not well identified, especially at long range, possibly because the assumed high beam height may be incorrect far from the radar. The training data did not include a lot of AP at long range so the echo top height may not well represent the climatology. The permanent echo, both ground and sea, is well identified thanks to the POD. Insect echo is reasonably well detected. The rare clear air types, smoke and chaff, are typically partly identified correctly. With the low prior probabilities these are unlikely to be completely identified. Their misclassification is mostly as precipitation.

5.1 Comparison with an existing method

Existing clutter removal code used for Bureau QPE provides a baseline for comparison. This existing method uses echo height thresholds and reflectivity thresholds. The method operates as follows. All reflectivity less than 5 dB is excluded from being precipitation. Areas where within the lower three scans the reflectivity in the upper scan is greater than 10 dB and the lower scan is less than 5 dB are classified as non-precipitation. For regions where the reflectivity is greater than 10 dB, the scan is found where in the vertical column the reflectivity passes below 5 dB. If the height of echo above 5 dB is less than 2 km, the echo in that column is classed as non-precipitation. In the bottom two scans, if the difference in reflectivity is greater than 10 dB, these regions are classed as non-precipitation. The remaining echo not excluded by these tests is classed as precipitation.

A comparison against this existing method is shown in Table 9. Since this method only classifies echo as precipitation or non-precipitation, and only for echo above 5 dB, all echo below 5 dB so has been excluded from the analysis. The Bayesian classification results have been collapsed to these two classes for comparison.

Table 9 Comparison of the Bayesian classification method with an existing threshold-based method, discriminating between precipitation and non-precipitation, for echo ≥5 dB.

		Bayesian				Thresholds	
		Automatic				Automatic	
		precip.	non			precip.	non
Manual	precip.	94.4%	5.6%	Manual	precip.	90.3%	9.7%
	non	23.5%	76.5%		non	46.7%	53.3%

Only using observations ≥ 5 dB results in better classification of precipitation (94.4% cf. 92.7% as indicated in Table 8) and a substantially lower false alarm rate. The thresholding method is almost as accurate at detecting precipitation but has a doubled false alarm rate. The results are biased by testing on the training set used to develop the NBC, so differences of a few per cent may not be significant, but the larger differences should be significant. Although thresholding is useful, there is not much to be gained by including it within the NBC because the role of thresholding is taken by the pdfs.

5.2 Suitability for applications

The success of the classification may also be judged by its suitability for different applications. This principally means that the balance of accurate precipitation detection to false alarm rate must suit the task. For QPE, maximising the detection of precipitation may have a higher priority, especially if combining with rain gauge observations. Excluding low reflectivity values results in a small improvement in the accuracy of precipitation classification, and also reduces the amount of clutter that can be classified incorrectly as precipitation, while having little effect on the precipitation amount. In practice, a reflectivity threshold of 15 dB may be used to exclude weak echoes from use in QPE.

For data assimilation, the priority is to reduce the amount of incorrect classification of bad observations as good; having no data is preferable to bad data. For example, a false observation of precipitation used to affect the model moisture would result in the model trying to force precipitation under where it shouldn’t exist. Therefore minimising the false alarms for precipitation detection is the priority. This can be managed by reducing the prior probability for precipitation and increasing the prior probability for clutter so that overall less of the scan will be classified as precipitation. However, the proportion of false alarm observations may not be affected greatly even if their number is reduced. It is apparent from the histograms of the various feature fields that there is not a great amount of separation between the classes.

For radial velocity assimilation, the observation usage will be range-limited to 100 km or so, because of the increase in error with range. This error is primarily the representativeness error that results from beam broadening, and the uncertainty in the beams location since standard atmospheric refraction is assumed. These errors can become significant at long ranges (Fabry 2010). Therefore, accurate identification of safe radial velocity observations is only critical near the radar. In general, classification near the radar is more accurate, as at long range most echo is classified as precipitation. From the examples in Fig. 23 to Fig. 26, it seems that much of the velocity observations might be acceptable based on the classification. The observation processing that checks observation minus background values, and removes isolated pixels prior to assimilation, will still be vital to filtering the observations. Assessment of Doppler radar classification

5.3 Value of Doppler parameters in classification

Doppler radars provide additional information for classification. The radial velocity field is used to provide an extra texture field. The radial velocity itself does not help with classification, especially since the zero-velocity filtering is already applied. From a growing number of Doppler radars the spectrum width is received, providing another field. The impact of these fields on the classification success is explored by examining the training dataset classified with and without these parameters. The only radar with all these parameters during most of the period used for the training dataset was Wollongong. Almost all classes, excluding AP ground clutter and second trip echo, were classified using the Wollongong radar volumes. Additionally, a few volumes containing chaff and shallow convection were classified from Newcastle, which also has spectrum width. The tables below show the classification results for these radars using the normal schema, and without spectrum width (no W). Spectrum width produces a small improvement in precipitation detection (Table 10) that is probably not meaningful Overall the differences are generally negligible for the combined classes. The effect on individual classes is detailed in Table 11 and large differences are marked. The benefits occur where spectrum width is useful for discriminating between classes, and detriments occur when spectrum width is similar for echo types and so the classifier does not perform as well.

Table 10 Classification with and without WAVG collapsed to three classes. The precipitation classes, the clear air (insects, smoke) classes and the clutter (all other) classes. For each manual class, the first row is the standard classification result, and the second row is the result with no spectrum width.

	precip.	clear air	clutter
precip.	**91.3%**	1.1%	7.6%
no W	**91.1%**	1.5%	7.3%
clear air	5.3%	**71.6%**	23.1%
no W	4.4%	**72.3%**	23.3%
clutter	33.8%	9.0%	**57.3%**
no W	33.5%	6.2%	**60.3%**

Table 11 Classification of the training dataset showing the benefit of using spectrum width. For each class there are two rows. The first is the result using the standard schema. The second row is the difference if WAVG is not used (no W), i.e. standard minus no W. Along the diagonal, positive changes show benefit to using W. Outside the diagonal, negative values show benefit to using WAVG. Large benefits are marked in green. Large detriments are shown in red. The final column shows the number of classed pixels.

	con	shc	str	ins	smk	chf	brd	pe	gc	ap	sl	Count
con	**70.9**	1.5	26.7	0.1	0.7	0.0	0.0	0.0	0.0	0.0	0.0	173446
no W	22.5	0.0	-22.2	0.1	-0.4	0.0	0.0	0.0	0.0	0.0	0.0	
shc	12.2	**48.5**	23.4	2.7	0.2	10.6	0.1	0.0	0.7	1.5	0.2	156740
no W	3.0	4.5	-8.9	0.8	0.1	-0.2	0.0	0.0	0.5	0.0	0.0	
str	28.6	3.7	**58.3**	0.5	0.5	6.6	0.1	0.1	0.1	1.2	0.3	1164280
no W	17.0	0.1	-16.9	-0.6	0.1	0.9	0.0	-0.2	-0.5	0.0	0.0	
ins	3.0	2.1	0.5	**74.6**	0.3	0.0	10.1	4.8	2.6	0.0	1.9	117919
no W	1.2	0.9	-1.4	-0.9	0.2	0.0	1.1	-1.3	0.2	0.0	0.0	
smk	0.3	3.6	0.0	52.1	**0.2**	0.0	31.7	11.6	0.5	0.0	0.0	20227
no W	0.0	2.3	0.0	-0.8	-0.1	0.0	7.1	-8.7	0.1	0.0	0.0	
chf	5.1	22.1	10.7	6.0	0.2	**46.1**	0.6	1.3	0.6	2.5	4.7	271513
no W	-1.1	0.1	1.1	-3.6	0.0	5.5	-0.1	-1.3	-0.4	0.3	-0.5	
brd	0.0	1.5	0.0	49.2	0.1	0.1	**45.0**	0.3	3.4	0.0	0.5	10636
no W	0.0	1.2	-0.1	-3.6	0.1	0.0	1.7	0.2	0.2	0.0	0.3	
pe	0.0	0.1	0.0	12.0	0.3	0.0	11.1	**75.4**	1.1	0.0	0.0	1514
no W	0.0	0.1	0.0	0.7	-0.5	0.0	1.9	-3.2	0.9	0.0	0.0	
gc	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0
no W	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
ap	0.9	8.6	11.2	0.0	0.1	30.5	0.0	0.0	0.0	**41.3**	7.6	18232
no W	0.3	1.3	-7.7	0.0	0.0	-4.3	0.0	0.0	0.0	9.8	0.7	
sl	0.2	23.1	0.5	0.0	0.4	10.8	0.0	0.0	0.0	6.9	**58.1**	61526
no W	0.1	-0.4	-0.2	0.0	0.2	-5.0	0.0	0.0	0.0	1.4	4.0	

The effect of using VTEX2D is similarly examined by classifying the Doppler radars with and without VTEX2D. The 'standard' method in this cased used all feature fields including VTEX2D but excluding WAVG. The second method excluded VTEX2D (and WAVG) to yield the result if the radar were not Doppler. Table 12 shows the difference between the two methods. Table 13 summarises the classification results in three classes.

Table 12 The result of classification of Doppler radar volumes from the training dataset with and without using radial velocity texture. For each manual class (listed in first column), the first row is the classification result using VTEX2D, and the second row is the difference when VTEX2D is not used, i.e. standard minus no VTEX2D. Large benefits of VTEX2D marked in green, large detriments marked in red.

