Detection of Weather Radar Clutter

Bøvith, Thomas; Nielsen, Allan Aasbjerg; Hansen, Lars Kai

Publication date:
2008

Document Version
Publisher's PDF, also known as Version of record

Detection of Weather Radar Clutter

Thomas Bøvith

Technical University of Denmark
DTU Informatics, Informatics and
Mathematical Modelling
and
DTU Space, National Space Institute
Department of Geodesy

Danish Meteorological Institute
Research Department
Remote Sensing and Geomagnetism

Kongens Lyngby 2008
IMM-PHD-2008-201

Technical University of Denmark
Informatics and Mathematical Modelling
Building 321, DK-2800 Kongens Lyngby, Denmark
Phone +45 45253351, Fax +45 45882673
reception@imm.dtu.dk
www.imm.dtu.dk

IMM-PHD: ISSN 0909-3192

Abstract

Weather radars provide valuable information on precipitation in the atmosphere but due to the way radars work, not only precipitation is observed by the weather radar. Weather radar clutter, echoes from non-precipitating targets, occur frequently in the data, resulting in lowered data quality. Especially in the application of weather radar data in quantitative precipitation estimation and forecasting a high data quality is important. Clutter detection is one of the key components in achieving this goal.

This thesis presents three methods for detection of clutter. The methods use supervised classification and use a range of different techniques and input data.

The first method uses external information from multispectral satellite images to detect clutter. The information in the visual, near-infrared, and infrared parts of the spectrum can be used to distinguish between cloud and cloud-free areas and precipitating and non-precipitating clouds.

Another method uses the difference in the motion field of clutter and precipitation measured between two radar images. Furthermore, the direction of the wind field extracted from a weather model is used.

The third method uses information about the refractive index of the atmosphere as extracted from a numerical weather prediction model to predict the propagation path of the radar's electromagnetic energy. This facilitates the prediction of areas of clutter caused by anomalous propagation of the radar's rays.

The methods are evaluated using a large independent test set, and to illustrate the performance on individual radar images three typical case examples are also evaluated. The results of the evaluation of the methods show that each method has good skill in detection of clutter with an average classification accuracy of 95 %. The methods thus have the potential for increasing the quality of weather radar data in their operational use.

Resumé

Vejrradarer giver værdifuld information om nedbør i atmosfæren, men på grund af virkemåden af radarer, observerer en vejrradar ikke blot nedbør. Radarekkoer fra andre objekter end nedbør, kaldet *clutter*, optræder jævnligt i data, hvilket medfører en forringet datakvalitet. Specielt i anvendelsen af vejrradar-data til kvantitativ bestemmelse og forudsigelse af nedbør, er en høj datakvalitet vigtig. Detektering af clutter er en af hovedkomponenterne i opnåelsen af denne.

Denne afhandling præsenterer tre metoder til at detektere clutter. Metoderne benytter *supervised* klassifikation og benytter række forskellige teknikker og datakilder.

Den første metode benytter ekstern information fra multispektrale satellitbilleder til at detektere clutter. Informationen i de synlige, nær-infrarøde og infrarøde spektre muliggør afskillelse af skydækkede og skyfri områder og skyområder med og uden nedbør.

En anden metode benytter forskellen i bevægelsesmønsteret for nedbør og clutter målt i mellem to radarbilleder. Derudover benyttes information om vindretningen fra en vejrmodel.

Den tredje metode benytter information om atmosfærens refraktive indeks, bestemt af en numerisk vejrmodel, til at forudsige udbredelsen af radarens elektromagnetiske energi. Dette muliggør forudsigelsen af områder med clutter forårsaget af anormal udbredelse af radarstrålerne.

Metoderne evalueres ved brug af et større, uafhængigt test-datasæt, og for at illustrere metodernes ydeevne i anvendelsen på enkelte radarbilleder, bliver tre typiske eksempelbilleder også evalueret. Resultaterne viser, at hver af de tre metoder har gode muligheder for at detektere clutter med en gennemsnitlig klassifikationsnøjagtighed på 95 %. Metoderne har således potentiale for at forbedre kvaliteten af vejrradardata i deres operationelle anvendelse.

Preface

This thesis was prepared at DTU Informatics, Department of Informatics and Mathematical Modelling (IMM), Technical University of Denmark (DTU) from October 2004 to December 2006 and at DTU Space, National Space Institute from January 2007 (when the Danish National Space Center merged with DTU) to June 2008, and at the Danish Meteorological Institute (DMI) in partial fulfillment of the requirements for acquiring the Ph.D. degree in engineering.

The project work was carried out at DTU in Lyngby and at DMI in Copenhagen and was supervised by Associate Professor, Ph.D., Allan Aasbjerg Nielsen (DTU) and co-supervised by Professor, Ph.D., Lars Kai Hansen (DTU), Chief Adviser, Senior Scientist, Ph.D., Rashpal S. Gill (DMI), and Senior Scientist, M.Sc., Søren Overgaard (DMI).

Funding for the project came from the Technical University of Denmark, the Danish Meteorological Institute, and the Danish Agency for Science, Technology and Innovation under the Danish Ministry of Science, Technology and Innovation through the research school CISP (Copenhagen Image and Signal Processing Graduate School).

Copenhagen, June 2008

Thomas Bøvith

Publication notes

Some of the work carried out during the Ph.D. project period was published elsewhere or is under publication.

Conference papers

Preliminary versions and first results of the data fusion method (Chapter 3) were presented at two conferences and published in the conference proceedings of these:

[19] T. Bøvith, A. A. Nielsen, L. K. Hansen, S. Overgaard, and R. S. Gill. Detecting weather radar clutter by information fusion with satellite images and numerical weather prediction model output. In *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium (IGARSS) 2006*, volume I, pages 511–514, Denver, Colorado, USA, August 2006. Conference paper.

[18] T. Bøvith, R. S. Gill, S. Overgaard, A. A. Nielsen, and L. K. Hansen. Detecting weather radar clutter using satellite-based nowcasting products. In *Proceedings of the Fourth European Conference on Radar Meteorology (ERAD 2006)*, pages 153–156, Barcelona, Spain, September 2006. Conference paper.

The presentation of [18] received the award for *Best Student Oral Presentation*.

Journal papers

A journal paper on the wave propagation method has been accepted for publication. The method is also presented in Section 5.4 of Chapter 5 of this thesis.

[14] H.-H. Benzon and T. Bøvith. Simulation and prediction of weather radar clutter using a wave propagator on high-resolution NWP data. *IEEE Transactions on Antennas and Propagation*. Accepted for publication.

This work was a joint effort of the authors, but Hans-Henrik Benzon specifically developed, modified, and applied the wave propagator, and Thomas Bøvith specifically extracted and preprocessed the input data (weather radar data, NWP data, and refractivity data) and carried out classification of radar data.

During the Ph.D. project period a contribution was made to a paper [1] on image fusion of high-resolution panchromatic images and low-resolution multispectral images. The contribution of Thomas Bøvith was on data preprocessing and quantitative quality assessment of the results of the image fusion methods:

[1] H. Aanæs, J. R. Sveinsson, A. A. Nielsen, T. Bøvith, and J. A. Benediktsson. Model-based satellite image fusion. *IEEE Transactions on Geoscience and Remote Sensing*, 46(5):1336–1346, 2008.

Acknowledgements

Many people have shown interest in my work and have provided help and support during the course of my Ph.D. study—I am grateful for this and in the following I would like to acknowledge them.

Firstly, I would like to thank the supervisors of my work, Allan Aasbjerg Nielsen (DTU), Lars Kai Hansen (DTU), Rashpal S. Gill (DMI), and Søren Overgaard (DMI), for their valuable input and motivational comments.

Henrik Steen Andersen (DMI) is thanked, for initiating the Ph.D. project in the first place—and as head of division for providing an excellent research environment in the former Ice Charting and Remote Sensing Division of the Observation Department, DMI, in the very first time of the project. My former colleges there are thanked as well. Søren Andersen, is thanked for input to the project in the early stages of the work as well as for help running Linux when I first started working at DMI in 2003. Morten Lind is thanked for help with nowcasting SAF products.

Also thanks to my current colleges in the Remote Sensing and Geomagnetism division of the Research Department, DMI. Georg Bergeton Larsen, head of the division, for providing a great and pleasant research environment for finalizing my thesis. Hans-Henrik Benzon, DMI, for good collaboration on the wave propagation method. Thanks also to the "Radar Group" at DMI.

Bent Hansen Sass and Bjarne Amstrup, Meteorological Model Systems division, DMI, are acknowledged for supplying the DMI-HIRLAM data.

Thanks to the staff and students of the Image Analysis, Computer Graphics, and Geoinformatics group at the Department of Informatics and Mathematical Modelling (IMM) at DTU, for creating a great academic and social frame for my Ph.D. studies. Rasmus Larsen is thanked for access to the optical flow software. Thanks to Henrik Aanæs for inviting me to contribute to the image fusion paper. Special thanks to Tove Jacobsen and Eina Boeck for kind and efficient help with all practical matters.

Also thanks to my new colleges at the Geodesy department of DTU Space, National Space Institute, DTU, lead by head of department, Per Knudsen.

Thanks to Adolfo Magaldi, Department of Astronomy and Meteorology, University of Barcelona, Spain, for discussions on the methods for clutter detection during my visit to Barcelona, 2006.

I would also like to thank Professor Peng Gong and the students and staff of the remote sensing lab at the University of California, Berkeley, USA, for providing an inspiring research environment during my six month visit from November 2006 to April 2007.

Acknowledgements are due to the makers of the Free/Open Source Software used for this project, in particular: the GNU/Linux project (the GNU operating system and the Linux kernel), Generic Mapping Tools (GMT) used for making the maps, and the LaTeX project for excellent typesetting software.

Finally, I would like express my gratefulness for the support, caring and patience from my family and my friends, both during the recent times of writing this thesis and in general—special thanks to my parents, Ingrid and Karsten, to my great-aunt, Bodil, and to my girlfriend, Monica.

"It's just a box of rain
wind and water
believe it if you need it
if you don't just pass it on ..."

—Robert Hunter

Quote from the song *Box of Rain*,
as performed by The Grateful Dead.
http://arts.ucsc.edu/gdead/agdl/box.html

Contents

Abstract i

Resumé iii

Preface v

Publication notes vii

Acknowledgements ix

1 Introduction 1

2 Weather radar, clutter and detection of clutter 5

2.1 Weather radar 5

2.2 Weather radar data 6

2.3 The DMI weather radars 8

2.4 Weather radar clutter 8

2.5 Clutter detection methods 13
2.6 Database of radar data 17

3 Data fusion method **23**

3.1 Introduction 23
3.2 Remote sensing work flow 24
3.3 Data: Meteosat Second Generation and Precipitating Clouds . . 25
3.4 Method: Data fusion method 34
3.5 Preprocessing 34
3.6 Feature extraction 37
3.7 Training data 38
3.8 Feature selection 40
3.9 Classification 50
3.10 Model training 54
3.11 Classifier evaluation 58
3.12 Reject option 63
3.13 Results 65
3.14 Summary and conclusion 69

4 Spatio-temporal method **71**

4.1 Introduction 71
4.2 Data 72
4.3 Feature extraction 73
4.4 Feature selection 91

4.5 Classification 92
4.6 Classifier evaluation 92
4.7 Results 95
4.8 Summary and conclusion 95

5 Propagation method **99**

5.1 Introduction 99
5.2 Radar propagation 100
5.3 Data 102
5.4 Method I: Wave propagation method 105
5.5 Method II: Azimuth-dependent propagation index 119
5.6 General summary and conclusions 129

6 General summary and conclusion **131**

6.1 Summary 131
6.2 Conclusion 135
6.3 Future work 136

A Radar image database **139**

B Feature histograms **143**

C Feature conditional probabilities **165**

D Feature selection results **187**

Bibliography **224**

Chapter 1

Introduction

Precipitation has great influence on the conditions of human activities and on the Earth's natural environment in general. For this reason, observation of precipitation is an important field in meteorological science and operational meteorology. The *weather radar* is an unique sensor for obtaining observations of precipitation in the atmosphere.

The information available from weather radars play an important role in forecasting, planning, and decision-making in many areas of society, e.g., in meteorology, hydrology, agriculture, and transportation (as illustrated in Figure 1.1). In hydrology, weather radars provide real-time input for prediction of flooding and the management of wastewater. In agriculture, decisions on harvesting and fertilization depend on the forecast of rain, and in transportation, the prediction of the risk of icy roads depends on observations of precipitation. In general, for all people, precipitation has a large impact on the conditions for many, mainly outdoors, activities.

Precipitation is highly variable in both time and space (typical scales are in the magnitudes of minutes and meters, respectively) and the weather radar is capable of providing information on precipitation on these time and space scales. In meteorology, weather radar data are an important source of information for very-short-term weather forecasting, also called *nowcasting*, which aims at precise forecasts for a few hours ahead in time. Especially in events of extreme and

severe weather precise maps of precipitation patterns and intensities are indispensable for issuing of warnings. Recent climate research predict an increasing frequency of heavy precipitation as result of the trend of global warming [40], which, of course, will put higher demand on remote sensing of precipitation as provided by weather radar.

Outside forecasting, the precise *observation* of precipitation itself is important, and the use of weather radar in combination with rain gauges (point measurements of precipitation at ground level) are used for keeping a record of the fall of rain, hail, and snow.

Figure 1.1: Weather radars provide information for weather forecasting (top left), planning of ourdoors activities (top right), traffic (middle and lower right) and agriculture (lower left). Photos: DR TV-avisen, other photos under Creative Commons licenses: http://flickr.com/photos/sekihan/2255090253/, http://flickr.com/photos/driek/1172780201/, http://flickr.com/photos/7232802@N06/623398013/, http://flickr.com/photos/markybon/212807141/.

Weather radars make use of the radar principle to measure the location and properties of precipitation in the atmosphere. Due to the way radars work, however, it is not guaranteed that only the phenomenon of interest (here: precipitation) is observed. On the contrary, it is very normal that other targets also show up in radar data (Figure 1.2). The unwanted radar echoes are called *clutter* or *false echoes* and their presence lower the data quality for all applications of the data. The poor data quality results in uncertainties which propagate through the data analyses to the final decisions made upon these. This is especially a problem in automated and quantitative uses of weather radar data, for example when used in hydrological models, nowcasting models, and numerical weather prediction models, but also for the qualitative use of weather radar data, in weather forecasting for example, clutter leads to unwanted uncertainty.

This motivates the development of methods for detection of weather radar clutter and hereby facilitate the removal of clutter. This is the main purpose of the work presented in this thesis.

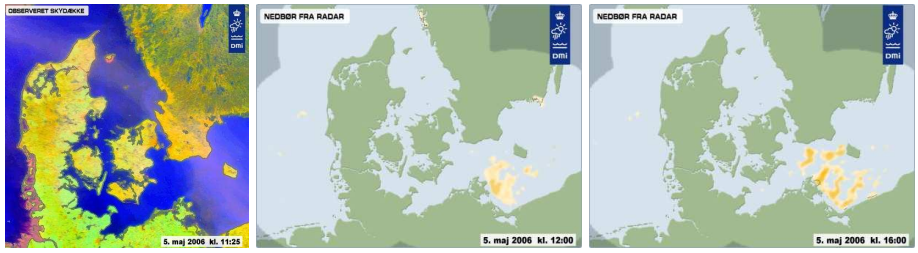

Figure 1.2: Clutter is a common unwanted feature of weather radar data. The satellite image on the left (from a NOAA polar-orbiting weather satellite) shows cloud-free conditions, while the two radar images (middle and right), as presented on the DMI website, show radar clutter (patches in yellow hues) in the Baltic Sea.

Thesis organization

The thesis is structured in the following way: In Chapter 2, the weather radar and data from weather radars are described. Clutter and the various types of clutter are described and a review of current available methods for detecting clutter is made.

Then follows the main chapters of the thesis which present three methods for clutter detection: Chapter 3 presents a data fusion method which uses external information on precipitation available in multispectral satellite data. In

Chapter 4, a spatio-temporal method is presented. The method makes use of the difference in the motion field of clutter and precipitation. The next chapter, Chapter 5, presents two methods both making use of data from numerical weather models to predict the meteorological conditions which lead to clutter. This is used to detect clutter in the radar data. In Chapter 6 the results are compared and summarized, conclusions are drawn and looks at future work are made.