	con	shc	str	ins	smk	chf	brd	pe	gc	ap	sl	Count
con	**43.6**	7.7	42.3	1.5	2.3	0.3	0.1	0.2	1.7	0.2	0.1	2732885
no V	-6.0	-0.3	6.0	-0.2	-0.1	0.0	-0.1	0.0	0.8	0.0	0.0	
shc	3.6	**60.5**	10.0	13.0	0.4	3.5	1.4	2.4	1.8	1.8	1.6	394778
no V	0.8	-1.1	-0.9	0.2	0.0	0.5	-0.2	0.1	0.4	0.2	0.0	
str	22.4	8.7	**62.1**	1.3	1.1	2.3	0.2	0.5	0.5	0.9	0.2	7406842
no V	5.7	0.7	-6.9	0.1	0.5	0.5	-0.6	0.1	0.1	-0.2	0.0	
ins	0.9	6.1	1.6	**73.6**	1.5	0.1	6.5	5.9	3.0	0.1	0.7	8149562
no V	0.2	-2.0	-0.2	12.5	0.4	-0.1	-5.2	-2.6	-3.1	0.0	-0.1	
smk	42.5	15.5	8.3	14.2	**11.7**	0.0	1.9	3.6	2.2	0.0	0.0	631508
no V	4.1	-1.3	-7.3	2.7	3.7	0.0	0.1	-1.3	-0.9	0.0	0.0	
chf	6.3	19.4	29.2	7.7	0.5	**25.4**	1.6	1.6	0.6	6.0	1.6	1306442
no V	1.4	0.1	-2.1	3.0	0.3	-0.2	-2.6	-0.5	-0.1	0.8	0.0	
brd	0.0	3.2	0.9	28.5	0.8	0.0	**51.0**	7.4	5.3	0.0	2.7	48201
no V	0.0	-7.5	-0.1	0.8	0.8	-0.2	8.4	-0.5	-0.7	0.0	-1.0	
pe	1.9	6.5	0.4	10.4	0.9	0.0	8.4	**68.2**	2.6	0.0	0.7	128116
no V	0.5	-2.7	-0.6	-1.9	0.3	0.0	3.0	0.1	1.3	0.0	0.1	
gc	14.1	5.7	16.5	10.8	16.1	14.8	4.2	15.8	**1.9**	0.0	0.2	35775
no V	-4.4	-2.7	4.5	5.2	-5.0	3.2	1.9	-3.7	0.9	0.0	-0.1	
ap	4.5	14.8	25.8	0.0	0.5	39.2	0.0	0.0	0.0	**13.3**	1.8	110602
no V	1.7	-2.5	-1.8	0.0	0.2	-0.4	0.0	0.0	0.0	3.0	-0.3	
sl	0.8	9.6	1.1	0.2	2.5	3.5	0.0	0.3	0.0	0.9	**81.1**	443922
no V	0.2	-1.7	0.1	0.0	1.0	0.6	0.0	-0.1	0.0	0.4	-0.5	
2tp	1.0	18.3	0.9	57.4	1.2	1.7	3.7	0.6	1.2	3.3	10.8	19176
no V	0.6	1.1	-0.5	7.8	0.9	-0.8	-4.2	0.1	-3.5	-0.1	-1.2	

A comparison of the collapsed classes with and without using VTEX2D shows that VTEX2D improves classification of clear air and reduces the false alarm rate. The biggest benefit to VTEX2D is improvement in the classification of aerofauna (insects, birds). However, it also reduces the accuracy of precipitation classification slightly. The small changes in accuracy may not be significant. The circumstances in which VTEX2D is most likely to exclude precipitation are when either there are small-scale variations in velocity associated with e.g. storms or tornadoes, or when the dealiasing contains errors, including from the dual-PRF dealiasing done on-site, which result in isolated discontinuities in velocity. For assimilation of radial velocity, excluding these artefacts is preferable.

Table 13 Collapsed view of the classification results for all Doppler training volumes, with and without using radial velocity texture.

		with VTEX2D					no VTEX2D		
		Automatic					Automatic		
		precip.	clear air	clutter			precip.	clear air	clutter
Manual	precip.	**92.7%**	3.1%	4.3%	Manual	precip.	**93.2%**	2.7%	4.1%
	clear air	12.7%	**71.5%**	15.7%		clear air	14.8%	**59.1%**	26.1%
	clutter	40.5%	8.1%	**51.4%**		clutter	41.8%	5.8%	**52.4%**

5.4 Radial velocity assimilation using clear air echo

The substantial availability of clear air echo on Australian Doppler radars encourages the exploration of assimilating radial velocity from clear air echo as well as from precipitation. This requires accurate identification of clear air with minimal contamination from clutter. From Table 12 we see that the identification accuracy of insect echo is 70% and the false alarm rate is fairly low. The detection of smoke is not so favourable, as it is rare, and more likely to be mistaken for insects if the smoke is constrained to the convective boundary layer, and mistaken for precipitation if it is a large fire.

Clear air would provide a large amount of additional wind observations, covering fine periods as well as during convective development, primarily in summer. Figure 27 show the number of clear air and precipitation observations that are available from several radars during a period in early 2013. Every night, there is a large increase in the number of clear air observations from the nocturnal insect bloom, and a smaller bloom of diurnal insects during the day. The number of observations from both clear air and precipitation is consistently higher than from precipitation alone. Some light weather echo is also mixed in with the clear air echo.

The classification values should be useful to select which observations to assimilate, and to inform how observations should be treated if clear air echo is handled separately from precipitation echo. The initial purpose of having the classification information is that the relative error of the observation sources can be assessed, by monitoring observation minus background statistics for example. Clear air echo assimilation needs to be assessed carefully to ensure that it will not damage the forecast in NWP.

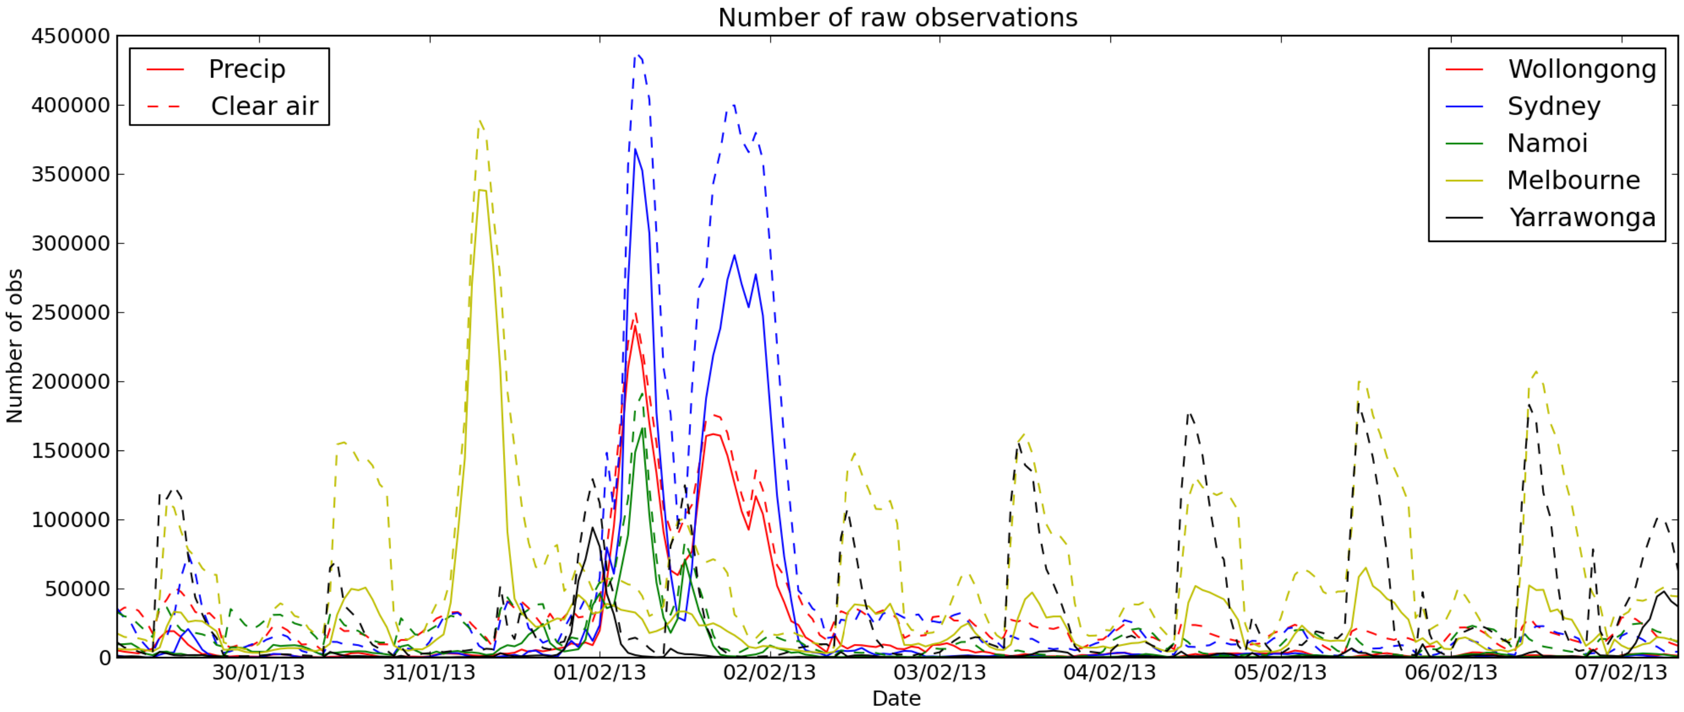

Fig. 27 The number of observations classified as precipitation (solid lines) or precipitation plus clear air (dashed lines) for a period during early 2013. The daily peaks represent the nocturnal insect bloom.

6 ADVANCES

As a first approach, the constraints when developing the classification system were to make it rely on as little external information as possible, i.e. that a scan can contain all the information needed to apply the quality control. This limits the dependencies of the radar software in case of malfunctions in other systems. However, there may be great benefit to using external information to aid in classification.

6.1 Probability of precipitation

Information on whether or not there is likely to be precipitation in the scan would be very helpful in excluding non-precipitation. If precipitation could be removed as a potential class then the scan must contain only clear air echo and various clutter types. This might not improve classification of cases where there are a mix of classes present, but will greatly improve classification where no precipitation is present, which benefits data assimilation applications.

One way to handle this is to reduce the prior probability for precipitation classes when the forecast probability of precipitation is low. Probability of precipitation forecasts may be available in advance of the scan time, and as an operational forecast product are likely to be available consistently.

To test the effect of using the probability of precipitation (PoP) to modify the prior probabilities, the PoP was acquired from Bureau's Operational Consensus Forecast. The PoP is in the format of a grid over the Australian region that gives the probability of precipitation over a given threshold in that grid space. The probability of precipitation over 0.2 mm (i.e. any measureable precipitation) was used. Values were interpolated to the radar locations for the training dataset, using a forecast lead time of 2–6 hours, presuming that the earliest a forecast may be available is 2 hours after analysis time. PoP forecasts are made for 3-hourly intervals, so for times within an hour of the validity time the value for that time was used. For scan times between two validity times, an average was used. The value at the radar location was presumed to be valid for the entire volume, since the grid spaces are quite large and probably cover much of the radar scan area.

The probability of precipitation values for the training dataset were extracted and stored. When the classification scheme was run over the dataset, the prior probabilities for the precipitation classes were dynamically altered using sed to edit the xml schema files. There were two different versions applied, one for the low reflectivity resolution radars and one for the rest. The system is as shown in Table 14.

Table 14 Modifications to prior probability based on probability of precipitation.

Probability of precipitation (p)	Prior probability		
	convective	shallow conv	stratiform
Most radars			
$p > 0.4$	0.4	0.25	0.5
$0.1 \leq p \leq 0.4$	0.2	0.12	0.2
$p < 0.1$	0.01	0.05	0
Low res radars			
$p > 0.4$	0.3	0.2	0.25
$0.1 \leq p \leq 0.4$	0.15	0.1	0.12
$p < 0.1$	0.02	0.05	0.02

The reasoning for the chosen prior probabilities is that all but one of the low resolution radars are in tropical regions where tropical convective precipitation is more likely. Shallow convection is also not well observed by these radars. If there is a low probability of precipitation, then it is more likely to be convective than stratiform precipitation because stratiform precipitation is well predicted and less stochastic, being forced by the large scale. For a very low probability of precipitation the prior probability of shallow convection is set slightly higher than the other precipitation classes. Note that the prior probability for precipitation of some variety is never zero, even if the PoP were zero.

The results from reclassifying the training dataset are shown in collapsed form in Table 15; full results are in Table 16 and the difference from the standard schema is in Table 17. The result of using PoP is a large improvement in the discrimination of clutter. Smoke and chaff particularly are better identified, as well as AP ground and sea clutter, which were responsible for the largest false alarm contaminations. In contrast, the precipitation identification is slightly reduced, which suggests the prior probabilities were set too low. The largest differences are for shallow convective and stratiform precipitation. This suggests these are the most difficult to separate from other classes.

Table 15 Collapsed class comparison of the standard classification versus using probability of precipitation to moderate the prior probabilities.

		Standard			Using probability of precipitation			
		Automatic				Automatic		
		precip.	clear air	clutter		precip.	clear air	clutter
Manual	precip.	**92.7%**	3.0%	4.3%	precip.	**90.9%**	3.9%	5.3%
Manual	clear air	12.8%	**71.5%**	15.7%	clear air	4.5%	**78.4%**	17.1%
Manual	clutter	36.9%	8.2%	**54.8%**	clutter	5.4%	16.5%	**78.1%**

Table 16 The classification result using probability of precipitation to modify the prior probability. Correct classifications are shown in bold.

man\aut	con	shc	str	ins	smk	chf	brd	pe	gc	ap	sl
con	**44.8**	**7.4**	**40.0**	1.6	3.1	0.5	0.1	0.2	1.7	0.4	0.1
shc	**3.2**	**51.2**	**5.7**	11.0	1.5	6.1	1.1	1.9	1.9	12.1	4.2
str	**24.5**	**8.1**	**59.5**	1.5	1.5	2.4	0.2	0.5	0.5	1.1	0.3
ins	0.2	1.2	0.2	**78.8**	2.4	0.1	6.5	5.9	3.6	0.2	0.8
smk	32.0	7.5	2.6	17.9	**25.1**	0.0	2.2	3.4	9.4	0.0	0.0
chf	0.1	6.8	0.0	8.7	10.7	**45.1**	1.7	1.4	2.9	20.6	1.9
brd	0.0	0.8	0.0	30.4	0.9	0.1	**51.6**	7.5	5.8	0.0	2.8
pe	0.1	1.6	0.0	12.8	2.4	0.0	8.6	**69.4**	4.1	0.0	0.8
gc	0.4	4.7	0.0	25.7	5.3	2.3	7.5	6.4	**47.4**	0.2	0.1
ap	0.4	3.3	0.7	0.0	9.5	15.6	0.0	0.0	0.0	**64.1**	6.3
sl	0.1	3.5	0.3	0.2	4.8	2.2	0.0	0.2	0.0	2.2	**86.4**
2tp	0.8	10.3	0.7	57.4	2.1	4.1	3.7	0.6	1.2	7.5	11.7

Table 17 The difference between Table 16 and the results of the standard scheme (Table 7). A negative value means fewer pixels were classed. Large improvements are marked in green and large reductions in accuracy are marked in red.

man\aut	con	shc	str	ins	smk	chf	brd	pe	gc	ap	sl
con	0.3	-0.3	-0.9	0.2	0.9	0.2	0.0	0.0	0.1	0.1	0.0
shc	-0.1	-15.7	-0.4	0.7	1.0	3.0	0.0	0.1	0.5	8.9	2.1
str	-0.2	-0.6	-0.3	0.3	0.5	0.0	0.0	0.0	0.1	0.2	0.0
ins	-0.7	-5.0	-1.3	5.2	0.9	0.1	0.0	0.0	0.6	0.1	0.1
smk	-10.5	-8.1	-5.7	3.7	13.4	0.0	0.0	0.0	7.2	0.0	0.0
chf	-5.8	-13.0	-29.1	1.7	10.2	18.6	0.1	0.0	2.4	14.6	0.3
brd	0.0	-2.7	-0.9	2.7	0.1	0.0	0.2	0.0	0.5	0.0	0.0
pe	-1.7	-4.7	-0.4	2.7	1.6	0.0	0.0	0.8	1.5	0.0	0.1
gc	-2.5	-23.8	-6.0	5.9	3.7	0.9	0.0	0.2	21.4	0.2	0.0
ap	-5.6	-13.9	-16.9	0.0	8.8	4.0	0.0	0.0	0.0	22.4	1.3
sl	-0.5	-3.7	-0.5	0.0	1.4	0.4	0.0	0.0	0.0	0.8	2.0
2tp	-0.2	-8.1	-0.2	0.0	1.0	2.3	0.0	0.0	0.0	4.2	0.9

The principal detriments of PoP on precipitation detection is that rainfall on days with low PoP, which is primarily shallow convective precipitation, is less well detected. Naturally an overall decrease in the detected precipitation would result from the prior probability being decreased compared to the standard scheme, and mostly this change is small. The increased classification of chaff as AP is not an issue. Overall there are large improvements in classification of clear air and AP clutter echoes, and the misclassification of these as precipitation is much reduced as a result. Analogously, predictors of other echo types might also improve classification. For example if conditions favourable for ducting were forecast, then the prior probability for AP could be increased. The classification of temporally predictable echoes such as those from dusk dispersals of aerofauna may also be improved by modulated prior probabilities. However, the classifier is not presently designed to incorporate temporally-varying information.

6.2 Continuity

Theoretically, improved results could be achieved by looking for temporal or spatial continuity of classification. For example, if there is a region of AP sea clutter correctly identified, then it is likely that echo nearby or beyond this will also contain clutter. Likewise a region of precipitation is likely to be not interrupted by smoke or insects or AP clutter. Some experimentation was done to spread classifications according to logic. However, the fundamental problem is that it is not possible to know what has been classified correctly in the first place, without some external information. One option is to replace isolated pixels with the surrounding class, e.g. a pixel of sea clutter surrounded on four sides by precipitation is probably precipitation. Testing this on the training dataset yielded little result as this affects few pixels overall. Continuity assurances propagate errors as easily as they propagate correct results.

A tool which might be of more use is a classification confidence that reflects the confidence of the Bayes classification result. A strength of classification index (Peter *et al.* 2014) has been developed for a two-class classifier by comparing the probability of a precipitation classification to a sea clutter classification. An analogous version may be developed to show how likely the classifier thought a pixel was its first class choice, compared to its second class choice. If the most certain classifications were more accurate, these could be used to 'correct' nearby pixels. However, this type of classification confidence indicator is not currently available and so has not been tested. To use such an index would require capturing several additional pieces of information and then deciding how these should be managed, which requires non-trivial changes to the system as well as a review of how the information could be useful.

7 CONCLUSIONS

A naïve Bayesian classification method has been developed to differentiate a variety of echo types including precipitation, various clear air, ground and sea clutter echoes. The classifier function is implemented for single polarisation radars and can operate on reflectivity-only radars; however, it performs better on Doppler radars where radial velocity and spectrum width are available. Theoretically, the classifier could be employed with dual-polarisation radars where the variety of polarimetric parameters would produce success comparable to dual-polarisation classification methods implemented elsewhere.

The classifier required training via a manually-classed dataset, which is used to produce the probability distribution functions for each class and feature field used by the classifier. This is a drawback as the production of a training dataset is laborious and time-consuming. However, once this is done the classifier has the potential to be a powerful tool in quality control of radar data. The version developed here examined twelve radar echo classes that are commonly seen with Australian radars: deep and shallow convective precipitation, stratiform precipitation, insects, birds/bats, smoke, chaff, permanent and AP ground clutter, side-lobe and AP sea clutter, and second-trip echo. The feature fields included reflectivity, texture and spin of reflectivity with different kernel sizes, echo top height using two thresholds, and vertical gradient of reflectivity. Additionally, texture of radial velocity and spectrum width were used when available. For each class and each feature field, a pdf was empirically fit from a range of standard and composite pdfs. The pdf parameters were entered into the classifier. Prior probabilities for each class were also generated, based on their frequency of occurrence, though rare echo types may have had a higher than climatological probability to ensure that they would be classed at all.

The classifier is not intended to accurately distinguish all these classes, since this is very difficult to do with single polarisation radar. It is not meant to discriminate between stratiform and convective precipitation, since there are existing methods that are used already. The classifier need only distinguish between precipitation and non-precipitation, or for some applications between precipitation, clear air and other echoes. For this purpose, using a single polarisation radar, the method is adequate. Average results from classifying the training dataset are that 93% of precipitation is correctly classified, and 19% of all non-precipitation is falsely classified as precipitation. Some further tuning of the prior probabilities could improve this result, depending on what application will use the classification information, and whether it is more important to capture all precipitation or to have a low rate of clutter contamination in the selected observations.

There are ways that the classification result could be improved, for example by informing the classifier when there is not rainfall expected and hence reducing the prior probability when the forecast probability of precipitation is low. This method substantially reduces the false alarm rate as anticipated; however, it does not improve the rate of correctly identifying precipitation. By implementing the ability to change the prior probabilities over time other external information, such as reports of fires or chaff, or conditions causing anomalous propagation, could be similarly used to improve the classification result.

REFERENCES

Achtemeier, G.L. 1991: The use of insects as tracers for "clear-air" boundary-layer studies by Doppler radar. J. *Atmos. Oceanic. Tech.*, 8, 746-765.

Australian Bureau of Meteorology: PC Based Rapic Transmitter Digital Receiver Option. Weather Watch Radar Software Documentation, 94 pp.

Bachmann, S. and Zrnić, D. 2008: Three-dimensional attributes of clear-air scatterers observed with the polarimetric weather radar. IEEE Geosci. Remote S., 5, 231-235.

Dixon, M., Kessinger, C.andHubbert, J. 2005: Echo classification and spectral processing for the discrimination of clutter from weather. 32nd International Conference on Radar Meteorology, Albuquerque, NM, *American Meteorological Society*, P4R.6.

Fabry, F. 2010: Radial velocity measurement simulations: common errors, approximations, or omissions and their impact on estimation accuracy. ERAD 2010 - The Sixth European Conference on Radar in Meteorology and Hydrology, Sibiu, Romania, 138-144.

Friedman, N., Geiger, D. and Goldszmidt, M. 1997: Bayesian Network Classifiers. *Machine Learning*, 29, 131-163.

Geerts, B. and Miao, Q. 2005: Airborne radar observations of the flight behavior of small insects in the atmospheric convective boundary layer. *Quant. Ecol*., 34, 361-377.