In the Appendix A to D are found figures and tables for completeness.

Chapter 2

Weather radar, clutter and detection of clutter

2.1 Weather radar

Weather radars are radar systems designed for observation of precipitation in the atmosphere [29], [25], [56]. Making use of the radar principle, they can determine the location and properties of hydrometeors (precipitation particles) at a given distance from the radar and height above the ground. A radar is an active sensor in that it transmits electromagnetic energy into its surroundings, the energy is reflected by objects in its path and a small part of the backscattered energy is finally received and recorded by the radar. Because the weather radar is a calibrated radar, the received power can be translated into a measure of reflectivity of the precipitation using the radar equation. The reflectivity (measured in dBZ) varies for different precipitation intensities (low reflectivity for light precipitation and high reflectivity for heavy precipitation) and for different precipitation types (rain, hail, and snow). This is due to the difference in drop sizes and drop shapes of the hydrometeors.

A weather radar is typically a ground-based radar, and sensing of the entire atmospheric volume around the radar is carried out by rotating the antenna around its vertical axis and changing the pointing angle of the antenna (typically once for each revolution of the antenna). This is illustrated in Figure 2.1.

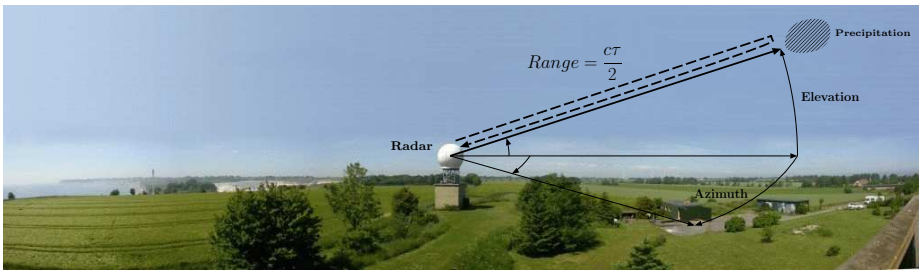

Figure 2.1: The weather radar is normally mounted on a tower or a building to avoid blockage of the radar rays by near-by buildings, trees and other obstacles. The range to a target is computed from the round-trip time, τ, and the speed of light, c.

The angle in the horizontal plane from geographical north in positive clockwise direction is called the *azimuth angle* and the tilt angle measured from horizontal and positive upwards is called the *elevation angle*. The *range* to a target is computed from the round-trip time of a pulse to the target and the speed of the electromagnetic energy (i.e., the speed of light).

A complete radar system, is a complex system comprised of many components [69]: From the hardware in the transmitter, antenna, and receiver, to signal processors and data analysis and visualization software, all these play together in the generation of weather radar data.

2.2 Weather radar data

The signal processor of the weather radar typically receives data from a number of pulses which are integrated and sampled to a polar coordinate system, one scan for each revolution and elevation angle. For a typical weather radar a 1 ° azimuth resolution is used and in the range direction, a resolution of 500 m or more is used (See Figure 2.2, left).

The resolution of weather radar images is a different concept than aerial or satellite images since the data are acquired in a polar coordinate system. The 1 ° beam width means that the resolution in the azimuth direction varies from meters close to the radar to kilometers at the maximum range of e.g., 240 km, while the range resolution is constant throughout the data. When converted to cartesian coordinates the data thus have a varying resolution even if the image has a specified pixel size.

Displaying a scan in cartesian coordinates is called a PPI image (Plane Position Indicator) as seen in Figure 2.2, right. Since the radar rays normally increase with distance to the radar, the observed precipitation of course is close to the ground near the radar in the center of the image and higher up in the atmosphere at far range from the radar. By using the data from several PPIs, a CAPPI (Constant-Altitude PPI) can be constructed which show the precipitation at a given height. Typically, a stack of CAPPIs are produced at intervals of 1 km. Other types of 2D image representations exist, for example what is called a *base product* which contains the lowest precipitation value from each PPI [33]. Finally, the 2D image products, CAPPI or base products, are often mosaiced into one image as seen in Figure 2.3.

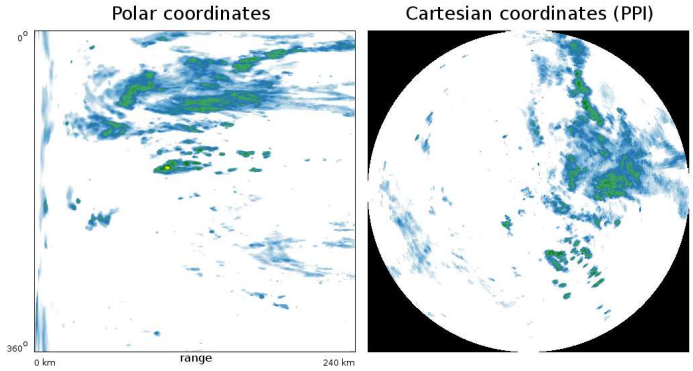

Figure 2.2: Polar coordinate scan, 1 °×500 m grid (left) and PPI cartesian product, 500 m×500 m pixel size (right), for the radar at Stevns on 2006–05–19 12:00 UTC.

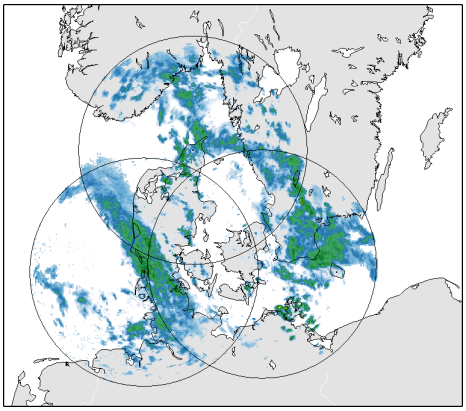

Figure 2.3: Mosaic of the image from Stevns shown in Figure 2.2 and the corresponding images from two other radars.

Most modern weather radars are *Doppler radars* which means that they are able to measure the phases of the transmitted and received signals. The shift in phase is used to derive the radial velocity of the precipitation particles. In meteorology this is used to map the wind speed but it is also used in clutter detection because echoes from non-moving targets are unlikely to be precipitation (this is described in greater detail in Section 2.4).

A new technology in the field of weather radar is dual-polarization radars. These are radars with the capability to transmit and receive electromagnetic energy in two polarizations. Normally, weather radars are operated with horizontal polarization only (because falling rain drops are flattened as they fall and the backscatter is greater in the horizontal than the vertical). Dual-polarimetric radars [21] provide a range of additional information besides radar reflectivity and the Doppler velocity. The differential reflectivity (the ratio between the horizontal (H) and vertical (V) power returns), the correlation coefficient (correlation between H and V power returns), and the differential phase (phase difference between H and V returns) enable improved hydrometeor classification (detection whether it is rain, hail, or snow) as well as better clutter/precipitation discrimination.

2.3 The DMI weather radars

For this project, data from the weather radars of the Danish Meteorological Institute (DMI) were used. The radar network is currently comprised of four C-band radars located in Rømø, Sindal, Stevns, and Bornholm (See Figure 2.4 for a map of the weather radar network). The first three are Doppler radars and the latter is a dual-polarization radar installed in 2008 as an upgrade to a non-Doppler radar. Because the data collection for this project was finished prior to 2008, the data from Bornholm were not used. The specifications of the DMI radars are summarized in Table 2.2 and are also presented in [34].

2.4 Weather radar clutter

Ever since the invention of radar in the beginning of the twentieth century, over its rapid developments during World War II to its meteorological applications after the war [6], *radar clutter* has been a challenge to the users and the uses of radar data. As mentioned earlier, there is no guarantee that the radar only observes the target of interest due to the design of the radar system, and when non-targets are observed *clutter* echoes appear in the data.

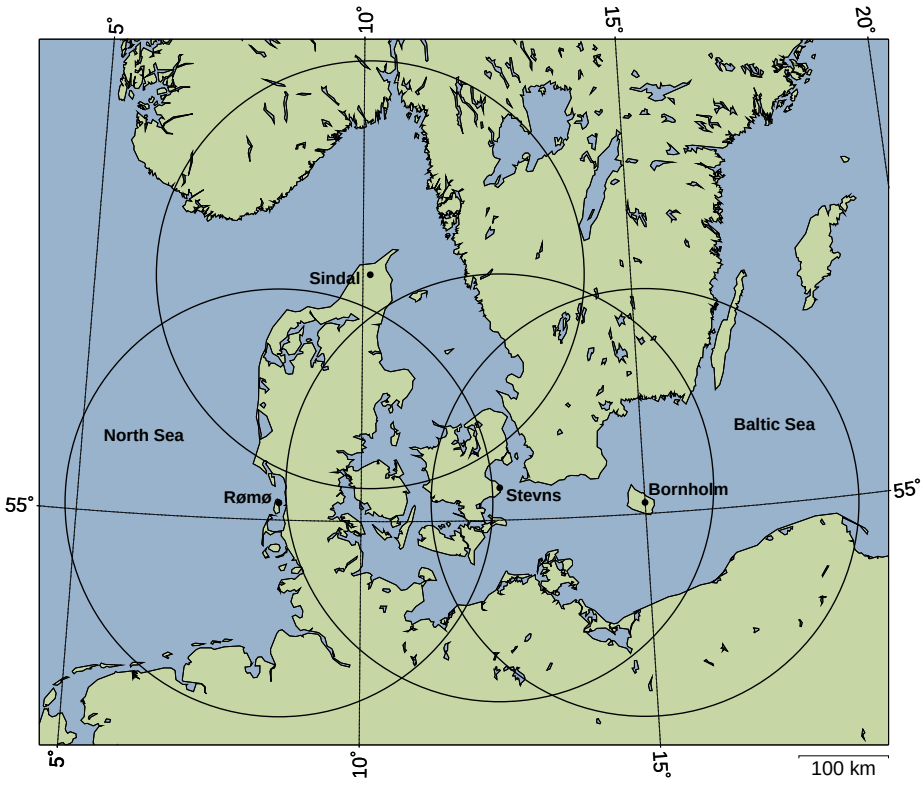

Figure 2.4: The DMI weather radar network. The range rings show the maximum range (240 km) of each radar. Radar locations and heights are listed in Table 2.1

Table 2.1: Geographical locations of the DMI weather radars.

Radar name and abbrev.	Geographical coordinates	Height above sea level
Rømø (EKXR)	55.173 °N, 8.552 °E	15.0 m
Sindal (EKSN)	57.450 °N, 10.136 °E	109.0 m
Stevns (EKXS)	55.326 °N, 12.449 °E	53.0 m
Bornholm (EKRN)	55.113 °N, 14.890 °E	171.0 m

Table 2.2: Parameters of the DMI weather radars.

Parameter	**Value**
Radar frequency	5.625 GHz (C-band)
Wavelength	5.3 cm
Radar peak power	250 kW
Beam width	1 °
Pulse length	2 µs
Radar maximum range	240 km
Pulse repetition frequency	250 Hz
Nyquist velocity	3.3 m/s
Rotation speed	20 °/s
Range gate size	500 m
Number of elevation angles	9
Elevation angles	0.5, 0.7, 1.0, 1.5, 2.4, 4.5, 8.5, 13.0, 15.0 °
Temporal resolution	10 min

Clutter is not the same for all applications of radar, of course. For example, in an airport, radars are used in multiple applications. One radar is employed for tracking airplanes and hence airplanes are the target of the radar application. To this radar application everything else but airplanes is clutter. Another radar might be used for observing precipitation, e.g., to avoid take-off during hail storms, and to this radar application everything but echoes from precipitation is clutter. In other words, and to rephrase a well-known proverb: *"One man's clutter is another man's target."*

Types of weather radar clutter

The types and sources of weather radar clutter are manifold and can be grouped by *where* they occur and by *what* is causing the clutter.

Using the first approach, there is *ground clutter* (with the subsets sea clutter and land clutter) and *airborne clutter*. This is illustrated by the sketch in Figure 2.5. Land clutter can be man-made objects (e.g., houses, towers, and bridges) or the natural environment (e.g., mountains, hills, or fields). Sea clutter is caused by backscatter from the surface of oceans or lakes, and ships also show up as sea clutter. Airborne clutter is caused by reflections from airplanes or other man-made objects in the air, or biological targets like birds and insects.

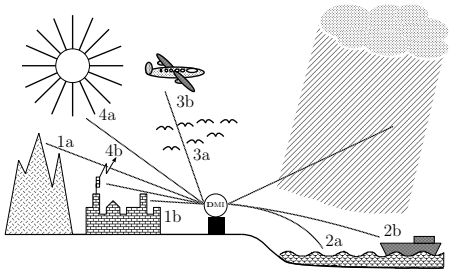

1 Land clutter from (a) mountains/earth surface or (b) buildings/windmills, etc.

2 Sea clutter from (a) sea surface or (b) ships

3 Airborne clutter from (a) biological targets (birds/insects) or (b) airplanes

4 Interference clutter from (a) the sun or (b) transmitting antennas

Figure 2.5: Types of clutter illustrated.

Using the second approach, to group clutter due to its causes, there is clutter caused by *beam blocking*. This occurs when objects or the topography are located in the normal path of the radars signals causing ground clutter. This is typically tall buildings or towers situated close to the radar site. In mountainous regions beam blocking from mountains is very common, but due to the low topography of Denmark this is not a major contributor to clutter there.

Another cause of clutter is related to the propagation path of the radars energy in the atmosphere. Normally, the radar rays increase in altitude with distance from the radar. This is due to the use of positive elevation angles of the antenna combined with the refractive index of the atmosphere which normally cause the rays to follow a curve which can be approximated with a sphere with a radius of 4/3 of that of the Earth.

Anomalous propagation, also called *non-standard propagation*, causes the radar rays to deviate from the normal path. When the rays are bent toward the Earth's surface they can produce clutter in the case of intersection with the land or sea surface. This type of ground clutter is also called *anomalous propagation clutter*, *anaprop*, or simply *AP*. Anomalous propagation is in spite of its name, quite

normal and occurs frequently in data from most weather radars. Anomalous propagation is described in greater detail in Chapter 5.

A third contributor to clutter is *interference clutter* caused by either solar radiation or radio signals transmitted by other antennas entering the radars receiver.

AP clutter and ground clutter caused by beam blocking is very common in weather radar data. In mountainous regions, land clutter from mountains is a challenge to mitigate and correct for, and in coastal regions, radars situated near the coast line are often prone to sea clutter.

The other types of clutter are mostly of less extent. The amount of airborne clutter from airplanes, helicopters, etc., is occasional, however, clutter from insects and birds are in certain regions quite severe. In Denmark, however, clutter from birds and insects is not a major issue. Interference clutter is an increasing problem in many places across the world due to the extended use of wireless communication networks and the increasing number of electromagnetic devices in the environment.

Examples of weather radar clutter

In Figure 2.6 three examples of clutter in weather radar images are shown. The image on the left shows AP land clutter on the mountains of Norway and Sweden caused by anomalous propagation of the radar's energy. Furthermore, land clutter close to the radar is seen. In the middle image, an example of interference clutter caused by solar radiation is seen and on the right image strong sea clutter caused by anomalous propagation in the Baltic Sea is seen.

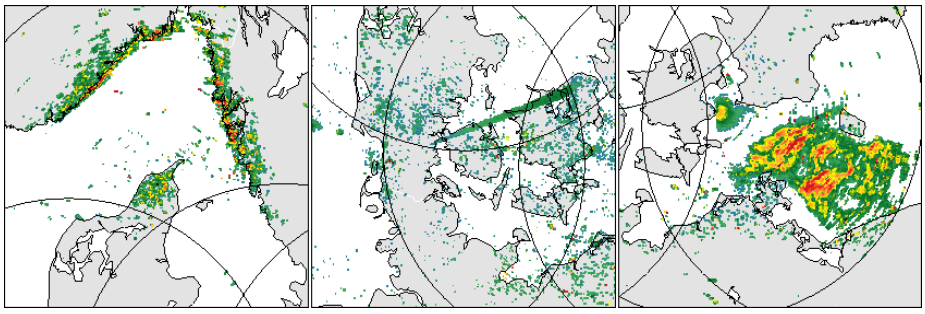

Figure 2.6: Three types of clutter: Land clutter (left), Interference clutter from the sun (the streak from the center pointing east, northeast) (middle), and sea clutter (right).