Hubbert, J.C., Dixon, M. and Ellis, S.M. 2009: Weather radar ground clutter. Part II: Real-time identification and filtering. *J. Atmos. Oceanic. Tech*., 26, 1181-1197.

Izenmann, A.J. 1991: Recent developments in nonparametric density estimation. *Journal of the American Statistical Association*, 86, 205-225.

Kessinger, C., Ellis, S., Van Andel, J., Yee, J. and Hubbert, J. 2004: *The AP ground clutter mitigation scheme for the WSR-88D*. 20th International Conference on Interactive Information Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology.

Koistinen, J. 2000: Bird migration patterns on weather radars. *Phys. Chem. Earth Pt. B*, 25, 1185-1193.

Koistinen, J., Mäkinen, T. and Pulkkinen, S. 2009: *Classification of meteorological and non-meteorological targets with principal component analysis applying conventional polarimetric parameters and their texture*. AMS 34th Conference on Radar Meteorology, Williamsburg VA, USA, 10A.13.

Lakshmanan, V., Fritz, A., Smith, T., Hondl, K. and Stumpf, G. 2007: An Automated Technique to Quality Control Radar Reflectivity Data. *J. Appl. Meteorol. Clim*., 46, 288-305.

Moszkowicz, S., Ciach, G.J. and Krajewski, W.F. 1994: Statistical Detection of Anomalous Propagation in Radar Reflectivity Patterns. *J. Atmos. Oceanic. Tech*., 11, 1026-1034.

Peter, J.R., Seed, A. and Steinle, P. 2013: Application of a Bayesian Classifier of Anomalous Propagation to Single-Polarization Radar Reflectivity Data. *J. Atmos. Oceanic. Tech*., 30, 1985-2005.

Peter, J.R., Seed, A., Steinle, P., Rennie, S.J. and Curtis, M. 2014: A Bayesian methodology for detecting anomalous propagation in radar reflectivity observations. *CAWCR Technical Reports* in press.

Rennie, S. J. 2012: Doppler weather radar in Australia. C*AWCR Technical Reports 55*, 52 pp.

Rennie, S.J. 2013: Common orientation and layering of migrating insects in south-eastern Australia observed with a Doppler weather radar. *Meteorol. Appl*., in press.

Rico-Ramirez, M.A. and Cluckie, I.D. 2008: Classification of Ground Clutter and Anomalous Propagation Using Dual-Polarization Weather Radar. IEEE T. *Geosci. Remote*, 46, 1892-1904.

Schuur, T., Ryzhkov, A. and Heinselman, P. 2003: Observations and classification of echoes with the polarimetric WSR-88D radar, 45 pp.

Steiner, M. and Smith, J.A.2002: Use of three-dimensional reflectivity structure for automated detection and removal of nonprecipitating echoes in radar data. *J. Atmos. Oceanic. Tech*., 19, 673-686.

Steiner, M., Houze, R.A.and Yuter, S.E.1995: Climatological characterization of three-dimensional storm structure from operational radar and rain gauge data. *J. Appl. Meteorol*., 34, 1978-2007.

Wilson, J., Weckwerth, T.M., Vivekanandan, J., Wakimoto, R.M. and Russell, R.W.1994: Boundary layer clear-air radar echoes: Origin of echoes and accuracy of derived winds. *J. Atmos. Oceanic. Tech*., 11, 1184-1206.

APPENDIX A

The training dataset. Table 18 lists the file names sorted by radar and date, what classes were classified in each and the number of classified pixels for each class. Table 19 contains the number of radars, days and volumes classed for each class.

Table 18 The number of pixels classified for each class in each file of the training dataset. The files are ordered by radar ID number and date.

RadarID/date_time as yyyymmdd_hhmmss	conv	sh con	strat	insects	smoke	chaff	birds	perm echo	AP GC	AP sea clutter	SL sea clutter	2nd trip
14/20121009_042001	0	0	196863	0	0	0	0	0	0	0	0	0
16/20120624_115001	0	0	0	0	0	0	0	0	24404	8185	3348	0
16/20120624_124001	0	0	0	0	0	0	0	0	45436	17795	3156	0
16/20120724_102005	0	0	0	0	0	0	0	0	40026	4046	1232	0
16/20120724_112005	0	0	0	0	0	0	0	0	57215	3091	1987	0
16/20120724_125005	0	0	0	0	0	0	0	0	49165	2493	8840	0
16/20120724_141004	0	0	0	0	0	0	0	0	49427	0	4233	0
16/20120724_171005	0	0	0	0	0	0	0	867	68539	7165	7995	0
19/20130409_000134	35726	0	0	0	0	0	0	0	0	0	0	0
2/20090207_080002	0	0	0	0	120171	0	0	0	0	0	0	0
2/20090207_083002	0	0	0	0	125390	0	0	0	0	0	0	0
2/20090207_090002	0	0	0	0	103964	0	0	0	0	0	0	0
2/20110915_044220	0	0	0	0	0	0	467	477	0	0	0	0
2/20110915_044820	0	0	0	0	0	0	542	201	0	0	0	0
2/20110915_050020	0	0	0	0	0	0	665	183	0	0	0	0
2/20110915_050620	0	0	0	0	0	0	543	0	0	0	0	0
2/20120224_110034	0	0	0	469450	0	0	0	380	0	0	0	0
2/20120226_070040	195468	0	0	0	0	0	0	1013	0	0	0	0
2/20120226_100020	291198	0	0	0	0	0	0	0	0	0	0	0
2/20120303_030021	0	0	1507814	0	0	0	0	0	0	0	0	0
2/20121109_023638	0	0	129850	0	0	0	0	0	0	0	0	0
2/20121126_220025	0	0	951575	0	0	0	0	0	0	0	0	0
2/20121128_134229	0	0	0	0	0	0	0	5099	143	0	0	0
2/20121128_141240	0	0	0	0	0	0	0	5154	163	0	0	0
28/20130902_042205	0	0	0	0	0	10184	0	0	0	0	0	0
28/20130902_060205	0	0	0	0	0	9955	0	0	0	0	0	0
29/20120723_175001	0	0	0	0	0	0	0	22	634	0	35283	0

RadarID/date_time as yyyymmdd_hhmmss	conv	sh con	strat	insects	smoke	chaff	birds	perm echo	AP GC	AP sea clutter	SL sea clutter	2nd trip
29/20120723_180002	0	0	0	0	0	0	0	0	847	0	33700	0
29/20121004_015001	0	0	0	0	0	0	0	0	0	0	46121	0
29/20121008_140001	0	0	0	0	0	0	0	0	0	115955	53912	0
29/20121008_144002	0	0	0	0	0	0	0	0	0	97004	54486	0
3/20120103_030005	0	0	0	14292	0	0	0	0	0	0	971	0
3/20120103_060005	0	0	0	0	0	0	0	0	0	8547	0	0
3/20120103_095005	0	0	0	0	0	0	2214	0	0	8677	935	0
3/20120118_094005	0	0	0	0	0	0	607	0	0	0	0	0
3/20120118_095005	0	0	0	0	0	0	2026	0	0	0	0	0
3/20120118_100005	0	0	0	0	0	0	2371	0	0	0	0	0
3/20120118_101005	0	0	0	0	0	0	1674	0	0	0	0	0
3/20120118_175005	0	0	0	58545	0	0	0	0	0	0	0	0
3/20120120_060005	39017	0	0	0	0	0	0	0	0	0	0	0
3/20120120_095005	77689	0	0	18520	0	0	1963	0	0	0	0	0
3/20120206_034005	5900	0	0	0	0	0	0	0	0	0	1158	0
3/20120206_040005	5034	0	0	0	0	0	0	0	0	0	1354	0
3/20120207_053005	6236	0	0	0	0	0	0	0	0	0	292	0
3/20120207_060005	7761	0	0	0	0	0	0	0	0	0	250	0
3/20120301_100034	0	0	437805	0	0	0	0	0	0	0	0	0
3/20120301_201035	0	0	126971	0	0	0	0	0	0	0	0	0
3/20120301_213041	0	0	104313	0	0	0	0	0	0	0	0	0
3/20120301_231035	0	0	198437	0	0	0	0	0	0	0	0	0
3/20120301_233044	0	0	261506	0	0	0	0	0	0	0	0	0
3/20120315_053041	2130	0	0	0	0	0	0	0	0	0	0	0
3/20120315_064034	2197	0	0	0	0	0	0	0	0	606	0	0
3/20120315_071034	2818	0	0	0	0	0	0	0	0	450	0	0
3/20120315_074035	3742	0	0	0	0	0	0	0	0	50	0	0
3/20120316_042034	22338	0	0	0	0	0	0	0	0	52	0	0
3/20120321_033040	0	0	0	0	0	56	0	0	0	0	253	0
3/20120321_040035	0	0	0	0	0	48	0	0	0	0	322	0
3/20120810_091006	0	0	15956	0	3718	0	0	1860	0	0	7736	0
3/20120810_131006	0	0	0	0	4085	0	0	0	0	0	3517	0
3/20120810_151006	0	0	0	0	3478	0	0	0	0	0	4147	0
3/20120810_173005	0	0	25230	0	2880	0	0	0	0	0	0	0
3/20120810_234007	0	0	0	0	3572	0	0	0	0	0	5741	0