In Figure 2.7 an example of clutter caused by a wind farm off the coast of Jutland is seen. The windmills resemble the small convective cells around it, but inspection of a radar animation reveals that the echoes from the wind farm do not move with the wind.

(a) Horns Rev wind farm (Photo: Promotional photo http://www.vattenfall.dk/).

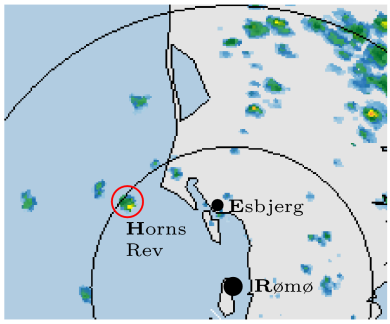

(b) Radar image from Rømø.

Figure 2.7: Example of clutter caused by windmills in a wind farm.

2.5 Clutter detection methods

The body of research into detection of weather radar clutter is quite extensive due to the relatively long time it has been dealt with and the complexity of the problem. Hence many different methods for detection of clutter have been proposed.

Looking at clutter detection research as a whole, the treatment of clutter caused by anomalous propagation seems to have received the most attention. This supports the statement in the previous section that AP clutter is probably the largest contributor to clutter and that it is also the most challenging to detect successfully. Most methods aim at solving the problem of clutter detection for one or a few clutter types, while few approaches deal with detecting multiple clutter types in an integrated data analysis framework [62].

Recent reviews of clutter detection techniques as found in [70] and [16], provide a good basis for grouping of methods. Grouping can be done according to the *level* at which clutter detection is performed and according to which *data* and *features* are used. Grouping based on the applied statistical *methods* can also be done.

Level grouping

Choice of radar site and hardware

Even before the radar is placed in the landscape, some precautions can be taken to lower the risk of clutter. So rather than being a way to *detect* clutter, this is a way to *avoid* clutter in the radar data in the first place. Many of the aspects of this are summarized in [70]: The choice of radar site is crucial to the amount of expected clutter. It is important to avoid close-by buildings, trees, and hills. This is achieved by finding a highly elevated site and by placing the radar on a tall tower or building. To this it can be added that a clear horizon is not always enough to avoid enhanced risk of clutter: A radar with no blocking, but placed near the ocean might result in increased clutter problems as well, due to increased clutter from the antenna sidelobes. Finally, a suitable choice of radar hardware parameters can help mitigate clutter problems. Of course, the choices of radar site and hardware are only made once per radar installation.

Low-level signal processing methods

These types of clutter detection methods are carried out in the radar's signal processor, i.e., before the actual radar (image) product is generated. A long list of methods are mentioned in [70], again. For Doppler radars, very efficient are filters using the Doppler velocity and spectrum width. Targets with zero or very low velocities are likely to be stationary ground targets (clutter). Precipitation, on the other hand, usually have non-zero velocities. Doppler clutter filters, however, are not efficient in detection of sea clutter echoes originating from the ocean surface which have non-zero velocities due to the motion of the ocean waves. Furthermore, clutter from other moving targets (e.g., windmills, birds, etc.) is difficult to detect using the Doppler velocity alone.

Pattern classification techniques

The modification of the radar's signal processor is an advanced task and is often only applicable using non-operational, research radars. Since many clutter detection methods are developed and applied in an operational setting with many users dependent on uninterrupted data delivery, these methods are based on using the radar data products in the step after the signal processor. This entails using the radar volumes, CAPPI products, mosaics of CAPPIs or similar products. From these data, various features are extracted which enable clutter/precipitation discrimination. This is described in the next section.

Data and features grouping

Radar data alone

These methods are based on analysis of the radar data itself. One of the simplest ways to deal with ground clutter is to generate a *clutter map*. This is done by observing ground clutter during times when no precipitation is present. The observations are integrated over some time and a mean clutter map is generated. This can subsequently be used to efficiently remove ground clutter, of course at the risk of eliminating precipitation echoes when mixed with clutter. Another shortcoming of this approach is that it is in practice only efficient in detection of land clutter. Sea clutter caused by anomalous propagation is more variable in time and space and a mean map of these can therefore not be generated.

A number of different features derived from the radar data itself are reviewed and implemented in [16]. Some methods make use of the difference in the spatial decorrelation time of clutter and precipitation signals [72], [71]. The decorrelation is faster for clutter than for precipitation. This is reported to be efficient in removing land clutter but not very efficient for sea clutter. Another group of methods successfully use texture as a feature for clutter/precipitation discrimination [49], [44], and [32] uses a 3D texture measure. In general, the 3D structure of the radar volume is a much used feature. The reason why this feature is used is the fact that AP clutter only affects the lower radar elevations (this is explained further in Chapter 5). Therefore geometric features of the vertical extent, the gradient, and the variability of the radar echoes are usefull [3], [70], [16], [24], and [48].

New radar technology

Recent advances in radar technology, especially dual-polarimetric radars [21], have provided new methods for clutter detection. Operating the radar in more than one polarization mode yields important information on the radar targets which can be used in identification of clutter as shown by [66], [65]. The correlation coefficient between horizontal and vertical power returns as well as the differential phase shift are used as good features. The use of dual polarization radars is still in its early days for operational use and single polarization radars will be the standard for years to come. So improved methods for single polarization radars are still needed.

Data fusion techniques

Another group of methods make use of data fusion in which the use of multiple and complimentary data sources improve the detection of clutter. Two ma-

jor groups of methods can be identified: 1) those using multispectral satellite observations to determine areas of precipitation/no precipitation and 2) those using meteorological parameters from weather models or radiosondes to predict areas of possible anomalous propagation clutter by modelling the propagation conditions for a given radar site.

1. Satellite observation methods

Using multispectral satellite observations in the visible, infrared and thermal infrared spectra to map areas of precipitation [54], [23], has provided new ways to detect clutter: Several methods using observations from geostationary satellites have shown promising results [60], [31], [61] and [58]. These methods use first generation geostationary satellites (Meteosat-7 and earlier, or GMS (Geosynchronous Meteorological Satellite). More recent work [19], [18], and [55] use Meteosat Second Generation with its improved resolution. Data fusion methods using multispectral satellite images are described in more detail in Chapter 3 where a new data fusion method is proposed and evaluated.

2. Propagation modelling methods

This group of methods use information from radiosondes or numerical weather prediction (NWP) models to predict when anomalous propagation occurs. From the parameters, temperature, air pressure, and humidity, the refractive index of the atmosphere can be computed which determine the electromagnetic propagation path. Some authors have used observations from radiosondes (weather balloons, which measure the vertical atmospheric profile of amongst other variables the temperature, pressure, humidity) [11] while others have used predictions of the meteorological parameters using NWP model output [10], [12], and [9]. Propagation methods are described further in Chapter 5 where two methods for clutter detection using propagation modelling are presented.

Statistical methods used

Clutter detection techniques employ a variety of techniques for deriving the class memberships of the echoes in radar data. Most methods use supervised classification techniques and the techniques range from simple thresholding on the input features [58], to application of classification tree methods, neural networks [48], [61],[35], [26] and fuzzy logic classifiers [16], [45], [24]. To compare the performance of the statistical methods in clutter detection applications as presented by the various authors is a difficult task. The configuration of the experiments differ too much in their combination of input data, the applied preprocessing, the amount of training and test data used. Furthermore, the

methods do not report their performance in a consistent away.

Thesis context

In the context of clutter detection methods, this thesis presents three methods from multiple of the above mentioned groups. In Chapter 3, a method using data fusion with satellite images is presented, in Chapter 4, a method using radar data alone and in combination with wind field predictions from a weather model is presented, and finally Chapter 5 presents methods using propagation modelling. The methods use the same weather radar dataset as input which makes it possible to compare the results of the methods with each other.

2.6 Database of radar data

A database of typical meteorological events was established for the development and testing of methods. From 15 days in 2005 and 2006 data were extracted from the three C-band Doppler weather radars of Sindal, Rømø, and Stevns. Data were extracted for four time points a day (at 01:00, 06:00, 12:00, and 18:00 UTC). The days showcase differing meteorological conditions with various combinations of clutter types and precipitation types (only rain but both convective and wide-spread precipitation). Table 2.3 describes each event and Appendix A shows thumbnail images of each of the 60 images.

The raw radar volume data with a resolution of $1^\circ \times 500\,\mathrm{m}$ were processed into 2D images using an in-house DMI algorithm, DMI Base [33], and subsequently mosaiced to a common stereographic map grid with a pixel spacing of 1000 m. The temporal resolution of the data is 10 minutes.

Three case examples

Three case examples are selected to provide the means to describe in more details the data as well as to evaluate the proposed methods on single images of the complete data set. The three examples show the cases of clutter only, clutter and precipitation, and precipitation only. *Case I* is a case of land and sea clutter practically without precipitation. *Case II* is a case of sea clutter with more precipitation and finally, *Case III* is a case of no clutter during an event of mixed precipitation.

Table 2.3: Metadata for radar data. Date and time and a short description of the type of meteorological event. See Appendix A for thumbnail images of from each day.

Date	Description
2005–07–11	Strong land clutter on Sindal radar and moderate to strong sea clutter on Stevns radar. Small amounts of precipitation on Rømø radar.
2005–07–13	Moderate land clutter on Sindal radar. Moderate sea clutter on Stevns radar. Minor precipitation on Rømø radar in the late afternoon and evening.
2005–07–15	Strong and extensive convective precipitation coming in from the south west. Minor land clutter on Sindal radar. Moderate sea clutter on Stevns radar.
2005–08–25	Widespread precipitation system coming in from the south west. Minor land and sea clutter on Sindal and Stevns.
2005–09–06	Extensive land clutter on Sindal and Stevns radars. Sea clutter on Rømø and Stevns radar. Minor precipitation on Rømø and Sindal radars.
2005–09–25	Strong sea and land clutter on Stevns radar. Precipitation fronts coming in from the south west.
2005–09–26	Strong sea and land clutter on Stevns radar. Precipitation fronts coming in from the south west.
2005–10–19	Strong and extensive land clutter on Sindal radar. Minor sea clutter on Stevns radar. Precipitation on Rømø radar coming in from the south.
2006–01–02	Minor precipitation event on Stevns radar. Minor land clutter on Sindal radar.
2006–03–24	Precipitation on Rømø and Sindal radar. Sea clutter on Stevns radar. Land clutter on Sindal radar.
2006–05–05	Strong sea clutter on Stevns radar. Land clutter on Sindal radar. Minor precipitation on Rømø radar.
2006–05–19	Extensive convective precipitation on all radars. Some land clutter on Sindal radar.
2006–05–24	Precipitation on all radars and some clutter on Stevns and Sindal radars.
2006–08–18	Sea and land clutter on Stevns radar. Land clutter on Sindal radar. Precipitation front on all radar coming in from the south west.
2006–08–29	Extensive precipitation on all radars.

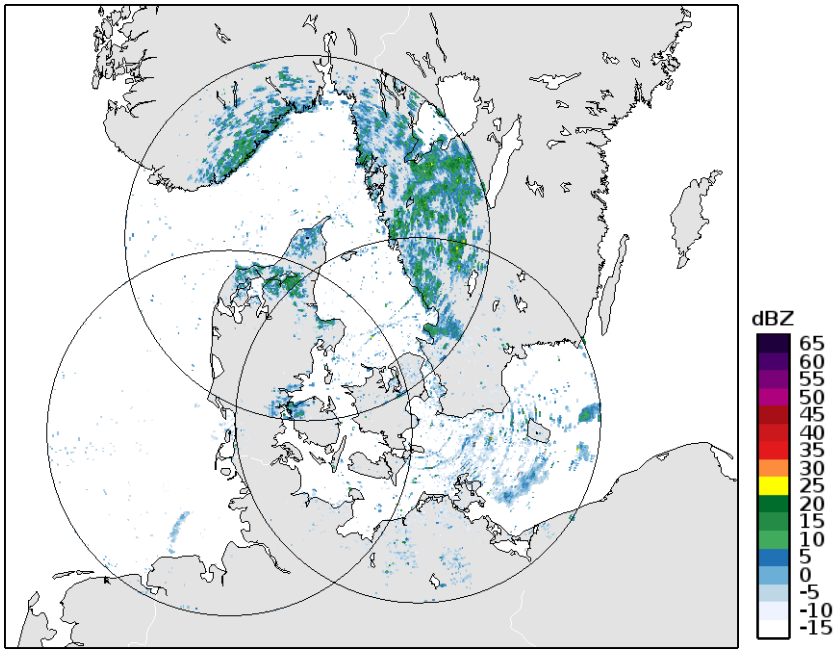

Figure 2.8: Case I. Radar reflectivity, 2005–05–11 06:00 UTC. Clutter, very little precipitation in the lower North Sea. Same colormap is used throughout the thesis for radar images.

Case I. Clutter, no precipitation. 2005-05-11 06:00 UTC

In the months of spring, summer and fall in Denmark, anomalous propagation conditions are quite frequent, giving rise to weather radar clutter due to superrefraction and ducting of the radar beams. Especially the radar at Stevns experiences AP clutter due to its location very close to the coast line of the Baltic Sea. Temperature inversions occur frequently here due the flow of hot and moist air masses over the colder ocean surface. Sea clutter is seen over a large part of the Baltic Sea. The radar at Sindal, at the time, was operating without applying the Doppler clutter filter. Therefore pronounced land clutter is seen on the land areas of Sweden and Norway. Some close range land clutter is also seen, especially on the northern most radar at Sindal. In the lower North Sea a little precipitation is seen.

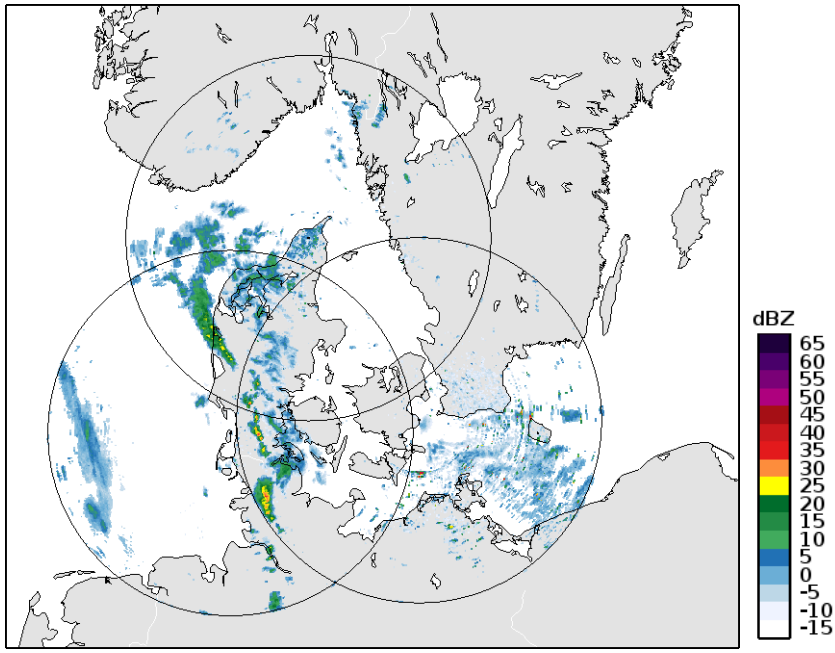

Figure 2.9: Case II. Radar reflectivity, 2005–09–25 18:00 UTC. Clutter and precipitation.

Case II. Clutter and precipitation. 2005-09-25 18:00 UTC

This exemplifies the case of AP clutter echoes and precipitation areas in the same radar image. Two precipitation areas are seen, one in the North Sea and one over Jutland. Sea clutter is observed in the Baltic Sea caused by anomalous propagation conditions caused by a temperature inversion which lasted throughout the day and into the night. In the North of Jutland, land clutter is seen close to the radar site.

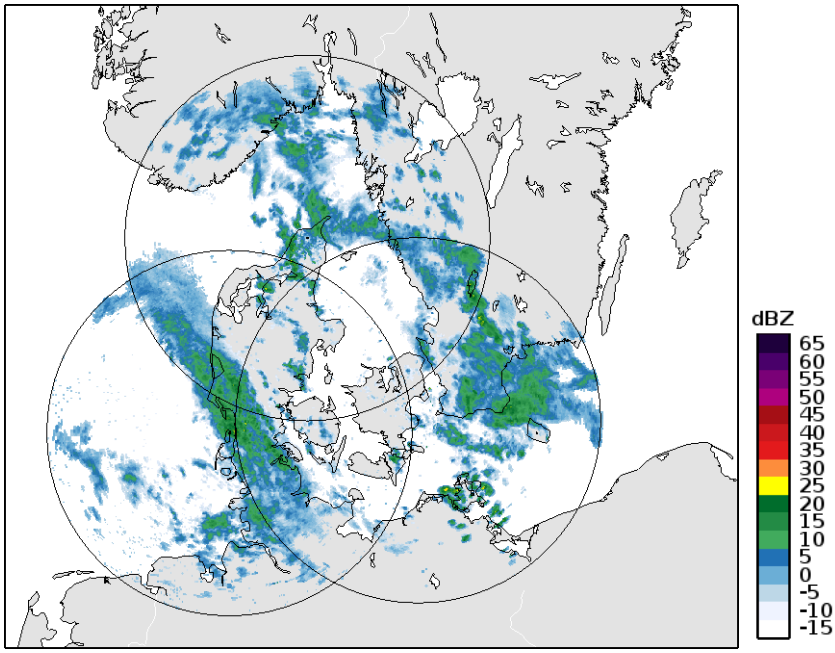

Figure 2.10: Case III. Radar reflectivity, 2006–05–19 12:00 UTC. Precipitation only.

Case III. Precipitation, no clutter. 2006-05-19 12:00 UTC

This case shows an event of mixed precipitation moving north-east extending over most of the radar coverage. No significant clutter was observed during this event. Only minor areas of close range ground clutter and some very minor mid to far range land clutter on the coasts of Norway and Sweden on the northern most radar at Sindal. This was only visible in a few images in the image sequence.

Chapter 3

Data fusion method

3.1 Introduction

Determining whether a radar echo is clutter or precipitation is not a trivial task and not possible using the radar reflectivity values themselves (at least not without making many errors). This can be realized by looking at the distributions of radar reflectivities for precipitation and clutter (Figure 3.1). The classes overlap, especially the land clutter and precipitation classes. This shows that other features are necessary to enable good discrimination of clutter and precipitation echoes, and in this chapter data from satellite images are used: 11 bands from the Meteosat Second Generation satellite and an operational nowcasting product derived from these called 'Precipitating Clouds'.

Previous studies using fusion of weather radar data and first generation Meteosat images have shown promising results. The improved spatio-temporal resolution of Meteosat Second Generation (commissioned in 2004) coupled with its increased number of spectral bands, is expected to make better estimation of precipitation possible [51], which potentially can yield improved clutter detection by also delineating areas of no precipitation better.

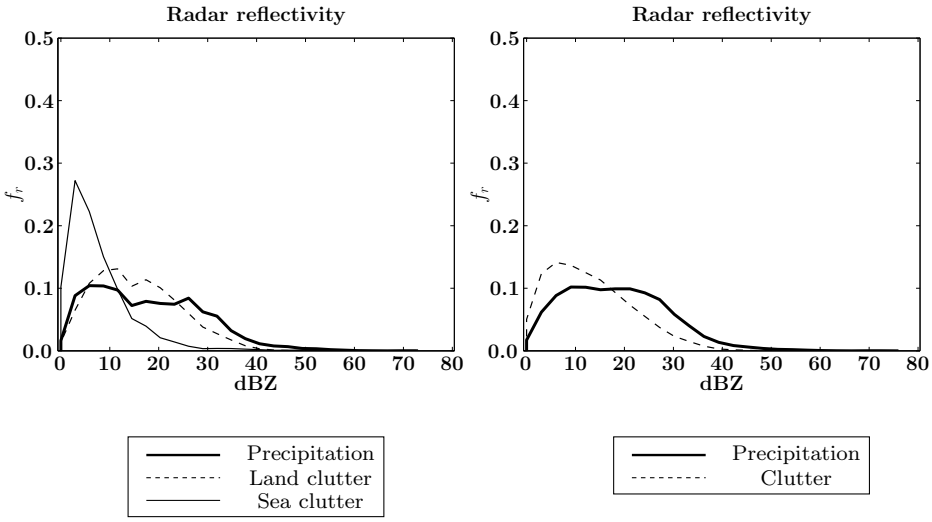

Figure 3.1: The distribution of radar reflectivity for clutter and precipitation classes. Based on a total of 1,140,697 samples. Two classes (left) and three classes (right).

3.2 Remote sensing work flow

For the development and the data analysis of the method, a remote sensing/image analysis work flow is used as illustrated in the flowchart in Figure 3.2. This chapter is structured to follow this flow:

Firstly, the input data (described in Section 3.3) are preprocessed (Section 3.5) for the extraction of features (Section 3.6). Then for the model design, training data is extracted (Section 3.7), the best sets of features are selected (Section 3.8), and the classification model is trained (Section 3.10). The data are then classified (Section 3.9) and the classification is evaluated (Section 3.11). Subsequently, in Section 3.13, the results of the classification are obtained.

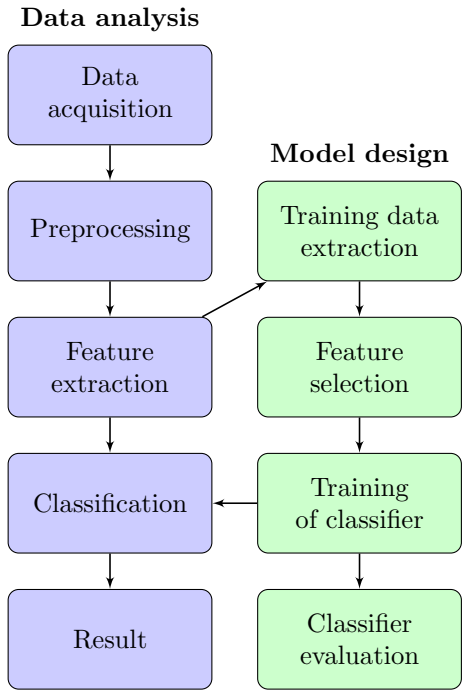

Figure 3.2: Flow chart of the general work flow for data analysis and model design.

3.3 Data: Meteosat Second Generation and Precipitating Clouds

Measuring and mapping of precipitation from space-borne platforms (in low-earth or geostationary orbits) can be performed using passive sensing of the visible and infrared parts of the electromagnetic spectrum as well as passive and active sensing in the microwave region (See [51] for an overview of the techniques). In this study, passive remote sensing in the visual and infrared spectrum from Meteosat Second Generation (MSG) was used in the form of 1) the raw 11 multispectral bands and 2) an operational meteorological product, Precipitating Clouds, derived from these images in combination with surface temperatures from a numerical weather prediction model.

For operational, real-time detection of clutter, the temporal and spatial resolution of the satellite imagery is of importance. The characteristics of low-earth orbiting platforms (typically in polar orbits of 800 km height) are a relatively high spatial resolution (e.g., values around 1 km for the NOAA and Metop meteorological platforms) but a low temporal resolution. For example, the temporal resolution of satellite imagery of the latitudes of Denmark from the NOAA polar-orbiting satellites is several hours. Such a low temporal resolution makes it impossible to detect clutter in radar images with their temporal resolution of 10 minutes. Therefore, a geosynchronous satellite platform is the only feasible choice for clutter detection: Geosynchronous platforms provide excellent temporal resolution (15 minutes for Meteosat Second Generation) comparable to the one of weather radars, but this is at the trade-off of a lower spatial resolution (a factor three or more lower than for low-earth orbiting satellites).

The SEVIRI (Spinning Enhanced Visible and Infra-Red Imager) sensor on board Meteosat Second Generation (Meteosat-8 onwards) [68] provides multispectral images from a geostationary orbit of 36.000 km height. Both the spatial and temporal resolution has been significantly improved compared to Meteosat-7. The spatial resolution at nadir is 3 km for the 11 visible, near infrared and thermal infrared channels and 1 km for the panchromatic channel. Meteosat-7's figures were 6 km and 3 km respectively for only three channels (visual, infrared, and water vapor). At the latitudes of Northern Europe (around 50° N) the resolution of MSG data decreases from the 3 km of the multispectral channels to approximately 5 km due to the imaging geometry. The temporal resolution of the data is 15 minutes. For this study only the 11 low-resolution bands were used. Their response curves over the electromagnetic spectrum can be see in Figure 3.3. The visual channels 1 to 3 are measured in radiance (units $\mathrm{mW}/(\mathrm{m}^2 \cdot \mathrm{sr} \cdot \mathrm{cm}^{-1})$), whereas the near infrared bands (channels 4 to 11) are measured in brightness temperatures in units of Kelvin.

In Figure 3.4 an example of the visual channel 2 can be seen and the 11 MSG bands can be seen in Figures 3.5, 3.6, and 3.7 for the three case examples. The bands can be combined and visualized as false-color images as shown in Figure 3.8, where the corresponding radar images for the same times are shown for comparison.

Both the visual and infrared bands of Meteosat Second Generation imagery provide very detailed information on clouds and precipitation in the atmosphere [51]. From the reflected sunlight in the visual/near-infred and the emitted thermal energy, several parameters relating to precipitation can be derived: cloud types, cloud water vapor, cloud heights, etc. Some of the uses of each MSG band for precipitation mapping as reported by [51] are summarized in Table 3.1.

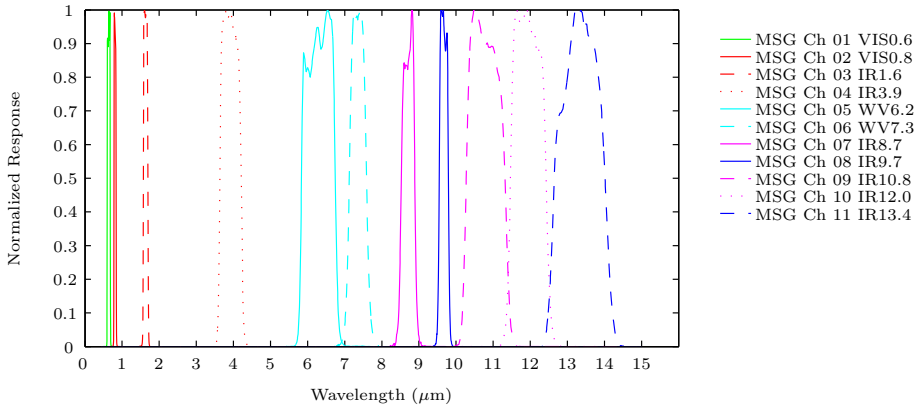

Figure 3.3: MSG spectrum response functions.

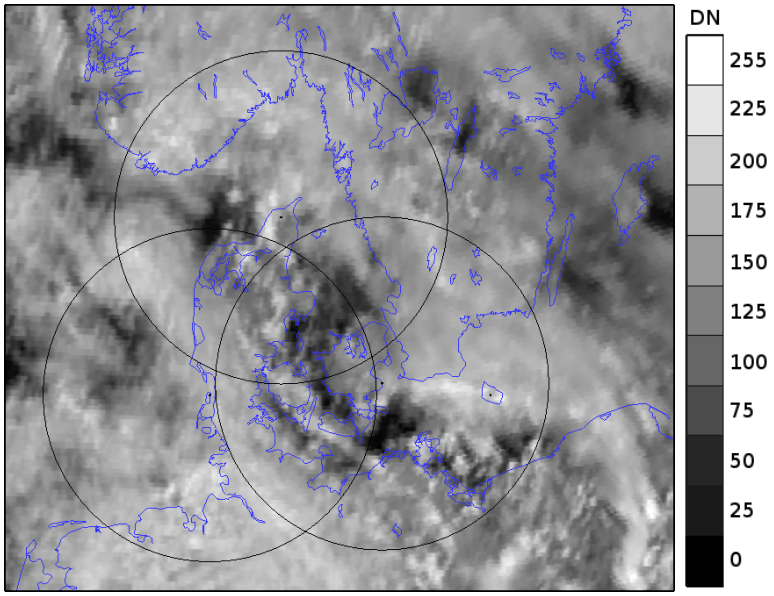

Figure 3.4: MSG, visual channel 2. Same colormap is used throughout the thesis for MSG images. The radar sites and coverages are also shown for easier comparison with the radar case examples.

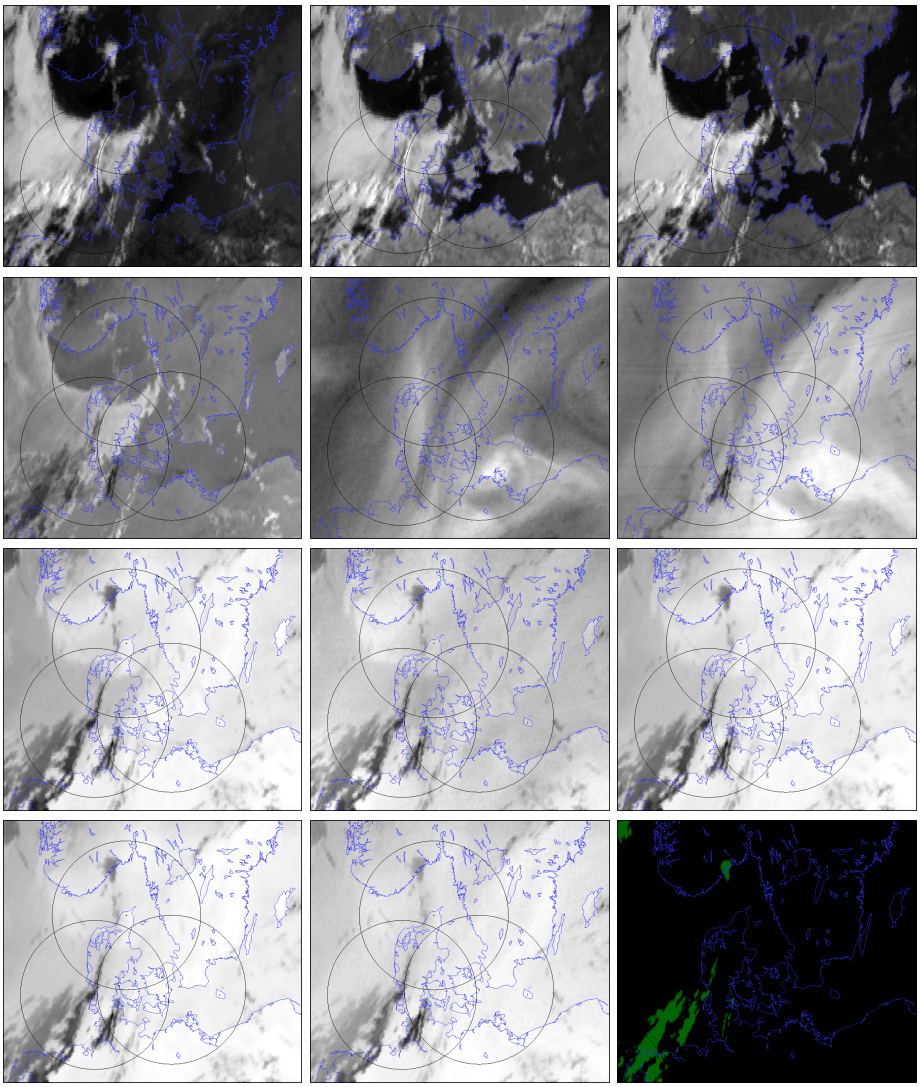

Figure 3.5: Channels 1 through 11 of MSG from upper left, row-wise. The lower right image is the Precipitating Clouds product. Case I, 2005–07–11 06:00.