RadarID/date_time as yyyymmdd_hhmmss	conv	sh con	strat	insects	smoke	chaff	birds	perm echo	AP GC	AP sea clutter	SL sea clutter	2nd trip
3/20120811_003005	0	0	0	0	3835	0	0	359	0	0	2654	0
3/20120827_132006	0	0	0	28907	0	0	0	0	0	0	1115	0
3/20120828_112005	0	0	0	0	0	942	0	0	0	0	1764	0
3/20120828_120005	0	0	0	0	0	1455	0	0	0	0	2171	0
3/20120828_123006	0	0	0	0	0	705	0	0	0	0	1991	0
3/20120828_133005	0	0	0	0	0	0	0	0	0	0	2986	0
3/20120829_131006	0	0	0	0	0	0	0	0	0	0	13167	0
3/20120829_142006	0	0	0	0	0	0	0	0	0	0	9500	0
3/20120829_233006	0	0	0	0	0	872	0	0	0	0	0	0
3/20121120_002008	0	0	0	0	0	5628	0	0	0	0	0	0
3/20121120_004004	0	0	0	0	0	6200	0	0	0	0	0	0
3/20121120_010005	0	0	0	0	0	7677	0	0	0	0	0	0
3/20121120_012007	0	0	0	0	0	8808	0	0	0	0	0	0
3/20121120_014005	0	0	0	0	0	10772	0	0	0	0	0	0
3/20121120_022007	0	0	0	0	0	12734	0	0	0	0	0	0
3/20121120_025005	0	0	0	0	0	13379	0	0	0	0	0	0
3/20121120_033005	0	0	0	0	0	13872	0	0	0	0	0	0
3/20121120_035005	0	0	0	0	0	14642	0	0	0	0	0	0
3/20121120_041005	0	0	0	0	0	12179	0	0	0	0	0	0
3/20121120_043005	0	0	0	0	0	13767	0	0	0	0	0	0
3/20121120_050005	0	0	0	0	0	16041	0	0	0	0	0	0
3/20130625_210040	0	80495	0	0	0	0	0	0	0	0	0	0
4/20130625_210308	0	76834	0	0	0	0	0	0	0	0	0	0
4/20130902_023034	0	0	0	0	0	43091	0	0	0	0	0	0
4/20130902_034835	0	0	0	0	0	56926	0	0	0	0	0	0
4/20130902_052434	0	0	0	0	0	33237	0	0	0	0	0	0
49/20090207_102203	0	0	0	0	2573	0	0	0	0	0	0	0
49/20090207_120203	0	0	0	0	3888	0	0	0	0	0	0	0
49/20090207_140203	0	0	0	0	13636	0	0	0	0	0	0	0
49/20090207_150203	0	0	0	0	9729	0	0	0	0	0	0	0
49/20120312_015219	0	0	0	391119	0	0	0	0	0	0	0	0
49/20120312_061219	0	0	0	303118	0	0	0	0	0	0	0	0
49/20120312_094220	0	0	2	505686	0	0	0	0	0	0	0	0
49/20120312_100220	0	0	0	505847	0	0	0	0	0	0	0	0
49/20120312_200220	0	0	0	308857	0	0	1067	901	0	0	0	0

RadarID/date_time as yyyymmdd_hhmmss	conv	sh con	strat	insects	smoke	chaff	birds	perm echo	AP GC	AP sea clutter	SL sea clutter	2nd trip
49/20120315_053219	5299	0	0	176120	0	0	0	866	0	0	0	0
49/20120315_094219	0	0	108775	484155	0	0	0	0	0	0	0	0
49/20120315_103220	0	0	196116	469643	0	0	0	343	0	0	0	0
49/20120315_143219	0	0	1081000	0	0	0	0	0	0	0	0	0
50/20120920_213142	45188	0	0	0	0	0	0	0	0	0	0	0
50/20130902_052152	0	0	0	0	0	5790	0	0	0	0	0	0
50/20130902_071152	0	0	0	0	0	15861	0	0	0	0	0	0
52/20121009_042107	0	0	6416	0	0	0	0	0	0	0	0	0
54/20120103_060001	0	0	0	7039	0	0	0	0	0	19943	28996	0
54/20120103_082401	0	0	0	2022	0	0	0	0	0	16662	24670	0
54/20120810_091202	0	0	0	0	3520	0	0	7830	0	0	41696	0
54/20120810_131804	0	0	0	0	4407	0	0	4434	0	0	41253	0
54/20120810_151202	0	0	0	0	6226	0	0	1573	0	0	56366	0
54/20120810_173002	0	0	0	0	1770	0	0	814	0	0	1642	0
54/20120810_233602	0	0	0	0	1167	0	0	1152	94	0	32932	0
54/20120811_003002	0	0	0	0	2310	0	0	2124	108	0	19293	0
54/20121120_004202	0	0	0	0	0	1777	0	0	0	0	0	1
54/20121120_015402	0	0	0	0	0	2130	0	0	0	0	0	1
54/20121120_042402	0	0	0	0	0	1504	0	0	0	0	0	0
54/20121120_053001	0	0	0	0	0	914	0	0	0	0	0	0
54/20121120_060002	0	0	0	0	0	523	0	0	0	0	0	0
54/20130626_020001	0	27112	0	0	0	0	0	0	0	0	0	0
63/20120313_064002	1101286	0	0	0	0	0	0	0	0	0	0	0
63/20120314_103002	766937	0	0	3	0	0	0	0	0	0	0	0
63/20130412_120002	0	0	0	0	0	0	0	0	0	0	7750	0
64/20121004_015109	0	0	0	134744	0	0	0	10694	6014	0	0	0
64/20130115_000105	0	0	0	0	0	0	0	0	0	0	4865	0
64/20130117_220102	0	0	0	0	0	0	0	0	0	0	28310	0
64/20130204_100102	0	0	0	0	0	0	0	0	0	0	24951	0
64/20130409_140103	0	0	0	0	0	0	0	7966	1070	6851	9041	0
66/20121004_011241	0	0	0	69319	0	0	0	0	8384	0	0	0
66/20130902_050623	0	0	0	0	0	103657	0	0	0	0	0	0
66/20130902_063623	0	0	0	0	0	363174	0	0	0	0	0	0
66/20130902_080034	0	0	0	0	0	490967	0	0	0	0	0	0
68/20130625_210009	0	28439	0	0	0	0	0	0	0	0	0	0

RadarID/date_time as yyyymmdd_hhmmss	conv	sh con	strat	insects	smoke	chaff	birds	perm echo	AP GC	AP sea clutter	SL sea clutter	2nd trip
68/20130626_000003	0	30194	0	0	0	0	0	0	0	0	0	0
69/20120312_020038	0	0	0	93532	0	0	0	0	0	0	0	0
69/20120312_090036	0	0	0	138496	0	0	0	0	0	0	0	0
69/20120312_103037	0	0	0	182936	0	0	0	0	0	0	0	0
69/20120312_225037	0	0	0	67082	0	0	0	0	0	0	0	0
69/20121004_015002	0	0	0	69615	0	0	0	0	0	0	0	0
69/20121031_053003	0	0	0	53158	0	0	0	0	0	0	0	0
69/20121031_101003	0	0	0	148140	0	0	0	0	0	0	0	0
69/20121106_102003	5859	0	0	152989	0	0	0	0	0	0	0	0
69/20121106_110003	5203	0	0	148868	0	0	0	0	0	0	0	0
70/20120319_074001	0	0	0	6678	0	4538	0	1430	0	0	0	0
70/20120319_075004	0	0	0	6173	0	4783	0	864	0	0	0	0
70/20120319_082002	0	0	0	5501	0	4599	0	1432	0	0	0	0
70/20130625_210001	0	15399	0	0	0	0	0	0	0	0	0	0
70/20130626_000002	0	29925	0	0	0	0	0	0	0	0	0	0
71/20120102_201904	0	0	0	0	0	0	0	3081	8172	0	0	0
71/20120103_022504	0	0	0	87980	0	0	0	0	0	3244	0	0
71/20120103_051904	0	0	0	113477	0	0	0	0	0	24552	4410	0
71/20120103_094904	0	0	0	81747	0	0	5141	621	1909	21393	5374	0
71/20120103_163704	0	0	0	208251	0	0	783	1595	10289	0	0	0
71/20120229_080712	0	0	1106132	0	0	0	0	0	0	0	0	0
71/20120303_214305	0	0	56012	91008	0	0	0	0	0	0	3158	8526
71/20120303_231902	0	0	69602	134640	0	0	0	0	0	0	2078	4849
71/20120304_012503	0	0	98695	162333	0	0	0	0	0	0	3325	0
71/20120308_040103	0	0	239724	0	0	0	0	0	0	0	0	0
71/20120312_085502	0	0	0	13504	0	0	985	0	0	0	1116	0
71/20120312_192502	0	0	0	0	0	0	2545	0	0	0	0	0
71/20120313_085504	0	0	0	0	0	0	1408	0	0	0	346	0
71/20120313_192504	0	0	0	0	0	0	2594	0	0	0	0	0
71/20120314_064305	0	0	0	0	0	0	0	0	0	0	0	3395
71/20120314_085502	0	0	0	0	0	0	1915	0	0	0	4649	2105
71/20120314_193103	0	0	0	0	0	0	2924	0	0	0	1697	537
71/20120321_030102	0	0	12647	72580	0	2078	0	317	0	0	4077	0
71/20120321_030704	0	0	0	88488	0	3356	0	263	0	0	3851	0
71/20120321_032502	0	0	0	71288	0	4293	0	123	0	0	3391	0

RadarID/date_time as yyyymmdd_hhmmss	conv	sh con	strat	insects	smoke	chaff	birds	perm echo	AP GC	AP sea clutter	SL sea clutter	2nd trip
71/20120321_034302	0	0	0	70657	0	4391	0	81	0	0	3950	0
71/20120321_040103	0	0	0	55031	0	5147	0	0	0	0	4191	0
71/20120321_041303	0	0	0	43283	0	5090	0	354	0	0	4985	0
71/20120321_043103	0	0	0	31287	0	4555	0	226	0	0	4436	0
71/20120321_044303	0	0	0	23044	0	4166	0	0	0	0	3605	0
71/20120322_050706	72145	0	0	0	0	0	0	0	0	0	0	0
71/20120323_005506	0	0	0	95196	0	0	0	47	1489	0	0	0
71/20120520_080704	32346	0	0	43183	0	0	0	466	0	0	0	0
71/20120605_044302	0	0	565842	0	0	0	0	0	0	0	0	0
71/20120607_011305	44478	0	0	24386	0	0	0	0	0	0	25515	0
71/20120810_091304	0	0	0	0	23352	0	0	2104	0	0	0	0
71/20120810_131905	0	0	0	0	7988	0	0	1184	0	0	0	0
71/20120810_150703	0	0	0	0	17419	0	0	2098	0	0	0	0
71/20120810_173103	0	0	64084	0	10351	0	0	2315	0	0	0	0
71/20120810_233705	0	0	0	0	8021	0	0	173	282	0	0	0
71/20120811_003102	0	0	0	0	16092	0	0	3078	0	0	0	0
71/20130626_000106	0	137844	0	0	0	0	0	0	0	0	0	0
72/20121009_041314	1223	0	0	5714	0	0	3068	708	0	0	0	0
72/20121102_093313	0	0	0	305107	0	0	0	0	0	0	0	0
72/20121103_094312	18502	0	0	271602	0	0	0	0	0	0	0	0
72/20121103_101314	22598	0	0	269151	0	0	0	971	0	0	0	0
72/20121103_103313	22675	0	0	263461	0	0	0	904	0	0	0	0
72/20121106_100313	0	0	0	128390	0	0	0	0	0	0	0	0
75/20120920_043630	0	0	0	11489	24232	0	0	0	0	0	0	0
75/20120920_052430	0	0	0	0	33216	0	0	0	0	0	0	0
75/20120920_083630	0	0	0	37995	36597	0	0	0	0	0	0	0
75/20120920_092430	0	0	0	47396	26086	0	8323	0	0	0	0	0
75/20120920_093630	0	0	0	8077	16424	0	5960	0	0	0	0	0
75/20120920_100030	0	0	0	275299	1843	0	0	0	0	0	0	0
75/20120920_120030	0	0	0	240407	3225	0	0	0	0	0	0	0
76/20120514_040659	0	0	57292	0	0	0	0	4940	0	0	13345	0
76/20120724_064233	0	0	0	0	0	0	0	9604	0	0	0	0
76/20120724_070032	0	0	0	0	0	0	0	9602	0	0	0	0
76/20120724_115432	0	0	0	0	0	0	0	10856	0	0	0	0
76/20120725_001233	0	0	19395	0	0	0	0	23015	0	0	0	0

RadarID/date_time as yyyymmdd_hhmmss	conv	sh con	strat	insects	smoke	chaff	birds	perm echo	AP GC	AP sea clutter	SL sea clutter	2nd trip
76/20130626_000032	0	76790	0	0	0	0	0	0	0	0	0	0
8/20121030_162253	12032	0	0	3284	0	0	0	2001	0	0	0	0
8/20121102_015252	0	0	0	0	0	0	0	6346	0	0	0	0

Table 19 The number of pixels, number of radar volumes, number of radars and number of days that contributed to each class in the training dataset.

	conv	shal con	strat	insects	smoke	chaff	birds	perm echo	AP GC	AP sea clutter	SL sea clutter	2nd trip
num pixels	2853025	503032	7638054	8399289	645165	1326463	49785	145111	373810	366761	735581	19414
num vols	31	9	26	61	32	44	22	53	21	20	64	5
num radars	10	7	7	12	6	8	6	12	7	6	8	1
num days	18	9	17	29	9	11	10	29	12	9	30	2

APPENDIX B

The figures show the fitted and raw pdfs for every feature field. Table 20 contains the typical pdf selected to represent each feature field.

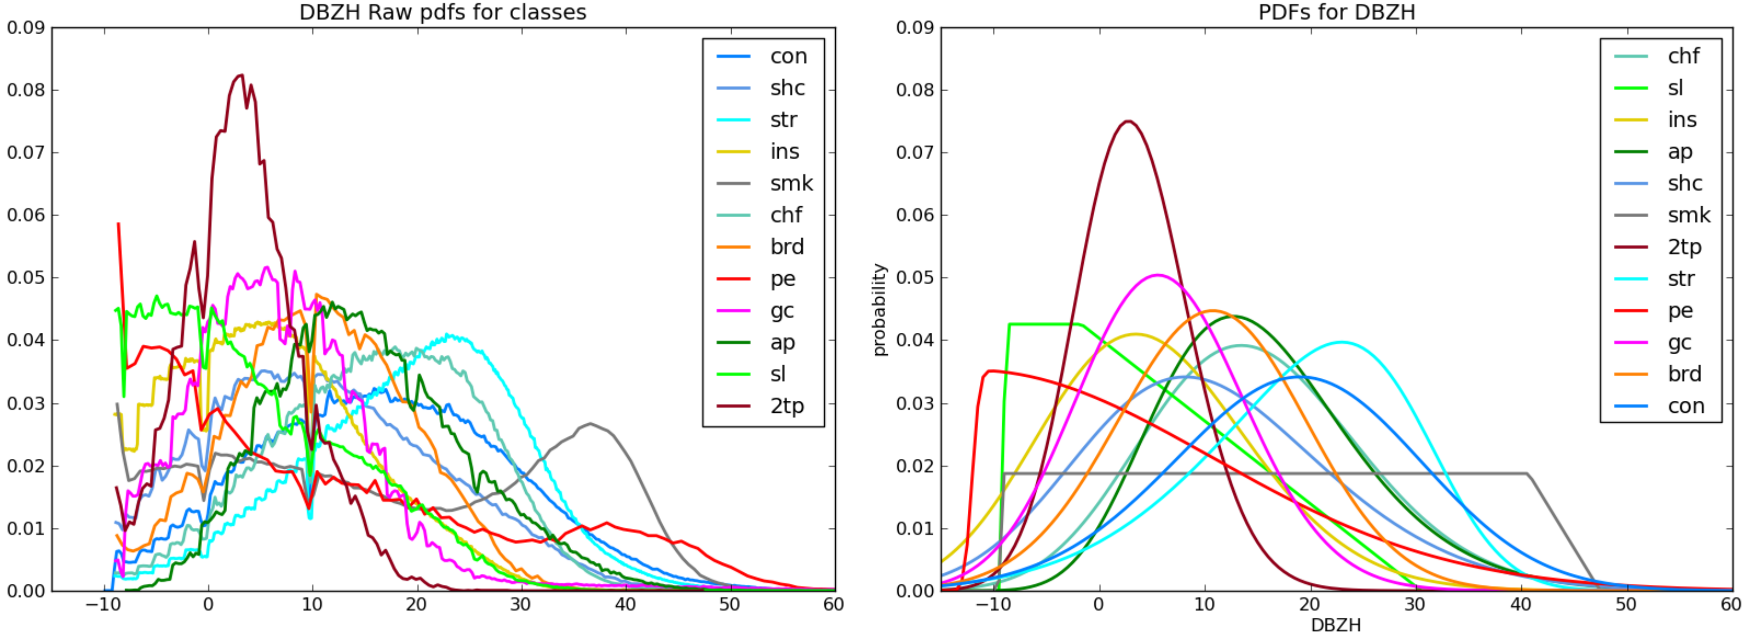

Fig. 28 Raw and fitted pdfs for reflectivity (DBZH).

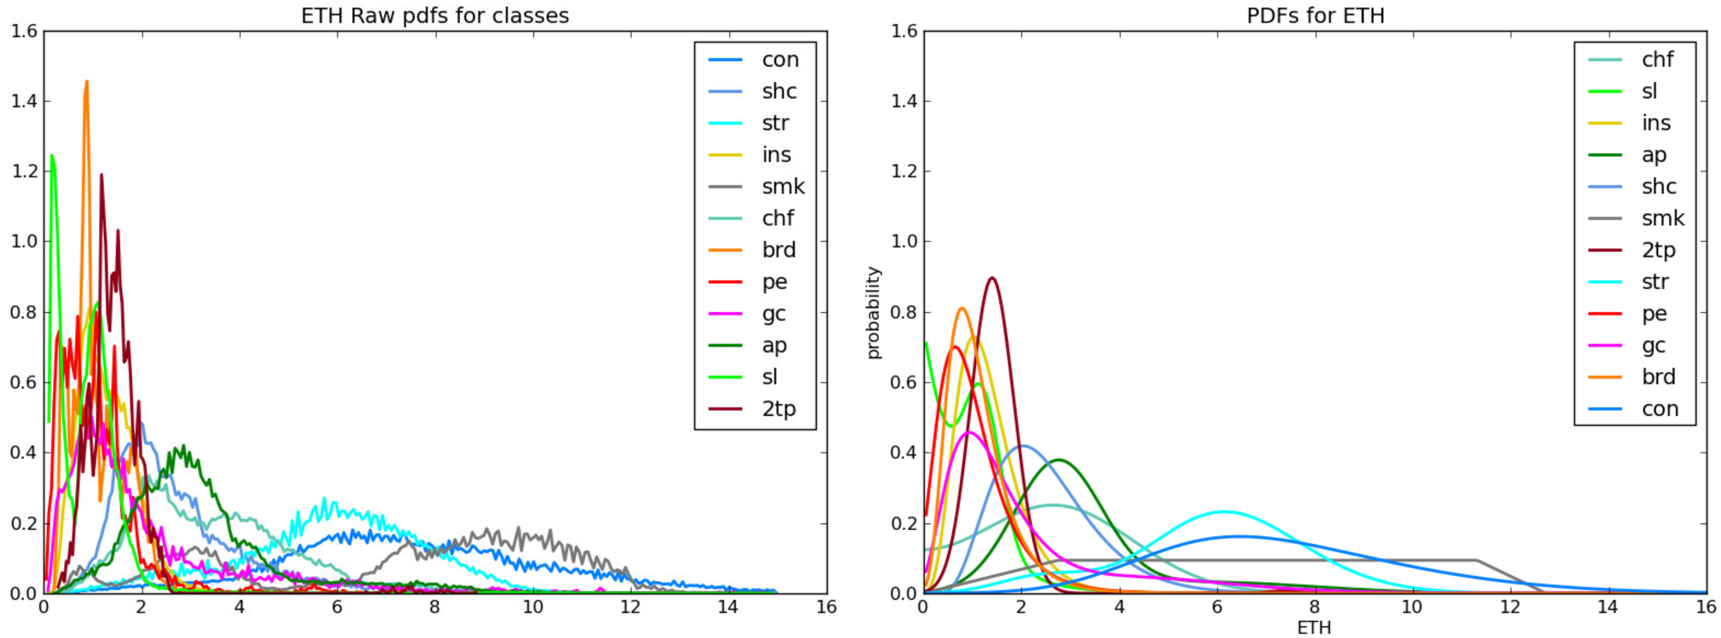

Fig. 29 Raw and fitted pdfs for echo top height using a 4 dB threshold (ETH).