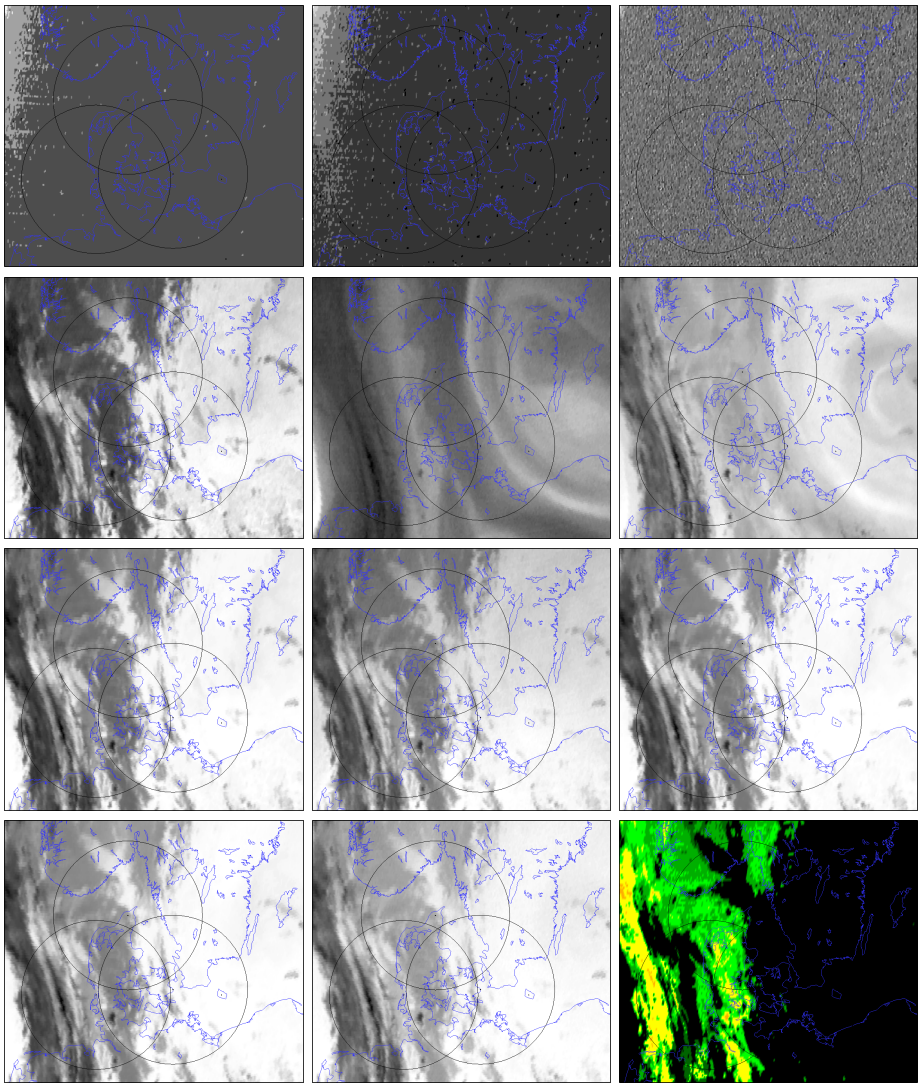

Figure 3.6: Channels 1 through 11 of MSG from upper left, row-wise. The lower right image is the Precipitating Clouds product. Case II, 2005–09–25 18:00.

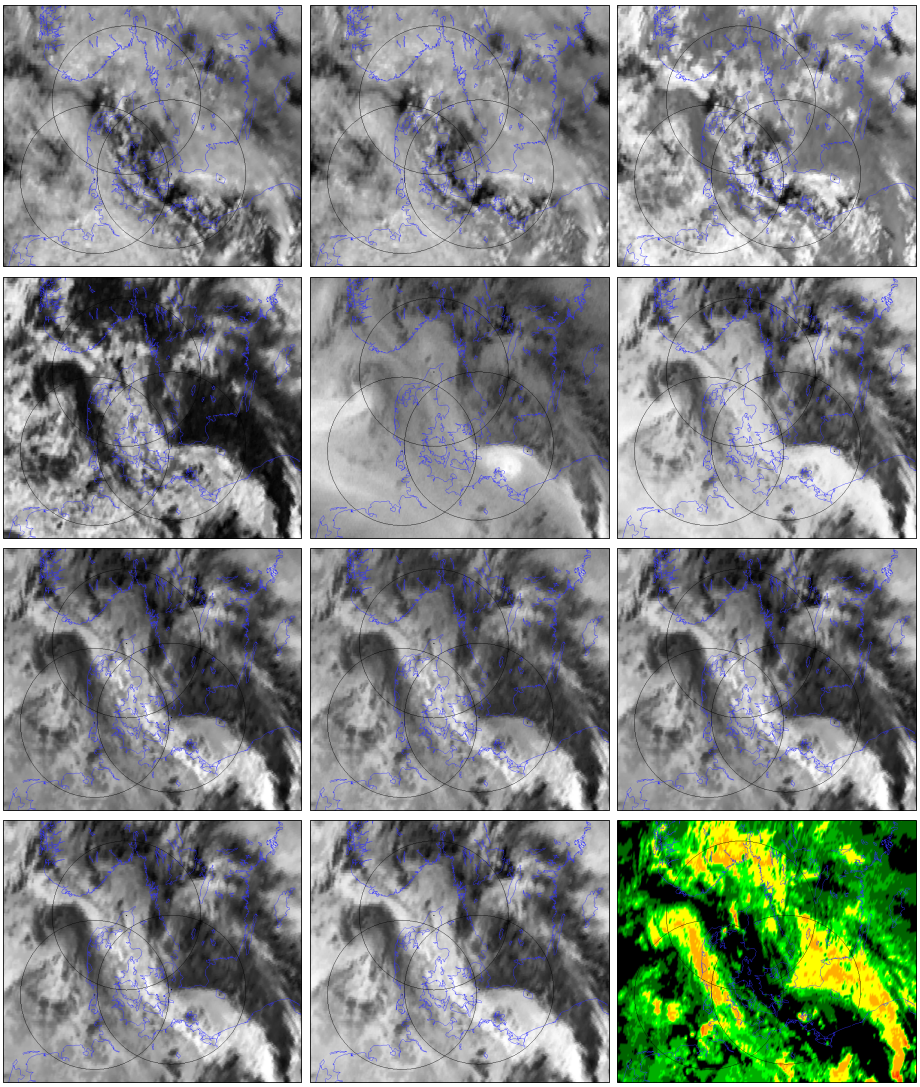

Figure 3.7: Channels 1 through 11 of MSG from upper left, row-wise. The lower right image is the Precipitating Clouds product. Case III, 2006–05–19 12:00.

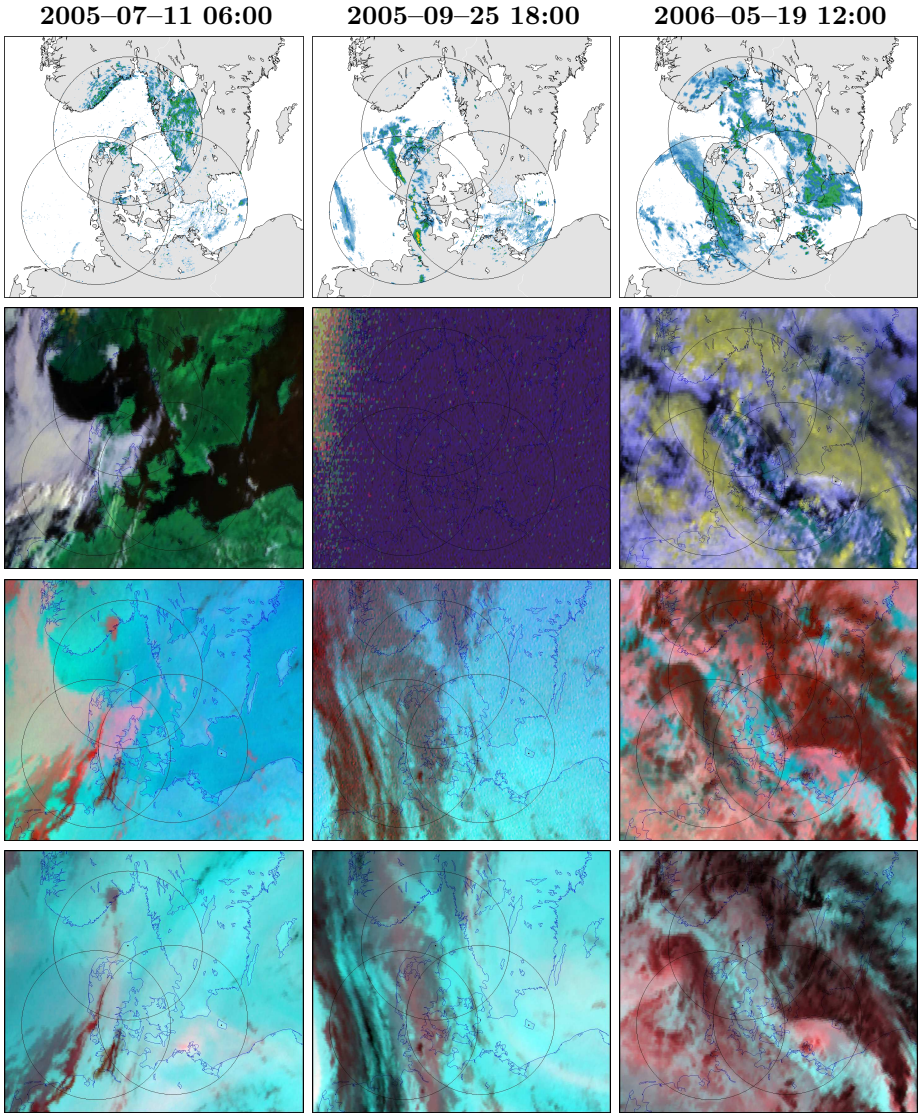

Figure 3.8: Radar reflectivity images (top row) and MSG false color composites (As red, green, and blue: Second row: channels 1, 2, 3. Third row: channels 3, 8, 9. Bottom row: channels 5, 9, 10.

Table 3.1: MSG bands. Channel number (Ch), center wavelength (λ), and use for precipitation and cloud mapping. After [51].

Ch	λ (µm)	Description	
1	0.6	Visible band (VIS)	Valuable for discriminating between clouds and no clouds in day time. The optical depth of clouds can be inferred from the visual channels.
2	0.8	Visible band (VIS)	—"—
3	1.6	Near-infrared band (NIR)	For discrimination between snow and cloud, and ice and water clouds. The IR reflectance holds information on the cloud particle's effective radius.
4	3.9	Infrared band (NIR)	Detection of low clouds and fog. Droplet mean radius. Clouds with small precipitation particles reflect more in this band than clouds with larger particles.
5	6.2	Water vapor band (WV)	Observation of water vapor in upper troposphere.
6	7.3	Water vapor band (WV)	—"—
7	8.7	Infrared band (IR)	Optical depth of thin cirrus clouds and discrimination between ice and water clouds.
8	9.7	Ozone band (O_2)	Observation of ozone. No immediate use for precipitation estimation.
9	10.8	Infrared band (IR)	Split window channel for cloud detection and cloud top temperature.
10	12.0	Infrared band (IR)	—"—
11	13.4	Carbon dioxide band (CO_2)	Usage in determining heights of transmissive cirrus clouds mentioned.

Precipitating Clouds product

Within the 'Nowcasting SAF (Satellite Application Facility)' of EUMETSAT[1] a range of meteorological products, e.g., cloud masks and cloud type products, are developed and produced for short-term weather forecasting. Of special interest for detection of weather radar clutter is the product named 'Precipitating Clouds' which provides the probability of precipitation for each pixel in the MSG pixel grid [59]. This is done by linear combination of the spectral bands of MSG together with surface temperatures extracted from a NWP model. The surface temperature helps distinguish between land surface and clouds especially in the winter time when snow covered areas resemble the cloud top temperature.

The 'Precipitating Clouds' algorithm consists of a day and a night part with the latter excluding the visual channels. The parameters of the model were developed and tuned using rain gauge data with the option of tuning to weather radar data. The version of the 'Precipitating Clouds' algorithm used to generate the data used in this study (version 1.1), however, was not tuned with radar data [59]. Figure 3.9 shows the 'Precipitating Clouds' product corresponding to Figure 3.4, 2006–05–19 12:00). The 'Precipitating Clouds' product is similarly shown for the case example dates in Figures 3.5, 3.6, and 3.7.

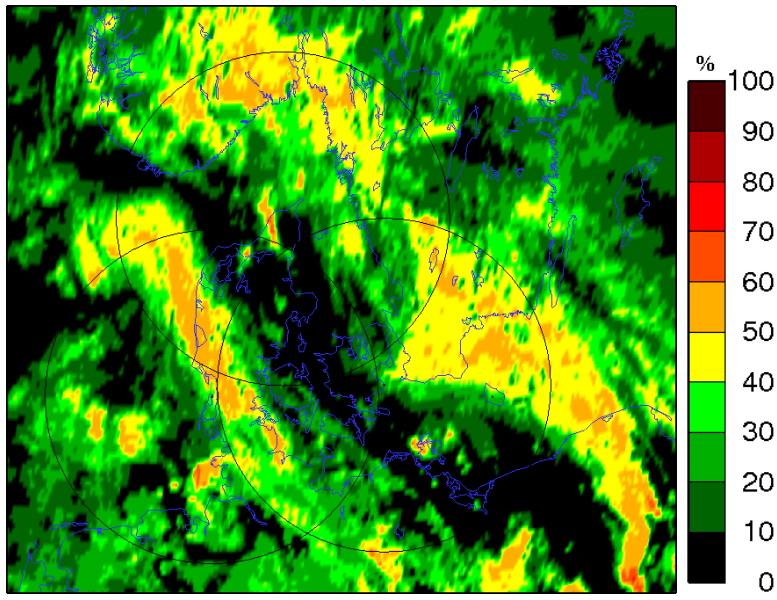

Figure 3.9: Precipitating Clouds product example. Same colormap is used throughout the thesis for this image type.

[1] European organization for collaboration on operational meteorological satellites.

3.4 Method: Data fusion method

In the following section, the proposed data fusion method for clutter detection is described. The method uses the above two sources of data 1) the 11 raw MSG bands and 2) the 'Precipitating Clouds' product to detect clutter in weather radar images. The method uses supervised classification as outlined in the flow chart in Figure 3.2. The MSG and 'Precipitating Clouds' data sets are evaluated separately because the 'Precipitating Clouds' dataset is derived directly from the other and to allow comparison of the results using the two data sets separately. Furthermore, the 'Precipitating Clouds' product is univariate and does not need the feature selection procedure.

3.5 Preprocessing

The preprocessing step for the method includes the fusion of the data. Since both radar and satellite data are images, the process is *image fusion*, which can be performed at various levels. A recent review of image fusion methods [63] groups the methods into the categories *pixel level*, *feature level* and *decision level*. For this study a *pixel level* image fusion was chosen. Fusing images on the pixel level requires resampling of the images to a common grid, and here a stereographic map grid with a grid spacing of 1000 m was chosen. This means oversampling the satellite data from approximately 500 m to 1 km. As described in the previous chapter the radar data has a varying resolution from below 500 m to over 4000 m and resampling to 1 km thus involves both over and undersampling. All of the datasets, the radar data, the MSG images, and the 'Precipitating Clouds' image were resampled to this grid using bicubic interpolation.

For remote sensing data fusion applications, the aspects of alignment, both spatially and temporally, is of concern. Using multi-sensor data, as is the case here, it is for multiple reasons difficult to achieve perfect alignment. Difference in imaging geometry, difference in data acquisition timing, and differences in sensing technique are the main reasons to the misalignment problems, as described in the following.

Imaging geometry

The spatial alignment of *the images* from satellite and radar can be carried out with good precision because the imaging geometry of the two sensors is well-known. This can be seen by visual inspection by superimposing the fused

images: the coastlines in the satellite image align well with observed land clutter along the coastlines. However, for both radar and satellite there are issues which can lead to misalignments between the precipitation and clouds observed by the two sensors (See Figure 3.10).

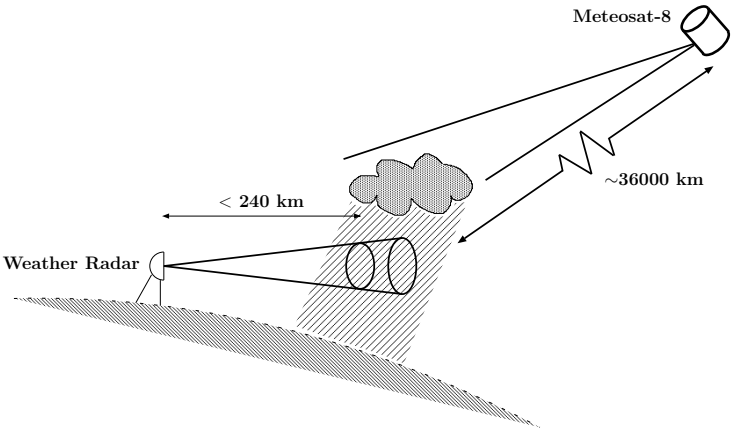

Figure 3.10: Sketch of the difference in sensing geometry of ground-based radar and space-borne multispectral observations.