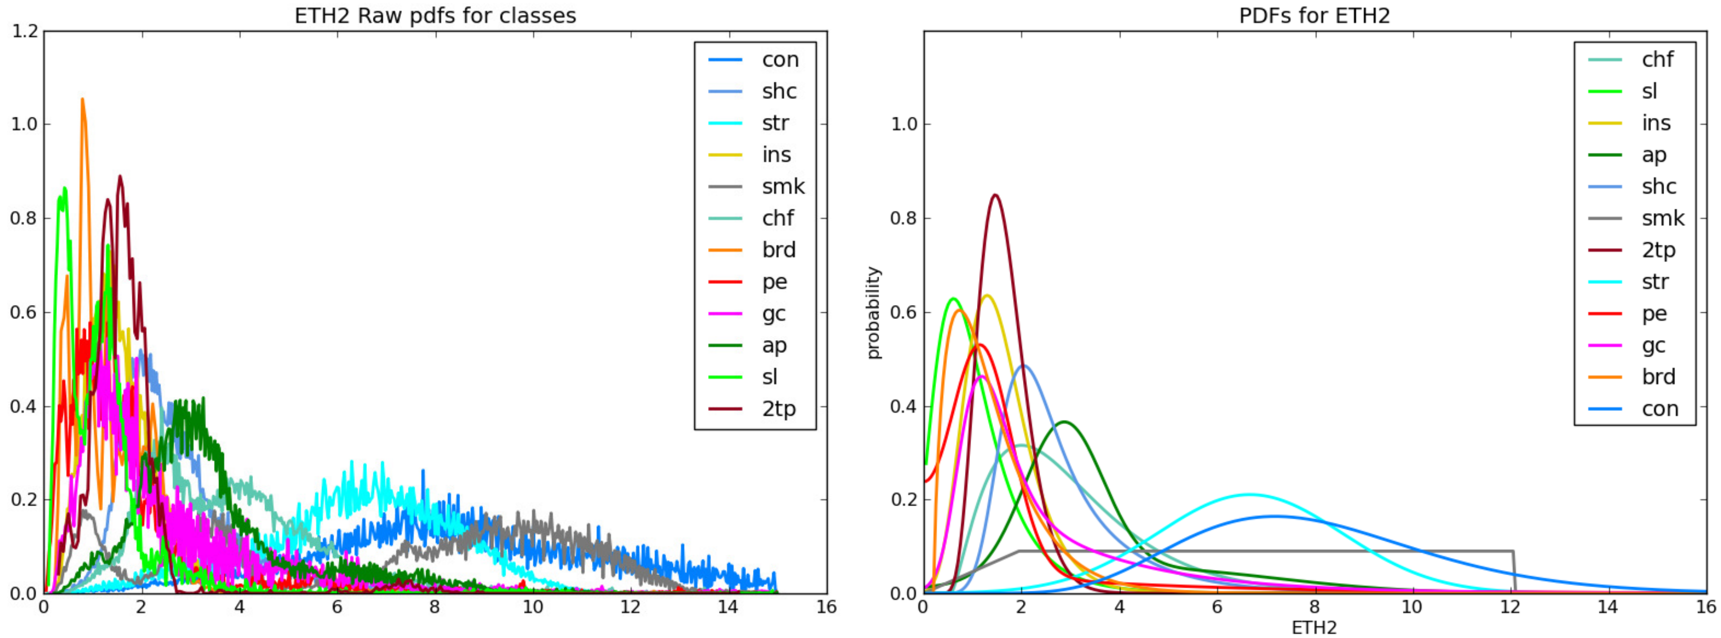

Fig. 30 Raw and fitted pdfs for echo top height using a –5 dB threshold (ETH2).

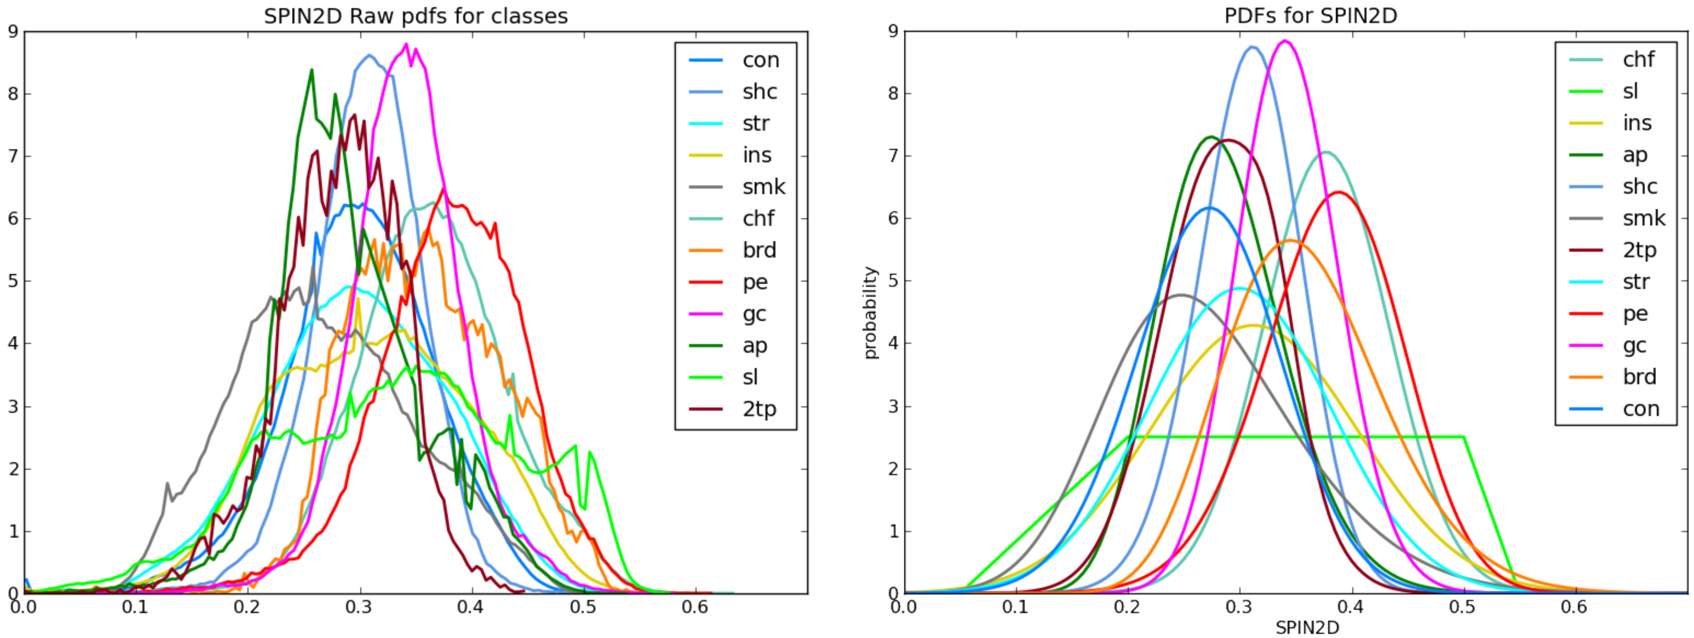

Fig. 31 Raw and fitted pdfs for spin of reflectivity (SPIN2D).

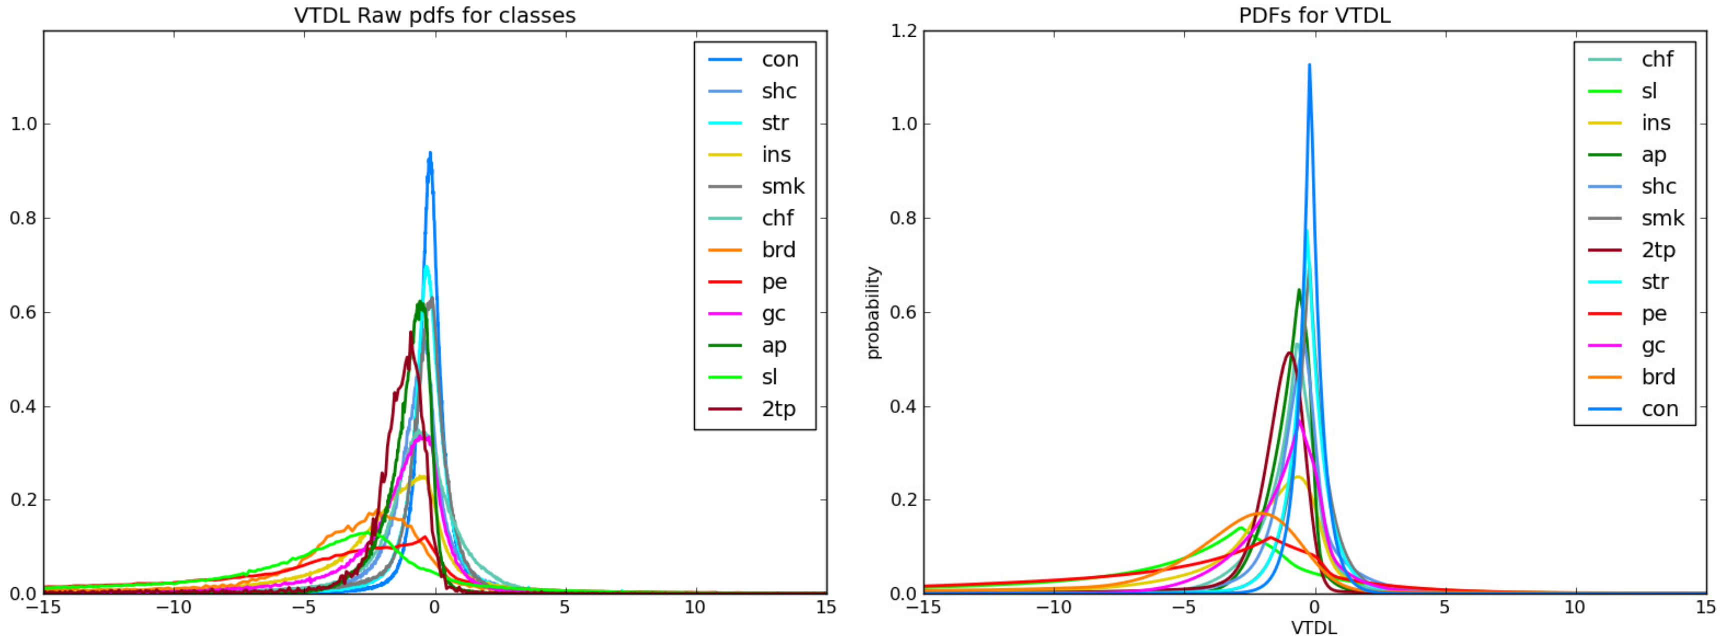

Fig. 32 Raw and fitted pdfs for VTDL, the vertical gradient of reflectivity.