For the radar, precise determination of the location of a radar echo depends on the propagation path of the radar rays. Under standard propagation, the propagation path is known, but non-standard propagation is quite common and the location the radar echo is less certain (especially the height determination is uncertain). However, the displacement error from non-standard propagation is generally considered to be small.

For a geosynchronous satellite platform, however, *parallax displacement* will cause the clouds to be offset quite severely from their correct position in the image. The displacement is a function of the height of the cloud and the latitude [42] and for 50° N the displacement is approximately 1.6 times the height of the cloud. For example, for a cloud of 2 km height the displacement is 3.2 km. Methods for correction of the displacement due to the parallax are widely used in the field of photogrammetry, but less widely used in remote sensing meteorology.

Temporal misalignment

Temporally, the datasets are not always perfectly aligned due to non-synchronous data acquisition. Therefore features in the images (precipitation and clouds) do

not align since they have moved in between the time points of data collection. The radar data are acquired every 10 minutes and the satellite data every 15 minutes as sketched in Figure 3.11. From the figure it can be seen that for real-time clutter detection, a temporal misalignment of 10 minutes will occur for two of the six radar images per hour. Another two will experience misalignment of 5 minutes and the remaining two will have no misalignment. For clutter detection using archived data the maximum misalignment can be lowered to 5 minutes.

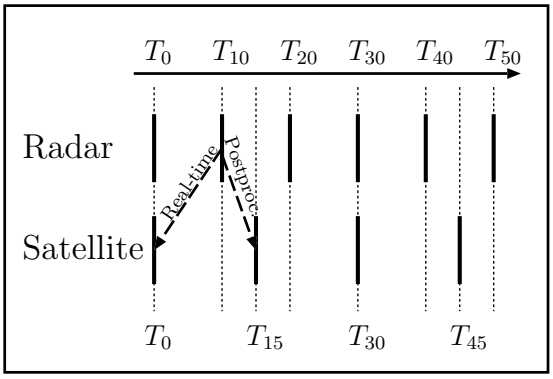

Figure 3.11: Time schedule for radar and satellite data. For real-time analysis the most recent satellite image must be used. For postprocessing of archived data, the image 5 minutes ahead in time can be used for better temporal alignment.

The actual misalignment of the clouds and precipitation in the data depends on the velocity of those features. For a velocity of clouds of $10\,\mathrm{m}\cdot\mathrm{s}^{-1}$, for example, the resulting mismatch between images that are 10 minutes misaligned will be 6 km. Correction of the temporal misalignment could be carried out by using techniques from image extrapolation (e.g., using optical flow). However, this was not pursued further.

Sensing differences

Besides spatial and temporal misalignment, also the fact that the two sensors observe different phenomena can contribute to the mismatching features in the images. The radar measures backscatter from the precipitation itself whereas the satellite measures the reflected and emitted energy from cloud top.

The combined effect of all these sources of misalignment cause a smaller or greater misalignment in the images which eventually can result in degraded performance of the clutter detection method.

Scale-space misalignment mitigation

A number of fused image sets of satellite images and radar images—both of perfect temporal alignment and of 10 minutes offset—were inspected visually and the observed precipitation does not always align well with the location of clouds in the MSG data or the estimated precipitation areas in the 'Precipitating Clouds' product. It is seen that the mismatch between radar and satellite data is most pronounced at the borders of precipitation areas. In the paper [19], a preliminary investigation of a pragmatic remedy to mitigate the border effects of misalignment was presented. Before classification, an expansion of the feature space by application of a scale-space was performed. In many applications of image analysis, a scale space approach can be fruitful for image segmentation, feature extraction and classification [53]. By including a scale-space representation of the satellite images (by smoothing of the images with a Gaussian kernel) into the classification it was possible to identify a unique scale which improved the classification accuracy. Although computed on a quite limited data set, the work outlines a method for mitigating the misalignment problem.

3.6 Feature extraction

The feature extraction step for the method is limited to extraction of the 11 raw MSG bands and the 'Precipitating Clouds' product. The MSG data were kept in their original values of digital numbers (DN) as stored in their 8 bit data format. Conversion of the MSG data into physical units of radiance and brightness temperatures could have been carried out by a linear scaling, however, this was omitted as it would have no influence on the classification results. The Precipitating Clouds product was converted into values of percentage from 0 to 100.

The visual and near-infrared MSG channels, of course, have little use in the night time (See Figure 3.6, top row). Therefore the features were grouped into the sets *day features* and *night features*. The day features contain all 11 bands while the night features excludes the visual and near-infrared channels.

Some methods, e.g., [58] and [48], use in addition to the satellite data, the air temperature at ground level from a numerical weather model. The motivation for this is for the algorithm to work in cold climates where a snow covered ground surface will have thermal infrared brightness temperatures comparable to that of cloud top temperatures. For this study, however, this situation was not taken into account, but extending the list of features with surface temperature is possible.

3.7 Training data

Supervised classification requires a training set of input features to learn the parameters of the classification model. From the database of radar data from 15 days of various meteorological events (listed in Table 2.3), four images per day at the times 01:00, 06:00, 12:00 and 18:00 UTC were chosen. The times for the training data (at the top of the hour) were selected to ensure the best possible temporal alignment of the radar and satellite data. This ensures the best training data for the model by eliminating the temporal misalignment.

The images were annotated by manual image inspection and areas of the three classes, precipitation, sea clutter, and land clutter were outlined as shown in Figure 3.12.

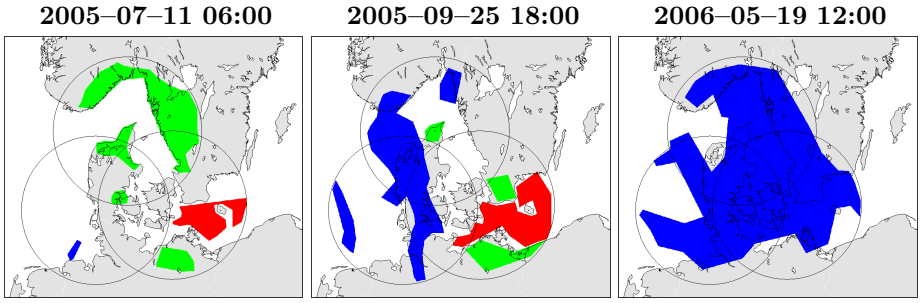

Figure 3.12: Example of manually delineated training areas. Red is sea clutter, green is land clutter, and blue is precipitation. Compare with radar images in Figures 2.8, 2.9, and 2.10

Training areas were drawn around as many pixels in every scene as possible and covered also the edges of the precipitation areas. This is considered important for the unbiased evaluation of the classification in using multiple datasets with the risk of misalignment (See Figure 3.13). Of course the extracted training data for building the classification models will be fitted to some erroneous data but the amount of these outliers is small compared to the total number of training samples and the models will be able to disregard these outliers. Ideally, the models should be built using training data from the interior of the features and evaluated on the entire training set.

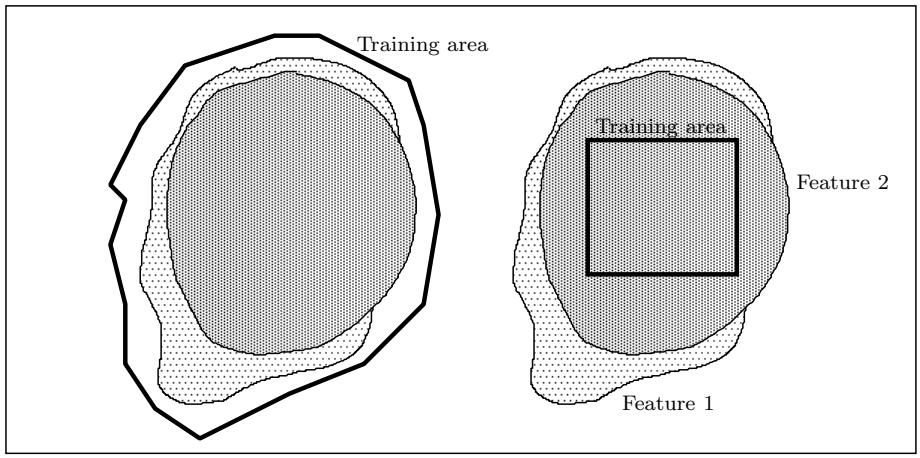

Figure 3.13: Left: The used approach to training areas includes also misaligned areas between the features. Right: The common procedure of selecting training areas includes only the overlap between the features.

The total number of training samples in the selected training areas is shown in Figure 3.14. Notice how the sea clutter and land clutter classes can be combined into a superset clutter class (for two-class classification), and that each class is finally grouped into weak and strong radar echo classes. This was done to enable evaluation of the classification methods for weak and strong precipitation. Weak echoes are echoes of radar reflectivities below 15 dBZ.

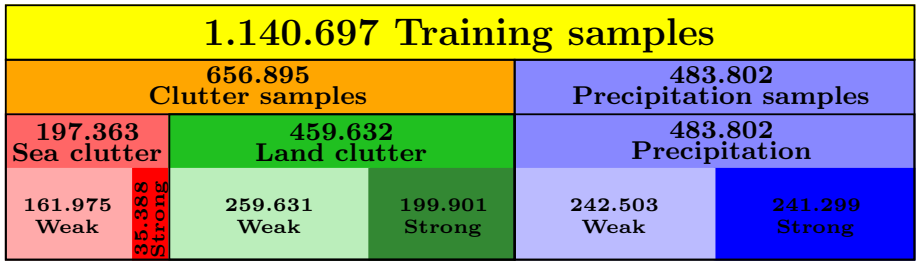

Figure 3.14: Number of training samples in total (top row), Two class case: clutter and precipitation classes (second row), Three class case: Sea and land clutter and precipitation (third row), and in bottom row the number of weak and strong echoes for these classes.

3.8 Feature selection

This section deals with selecting the best combinations of the MSG bands for classification of clutter. Firstly, it is customary to construct histograms for the features and classes to get insight in the distribution of the data.

3.8.1 Feature histograms

Using the training areas, training data were extracted from the satellite data for the precipitation and clutter classes. For each band in the MSG data and the Precipitating Clouds product, relative frequency histograms were made. In Figures 3.15 and 3.16, histograms for all 11 MSG and the 'Precipitating Clouds' product can be seen for the three class and two class case for, day features, and weak and strong echoes combined (called *all echoes* in the following). For the two class case, the land and sea clutter classes are simply treated as one class.

It is seen how the clutter classes generally have low radiance in the visual and near-infrared (channels 1 to 3) and high temperatures in the infrared (channels 4 to 11) while the opposite pattern is seen for precipitation. This is because the clutter class mostly occurs in areas of no clouds, which means it is the reflection and transmitted energy from the land or sea surface which are warmer and darker than clouds. Noticed is also how some features, channel 1 for example, show good separation of the classes while some overlap more (channel 5, for example).

The 'Precipitating Clouds' product histograms in the lower right of the figures, show distinct spikes which are the result of the original data being quantized into intervals of 10 percent.

Class conditional probabilities were computed from the feature histograms, see examples in Figures 3.17 and 3.18. These figures show the probability of a given sample value of a particular feature being precipitation or clutter. For example, in Figure 3.18, the top left plot shows class conditional probabilities for MSG Ch 01. The plot shows that for values in MSG Channel 1 below approximately 60 the probability of clutter is highest and for values above 60 the probability for the precipitation class is highest.

A complete collection of feature histograms and conditional probability curves and can be found in Appendix B, Figures B.1 to B.14 and Appendix C, Figures C.1 to C.14, where plots for all combinations of day and night features and all, weak, and strong echoes, and the two and three class cases are shown.

3.8.2 Feature selection

The contents of the feature histograms support the statements summarized in Table 3.1 that some bands seem to provide more information on precipitation than others. In this study, however, it was chosen to not make any presumptions of the usefulness of each feature. Instead, all were treated as potentially useful and an automated process of feature selection was employed.

The approach to simply lump all the features together into one large features space might be tempting for reasons of ease. However, this might create an unnecessarily complex model and it is often seen that using too many features degrades the classification accuracy. This is known as the *Hughes Phenomenon* which is the degradation of the classifier performance with increasing data dimensionality [64]. Therefore features which do little to improve—or even degrade—the classification accuracy should be left out of the final classification model. This can be accomplished through *subset selection* which aims at selecting the best subset from a full set of features. A wide range of methods for subset selection exists, see e.g., [36] and [64]. In this work, three methods for feature selection were applied: forward selection, backward elimination, and exhaustive search.

Forward selection

Forward selection is carried out by starting out with selection of the single best features and then adding features one by one, choosing the one which improves a selection criterion. Here the selection criterion was chosen to be the classification accuracy from using a quadratic discriminant function as classifier. The model was trained on 25,000 training samples for each class and evaluated for the same number of test samples for each class. It is not guaranteed that forward selection yields the best subset, because even if a feature is the best one to use on its own, it is not necessarily the best one to use in combination with the other subsequently added features.

Backward elimination

This method works the other way around. First, the classification accuracy is computed using the full set of features. Then, one by one, the features which decreases the accuracy the least (i.e., contribute the least to the classification accuracy) are removed. Like in the case of forward selection, backward elimination does not guaranteed finding the subset with the highest classification accuracy. Backward elimination was performed with the quadratic discriminant function classifier as selection criterion.

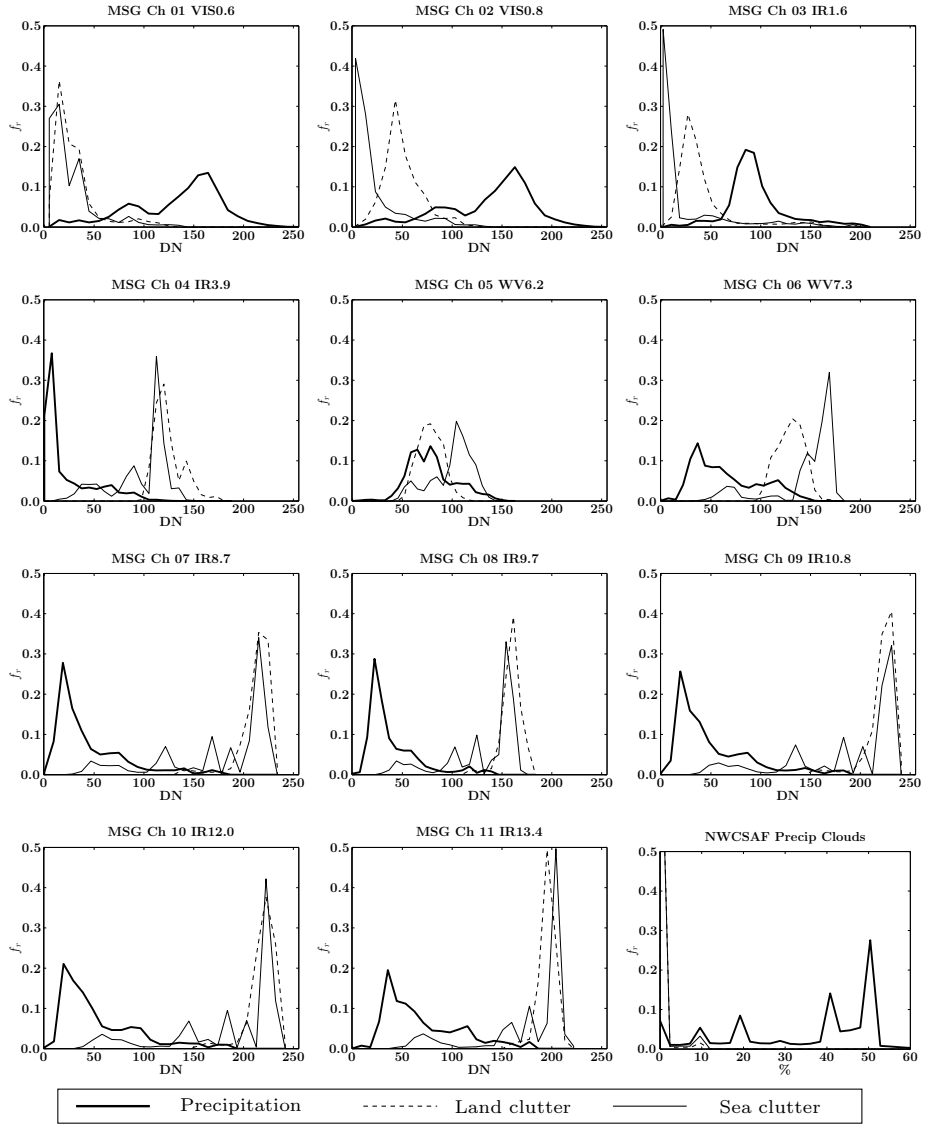

Figure 3.15: Data Fusion feature histograms. Day features and all echoes. Three classes.

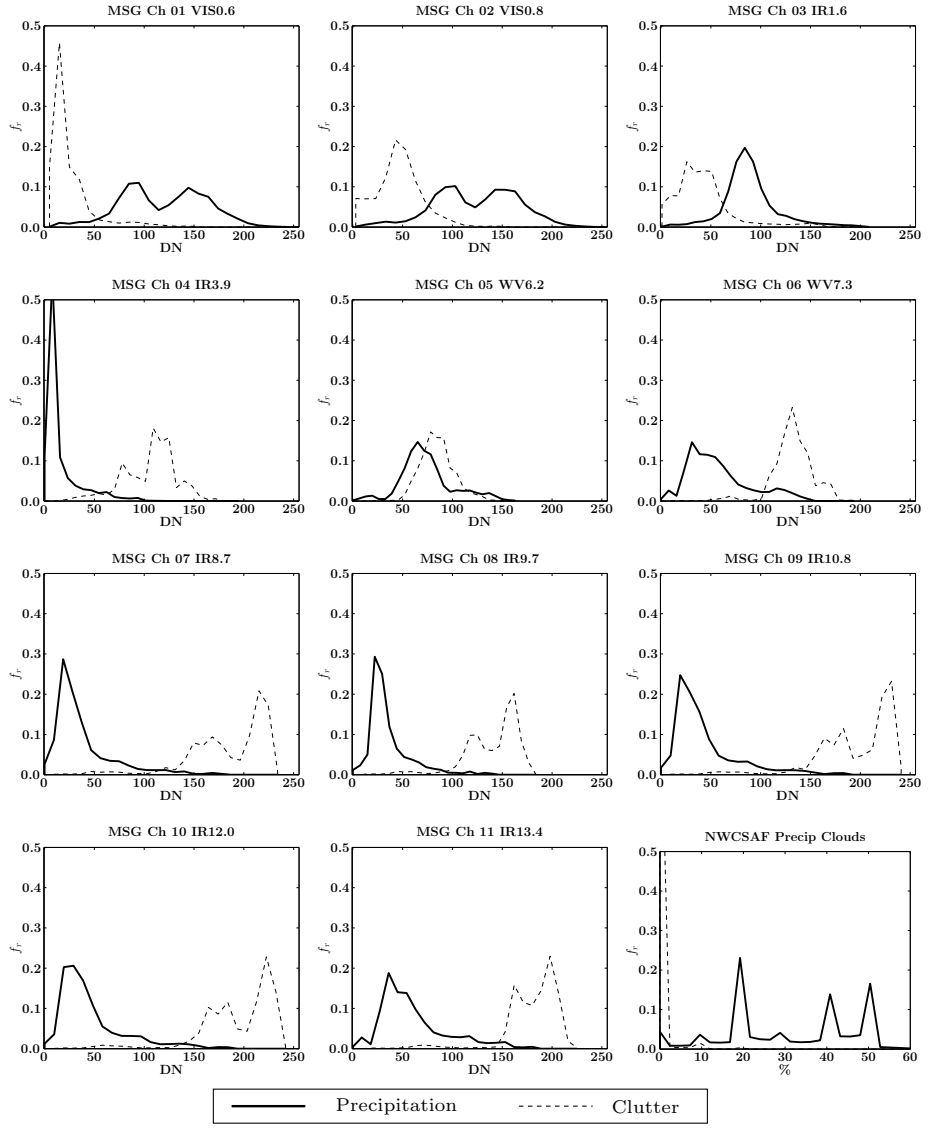

Figure 3.16: Data Fusion features. Day features and all echoes. Two classes.

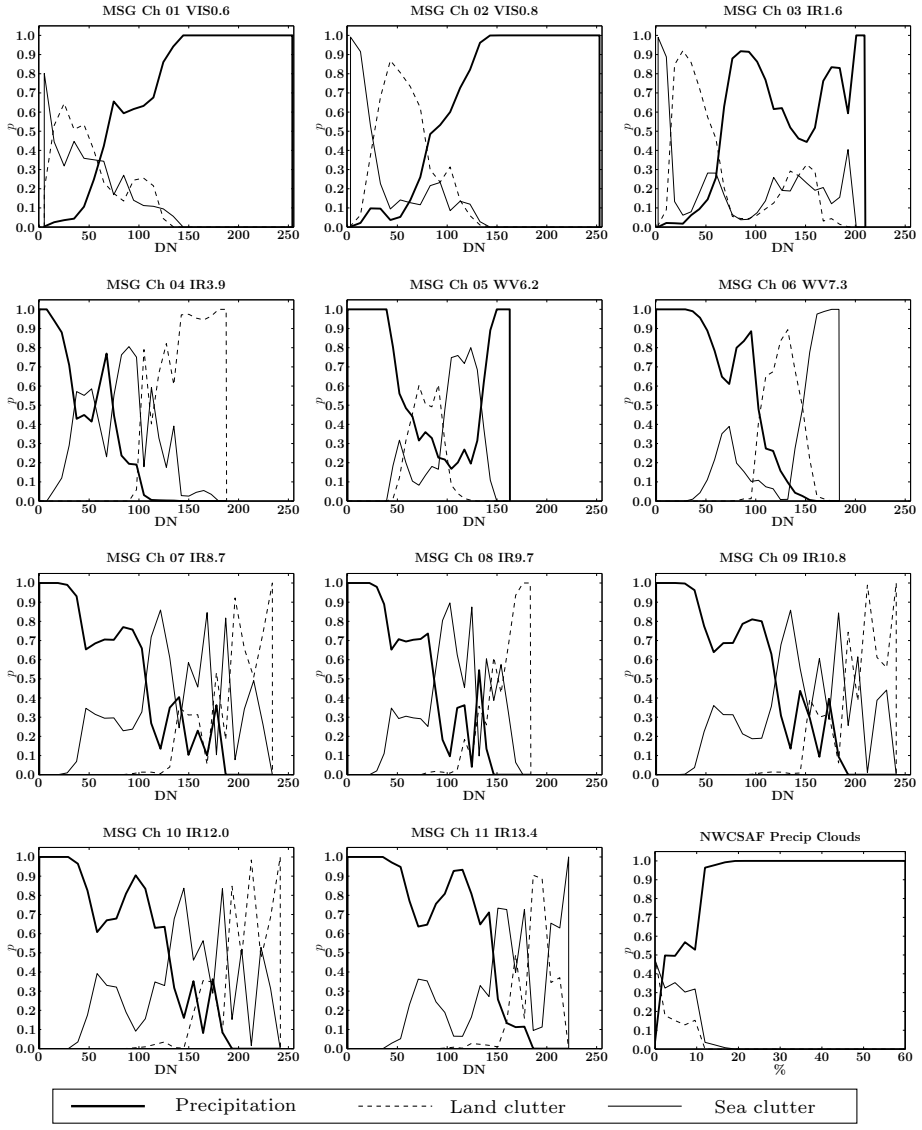

Figure 3.17: Data Fusion features. Day features and all echoes. Three classes.

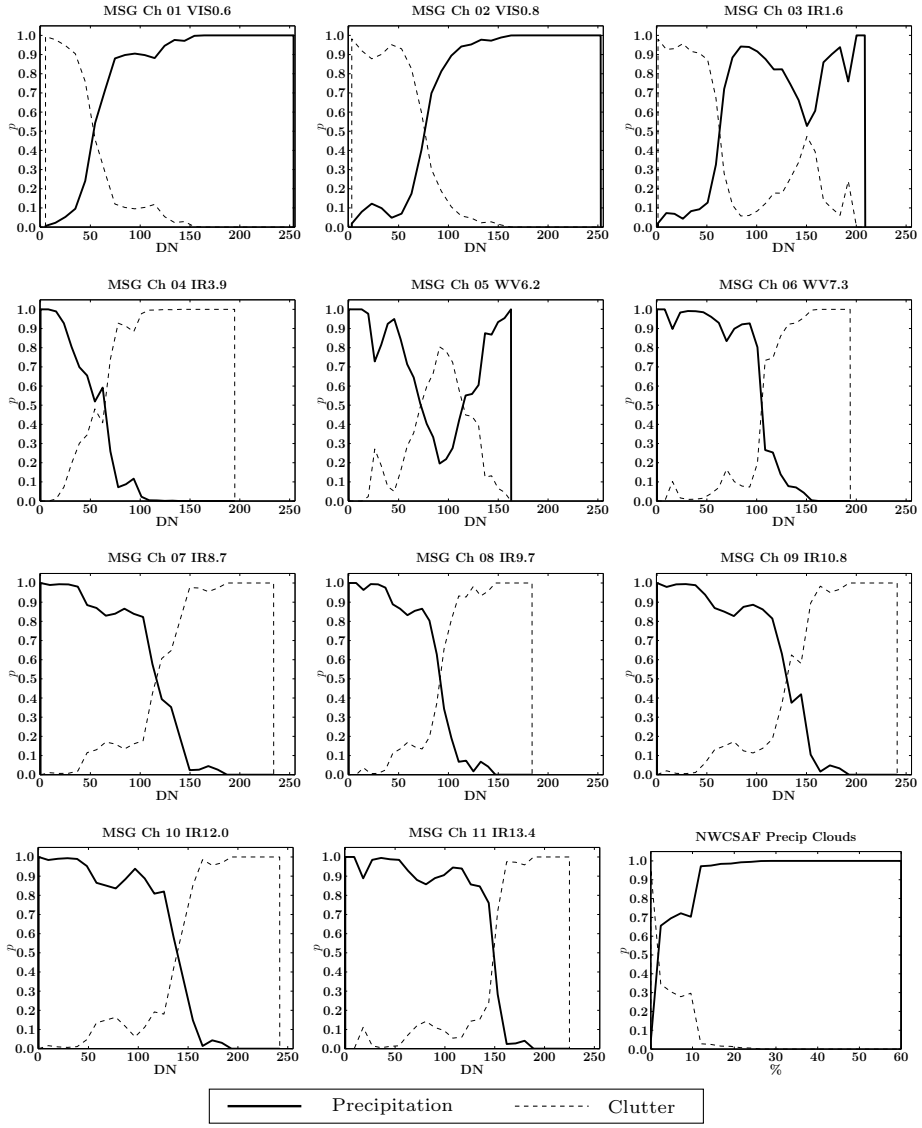

Figure 3.18: Data Fusion features. Day features and all echoes. Two classes.

Exhaustive search

This method for feature selection simply evaluates all possible combinations of features. Exhaustive search guarantees to find the optimal subset of features, however, the method is very computationally heavy for a high number of features since all possible combinations of all numbers of features are computed.

As selection criteria, the quadratic discriminant function was used as described above as well as the Jeffries-Matusita distance:

Jeffries-Matusita distance

The Jeffries-Matusita distance [64] expresses the average distance between two class density functions and is a measure of the separability between the two features. For each number of features, the pairwise distances are computed and the ones with the greatest distances are chosen. In comparison to using the quadratic discriminant function as criterion, no classification of the data is carried out using the Jeffries-Matusita distance. However, as with the use of the discriminant functions, a Gaussian distribution of the data is assumed.

Results

The three feature selection methods were applied to the training data for the MSG bands and for the combinations of day features and night features, and all, weak, and strong echoes. Furthermore, feature selection was performed for the three class case (treating land and sea clutter as separate classes) and for the two class case (where land and sea clutter are combined into one class).

An example of feature selection results for the data fusion method is shown in Table 3.2. The top table lists the highest classification accuracy in separation of clutter and precipitation for increasing number of features. In italics are highlighted the best classification accuracy obtained and in bold is the *chosen* accuracy highlighted which is either the highest accuracy or the feature combination with the fewest number of features closer than 0.2 % to the highest accuracy. The numbers in parenthesis are the features listed in the lower left table (in this table the feature numbers which were chosen are shown in bold).

The finally chosen subset, which is used in the further classification, is selected from the exhaustive search method and again the chosen subset selection is the subset of features which is either the one with the single highest accuracy or the one with the fewest features closer than 0.2 % to the highest accuracy. The graph on the lower right shows the feature selection curves of each feature selection method and can be used to compare the forward selection, backward elimination, and exhaustive search methods.

In Appendix D, Tables D.1 to D.12, the feature selection tables for all combinations of day and night features and the different echo strengths are shown.

Table 3.2: Feature selection results. Data fusion method, day features, all echoes, two classes.

# of	Classification accuracy (Feature combination)			
feats.	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	95.6 % (8)	95.0 % (10)	(8)	95.6 % (8)
2	96.2 % (5 8)	96.1 % (5 10)	(1 10)	96.2 % (5 8)
3	96.2 % (5 8 9)	96.5 % (5 9 10)	(1 2 10)	96.5 % (5 9 10)
4	96.6 % (5 8 9 10)	96.8 % (5 6 9 10)	(1 2 10 11)	96.8 % (5 6 9 10)
5	**96.8** % (5 6 8 9 10)	**96.9** % (1 5 6 9 10)	(1 2 5 6 10)	**96.9 % (1 5 6 9 10)**
6	96.8 % (1 5 6 8 9 10)	*97.0* % (1 3 5 6 9 10)	(1 2 3 5 6 10)	*97.0 % (1 3 5 6 9 10)*
7	*96.9* % (1 4 5 6 8 9 10)	97.0 % (1 2 3 5 6 9 10)	(1 2 3 5 6 7 9)	97.0 % (1 2 3 5 6 9 10)
8	96.8 % (1 4 5 6 7 8 9 10)	97.0 % (1 2 3 5 6 8 9 10)	(1 2 3 4 5 6 7 9)	97.0 % (1 2 3 5 6 7 9 10)
9	96.7 % (1 4 5 6 7 8 9 10 11)	97.0 % (1 2 3 5 6 8 9 10 11)	(1 2 3 5 6 7 9 10 11)	97.0 % (1 2 3 5 6 8 9 10 11)
10	96.7 % (1 3 4 5 6 7 8 9 10 11)	96.9 % (1 2 3 5 6 7 8 9 10 11)	(1 2 3 5 6 7 8 9 10 11)	96.9 % (1 2 3 5 6 7 8 9 10 11)
11	96.8 % (1 2 3 4 5 6 7 8 9 10 11)	96.8 % (1 2 3 4 5 6 7 8 9 10 11)	(1 2 3 4 5 6 7 8 9 10 11)	96.8 % (1 2 3 4 5 6 7 8 9 10 11)

Feat. #	Feature name
1	MSG Ch 01 VIS0.6
2	MSG Ch 02 VIS0.8
3	MSG Ch 03 IR1.6
4	MSG Ch 04 IR3.9
5	MSG Ch 05 WV6.2
6	MSG Ch 06 WV7.3
7	MSG Ch 07 IR8.7
8	MSG Ch 08 IR9.7
9	MSG Ch 09 IR10.8
10	MSG Ch 10 IR12.0
11	MSG Ch 11 IR13.4

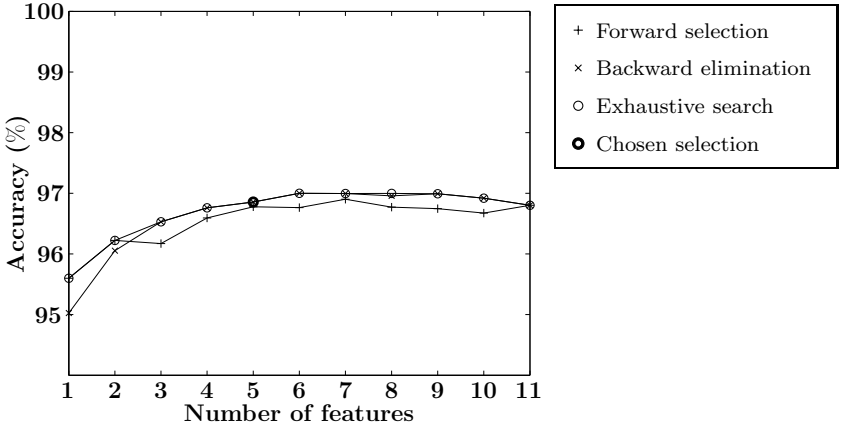

The points made in the above sections on the differences between the feature selection methods are demonstrated in the figures. For example, in Table 3.2 the selected feature in forward selection for using one feature only is feature 8. This feature is not in the finally chosen subset from exhaustive search (features 1 5 6 9 10). In fact, feature 8 is only good up to two features and then it is not introduced until the use of nine features. The Hughes Phenomenon is present in the plot of number of features against classification accuracy (the accuracy increases until six features are used and stays constant until it drops off slightly at 10 and 11 features).