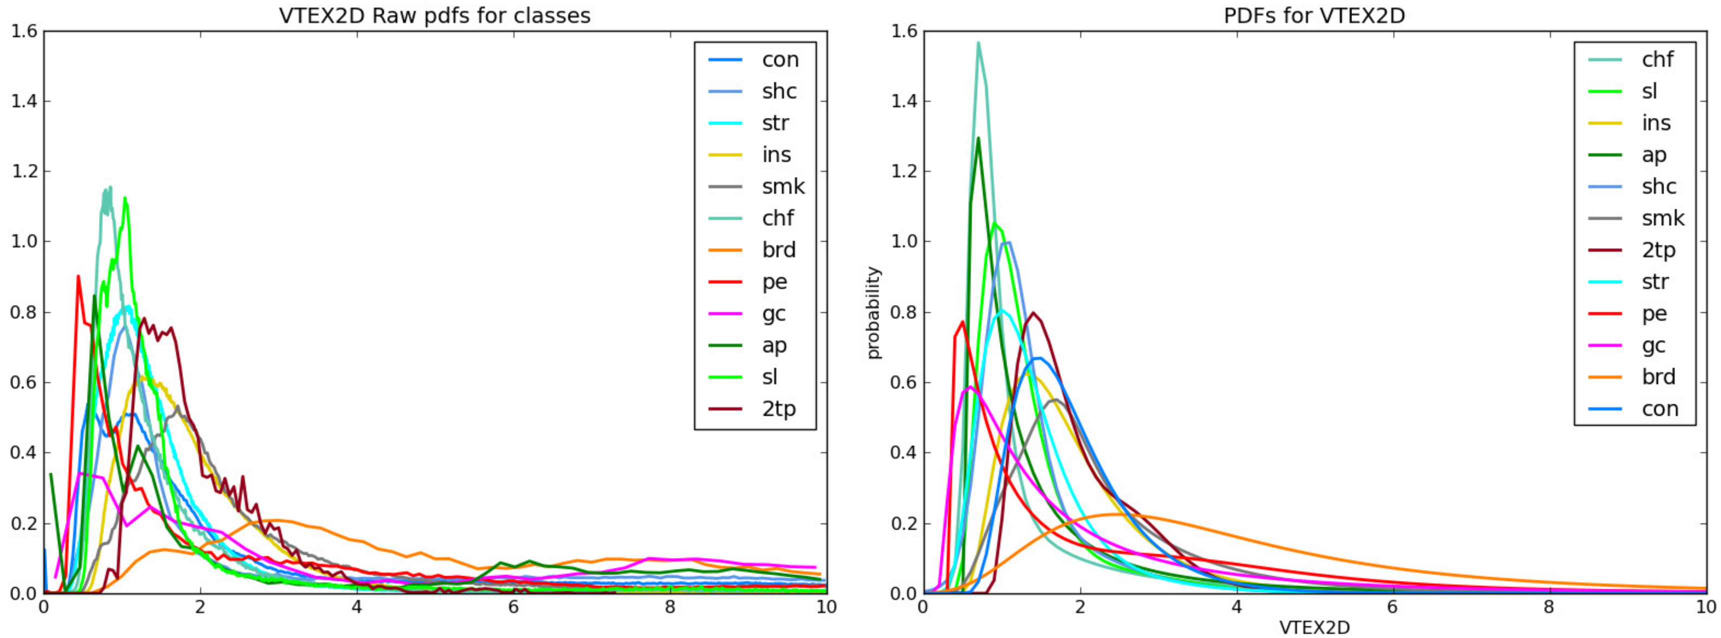

Fig. 33 Raw and fitted pdfs for VTEX2D, the texture of radial velocity.

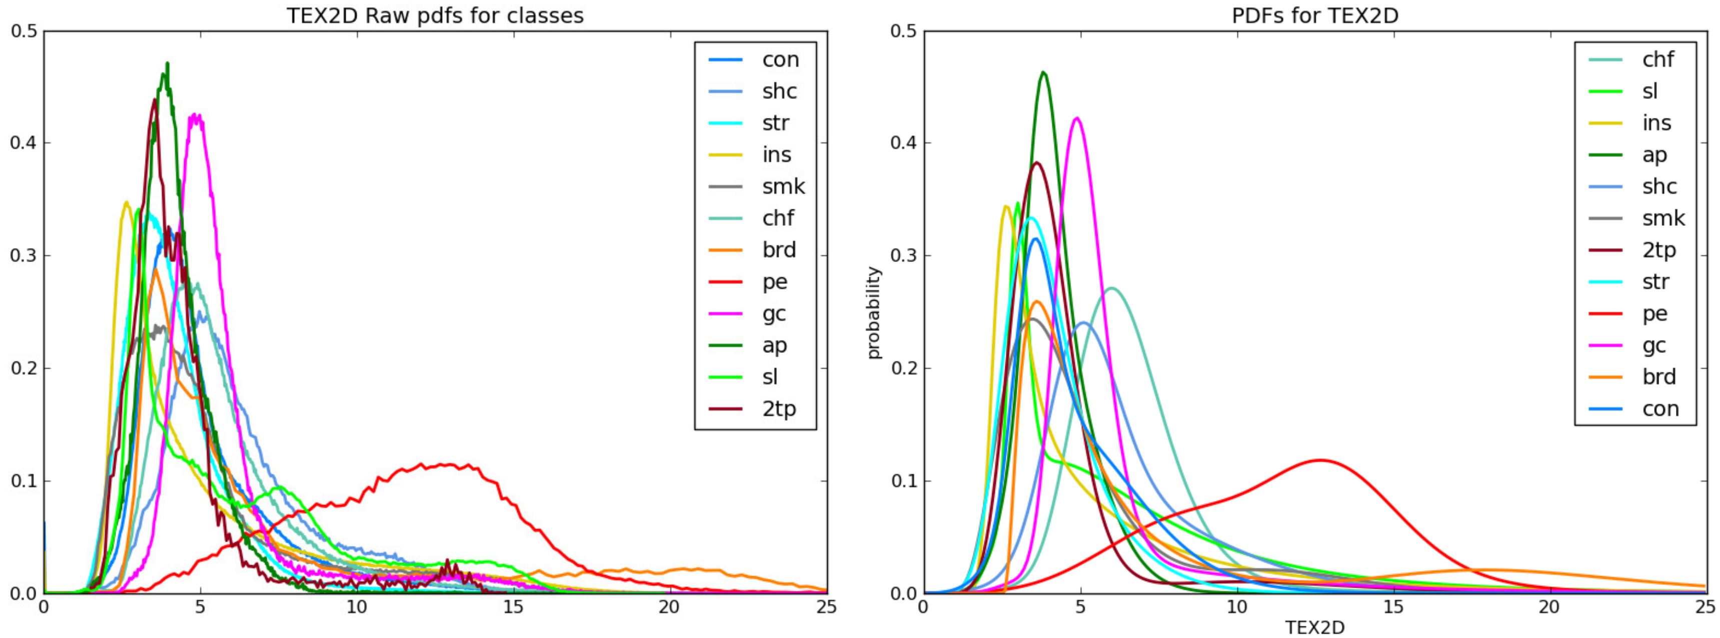

Fig. 34 Raw and fitted pdfs for texture of reflectivity (TEX2D).

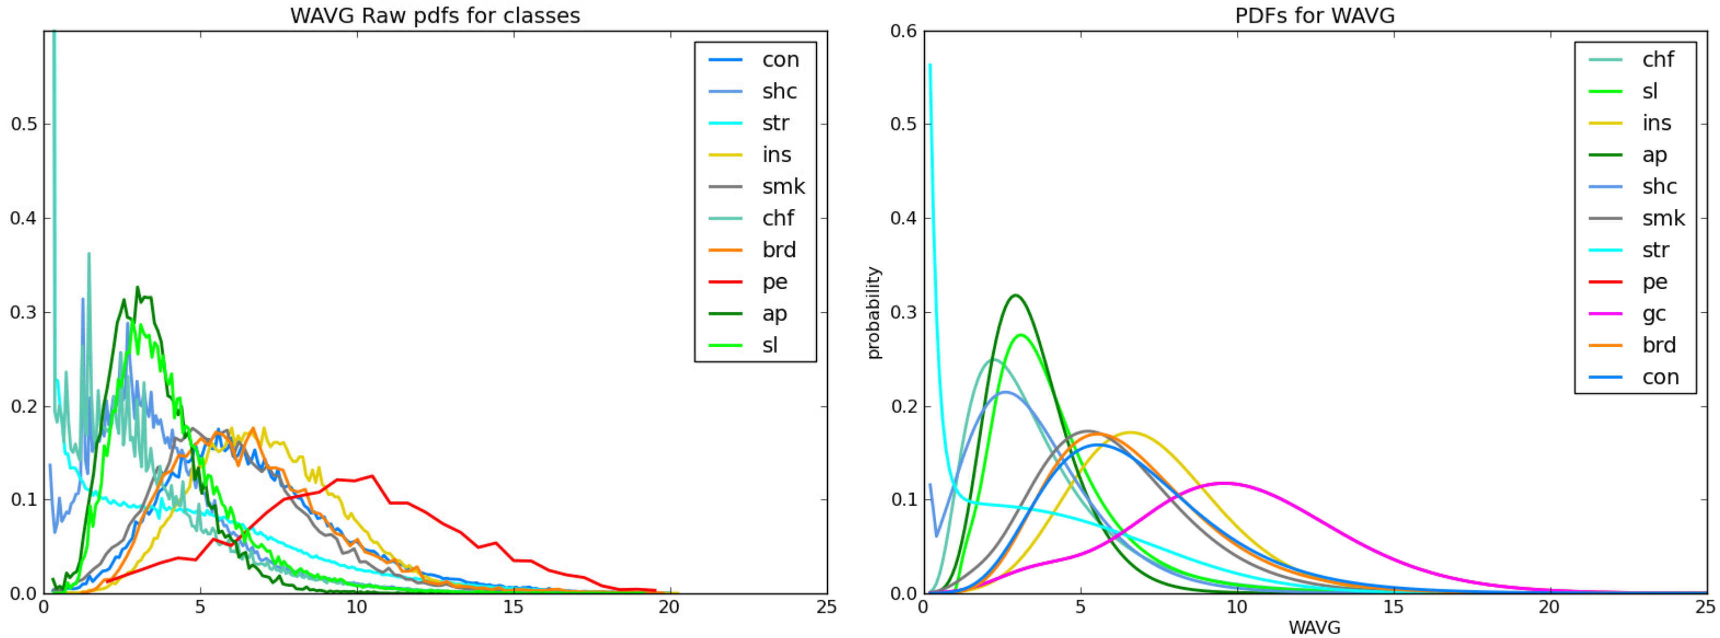

Fig. 35 Raw and fitted pdf for spectrum width, averaged with a Gaussian kernel filter (WAVG). Note that pe and gc have the same pdf because there is no raw pdf for the gc class.

Table 20 Typical pdf choice for each class and feature field.

class	*DBZH*	*ETH*	*ETH2*	*VTDL*	*SPIN2D*	*ZTEX2D*	*VTEX2D*	*WAVG*
con	skew normal	skew normal	gamma	laplace skewnorm	skew normal	log binormal	inverse normal	log binormal
shc	skew normal	gamma	log binormal	laplace skewnorm	skew normal	log binormal	log binormal	log binormal
str	skew normal	log binormal	log binormal	laplace skewnorm	skew normal	log binormal	log binormal	laplace normal
ins	inverse normal	log binormal	inverse gamma	laplace skewnorm	skew normal	log binormal	inverse normal	log binormal
smk	trapezoid	trapezoid	trapezoid	laplace skewnorm	skew normal	log binormal	log binormal	log binormal
chf	inverse normal	laplace normal	gamma	laplace skewnorm	laplace normal	log binormal	log binormal	log binormal
brd	skew normal	inverse gamma	gamma	laplace skewnorm	gamma	log binormal	inverse gamma	log binormal
pe	skew normal	inverse gamma	laplace normal	laplace skewnorm	skew normal	log binormal	log binormal	log binormal
gc	inverse normal	log binormal	log binormal	laplace skewnorm	skew normal	log binormal	inverse normal	log binormal
ap	inverse normal	log binormal	log binormal	laplace skewnorm	gamma	log binormal	inverse normal	log binormal
sl	trapezoid	laplace normal	inverse gamma	laplace skewnorm	trapezoid	log binormal	log binormal	log binormal
2tp	skew normal	log binormal	gamma	laplace skewnorm	log binormal	log binormal	log binormal	none

The Centre for Australian Weather and Climate Research is a partnership between CSIRO and the Bureau of Meteorology.