Looking at all the feature selection results it can be seen that from the chosen *day features*, that the MSG channels 1, 5, and 8, are selected in almost all of the combinations of echo types and two/three class. It is interesting to note that channel 5, when seen in the relative feature histogram (Figure 3.16), shows a large overlap of the clutter and precipitation classes. Thus, on its own channel 5 would not have a lot of skill at separating clutter and precipitation. However, when using two features, it is the best combination together with channel 8. To illustrate this, a plot of the precipitation and clutter classes in this feature space has been made (Figure 3.19).

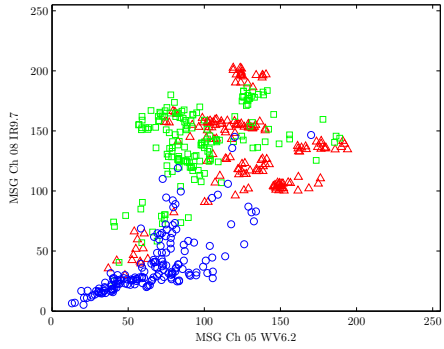

Figure 3.19: In 2D feature space MSG Ch 5 helps in the separation between clutter and precipitation. Feature space for MSG, Ch 5 vs. Ch 8. Red symbols are sea clutter, green symbols are land clutter, and blue symbols are precipitation.

For the chosen *night features*, the chosen feature combinations are dominated by MSG channels 6, 7, 8, 10, and 11. It is noteworthy that channel 8 is selected for both day and night features, and that it is one of the more significant features. As listed in Table 3.1, this is the ozone band of MSG and is described to have no immediate use for precipitation estimation [51]. However, from the feature selection it is shown that the band has use for precipitation and clutter discrimination.

3.9 Classification

Classification is the procedure of labelling observations of features, like the ones described in the previous sections, into two or more classes [30], [64]. Two main groups of classification exist, supervised and unsupervised classification, which differ in that for supervised classification a set of input features is available for which the class membership is known. These training data are used to estimate the parameters of the classifier and evaluate the performance of it. For unsupervised classification, a clustering process must be employed to discover patterns in the data prior to classification.

For this study, supervised classification was chosen because detailed knowledge on the classes and their appearance is available. This was used in Section 3.7 to collect the training data for the classification.

Selection of the optimal classifier for a given classification problem is not trivial. Careful study of the input features and their histograms might hint at an appropriate model. Another approach is to make use of *ensemble methods*, also known as multiple classifier systems, where a range of classifiers are applied in parallel and their results combined. Ensemble methods [46] have the potential to outperform the single classifiers of the ensemble. This happens if each classifier makes mistakes in different domains of the input feature space and when combined, the mistakes "cancel out". Another advantage of using an ensemble classifier is that the performance of each classifier can be accessed and compared at the same time.

In this work, five different classifiers were used and combined using a majority vote between these to form the ensemble classifier. The five classifiers used to classify the data were: linear and quadratic discriminant analysis, a decision tree, k-nearest neighbor classification, and a support vector machine. The criterion for the choice of classifiers was that the classifiers should be from different families of classifiers, e.g., linear discriminant functions are parametric classifiers which assume Gaussian distributions of the classes, while the k-nearest neighbor method is non-parametric.

Linear and quadratic discriminant analysis

Linear and quadratic discriminant analysis for classification involves estimating the parameters of a discriminant function which is subsequently used to map each input vector to a class label [17], [30].

Linear discriminant functions divide the feature space by linear decision boundaries while quadratic discriminant functions use quadratic decision boundaries.

This is illustrated in Figure 3.20, where a subset of 200 training samples (100 clutter samples (red) and 100 precipitation samples (blue)) are plotted for MSG Channels 1 and 4, and Channels 4 and 5, pairwise. The decision boundaries are plotted with black lines and the regions belonging to the clutter and precipitation classes are colored in light red and light blue colors, respectively.

An important feature of the linear and quadratic discriminant functions is that the classes are modelled by Gaussian distributions which is not always a valid assumption.

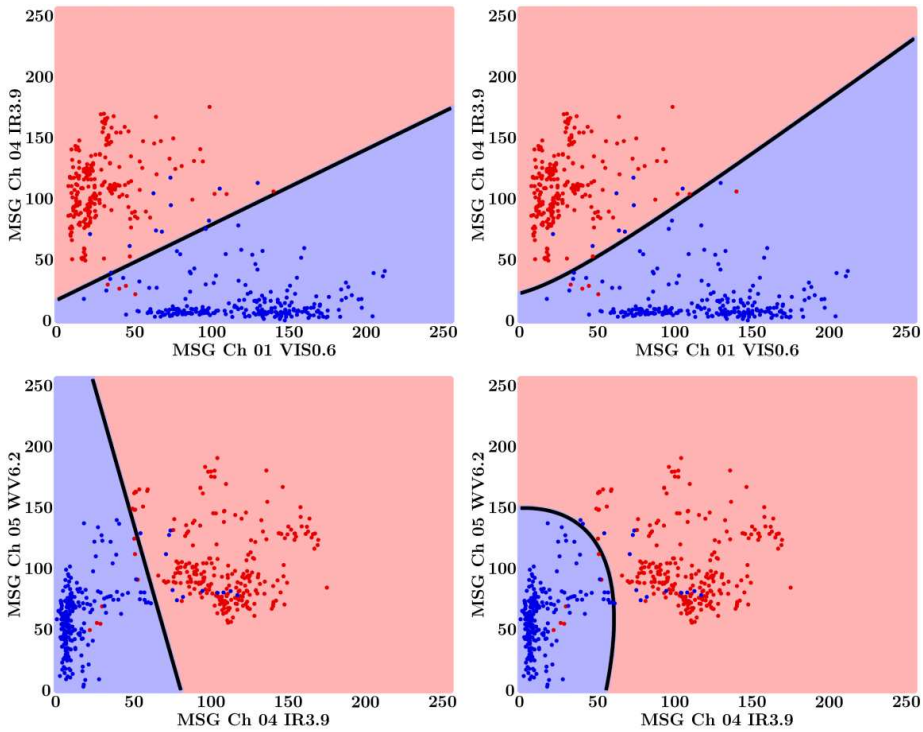

Figure 3.20: Disciminant analysis. Left column: Linear discriminant function, Right column: Quadratic discriminant function.

Decision tree

A decision tree [30], also called a classification tree, is a different approach to supervised classification. The class membership of an input vector is found by following a path through a tree structure based on thresholds on the values of the input features. An example of a decision tree is shown in Figure 3.21, where the first node partitions the features space on the threshold of 48.8082 on MSG

Channel 4. This way an input feature is assigned to its final class membership at the bottom of the tree. A decision tree thus divides the features space using planes which are perpendicular to the axes of the feature space. Classification trees are often used in remote sensing image classification because of the way the thresholds on the input features can be used to efficiently convey any expert knowledge on the domain of interest.

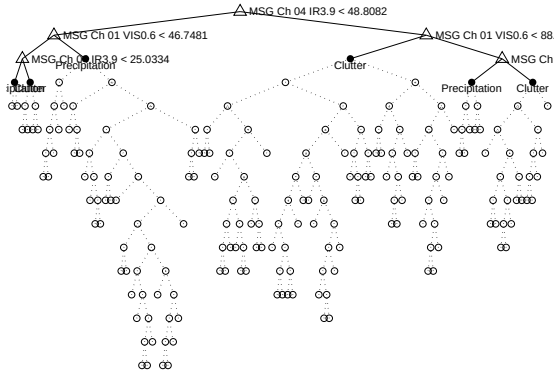

Figure 3.21: Decision tree example. Only the values of the first three tree nodes are shown to in the attempt to increase the clarity of the figure.

K-nearest neighbor

The k-nearest neighbor (KNN) classifier [17] is non-parametric method for classification. As its name explains, it classifies input vectors to the class of the k nearest training vectors. This is illustrated in Figure 3.22 where classification is carried out using values of k of 1 (left) and 11 (right). As it can be seen the value k controls how smooth the class regions are and how well the classifier can capture small-scale structures in the feature space. However, low values of k makes the method more sensitive to noisy data and outliers in the training data. Still, for further analysis a low value of $k = 1$ was chosen. Several higher values of k were applied without much change in the classification accuracy but at an increase of computation time.

Support vector machine

The last applied classifier, was a support vector machine (SVM) classifier. A support vector machine divides the feature space by linear hyperplanes through a process of maximizing the distance from the hyperplane to a subset of the training points called support vectors. This is illustrated in Figure 3.23.

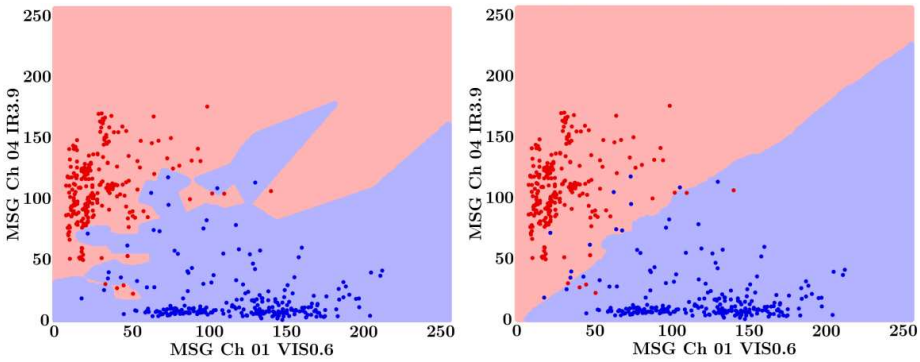

Figure 3.22: K-nearest neighbor classification. $k = 1$ is shown on the left, and $k = 11$ on the right. The color codes and features are the same as in Figure 3.20.

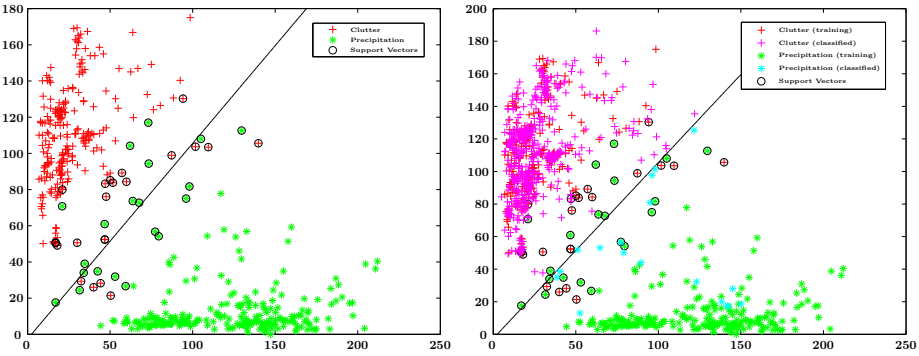

Figure 3.23: Support vector machine. MSG Ch 01 vs. MSG Ch 04.

Ensemble classifier

Following the classification using each of the five classifiers described above the classification results are combined to yield the ensemble classifier result. Various rules exists for the combination [46]. Here a majority vote rule was chosen: the class selected by the majority of the classifiers is chosen as the final output class. Since an odd number of classifiers were used, no ties needed no be broken.

3.9.1 Precipitating Clouds classification

The classification method used for the 'Precipitating Clouds' image product follows a different procedure than for the MSG data since the 'Precipitating Clouds' product contains only one variable. An optimal threshold was found by incremental adjustment of a threshold on the 'Precipitating Clouds' value and maximizing the classification accuracy (See Figure 3.24).

Different optimal thresholds were found for day and night parts of the algorithm. For the day algorithm a threshold on $\leq$0.0 % was found optimal for clutter detection of all, weak, and strong echoes. For the night algorithm, however, the threshold takes the values $\leq$28 %, $\leq$18 %, $\leq$38 % for all, weak, and strong echoes respectively. The results are shown in Tables 3.9 and 3.10.

As it can be seen from those figures, the probability of precipitation as contained in the 'Precipitating Clouds' product is not calibrated to divide at 50 % for distinguishing precipitating and non-precipitating events. [59] mentions a value of 20 % or 30 %, which corresponds well with the results for the *night features* achieved above (thresholds between 18 % and 38 %), but not very well with the results from the *day features* (optimal threshold at 0 %).

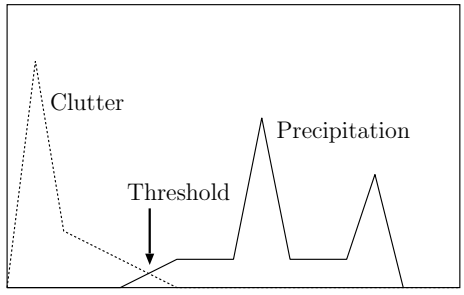

Figure 3.24: The optimal threshold for 'Precipitating Clouds' method is found by incremental search.

3.10 Model training

After training the classifiers, the performance of the classifiers are computed using an independent *test set* taken from the original training samples. If the data used for training the classifier was used also to access the performance of the classifier, unrealistic high performance measures would be obtained, because

the classifiers are fitted to those data.

In many applications of supervised classification, limited amounts of training data are available, however, in this study this was not the case. Often, when enough training data is available, the procedure is simple to split training data into two equal parts, one for training and one for testing. Nevertheless, for this work, it was found a study worth to investigate how the number of training samples affects the classification accuracy.

Varying amounts of training data was used (from 30 to 4000 samples per class) and the data were classified using the methods described above. For each number of training samples the classification was carried out 11 times, which makes it possible to study the standard deviation in the classification as results of the number of training samples. The test set was 50,000 training samples in total, i.e., not the full remaining training set.

When the classification accuracy is plotted as a function of the number of training samples, *learning curves* are produced. A couple of examples of these are shown in Figures 3.25 and Tables 3.3 and Figure 3.26 and Table 3.4. Shown in the table are the mean and standard deviation of the classification accuracy for each of the five classifiers and their majority vote combination. In the graph, the thick lines are the mean classification accuracy and the thinner lines indicate one standard deviation from the mean.

Learning curves were produced for all combinations of day and night features, and in general, all learning curves flattened out at between 1000 and 4000 training samples per class, as illustrated on the two figures. Also noticed was that the standard deviation of the classifiers decrease to very low values (around 0.1 %) for 4000 samples. Thus this was chosen to use this number of training samples per class in the further analysis.

Regarding the individual classifiers, k-nearest neighbor and the decision tree methods were seen to achieve highest accuracies, followed by the support vector machine and the discriminant analysis.

Table 3.3: Learning curve results. Data fusion method, day echoes and day features, strong/weak echoes, three classes.

#	LDA		QDA		Dtree		KNN		SVM		Majority	
samples	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std
30	96.6	0.88	95.6	1.87	92.0	2.97	95.9	1.38	96.4	0.78	96.8	0.64
50	96.8	0.61	96.8	0.68	92.8	4.15	96.8	0.67	96.8	0.68	97.2	0.36
100	97.4	0.16	97.5	0.46	96.5	0.82	97.0	0.89	96.8	0.55	97.6	0.05
500	97.3	0.10	98.1	0.19	97.7	0.33	98.8	0.34	97.7	0.11	98.2	0.20
1000	97.4	0.06	98.1	0.12	97.9	0.52	99.3	0.10	97.7	0.16	98.3	0.10

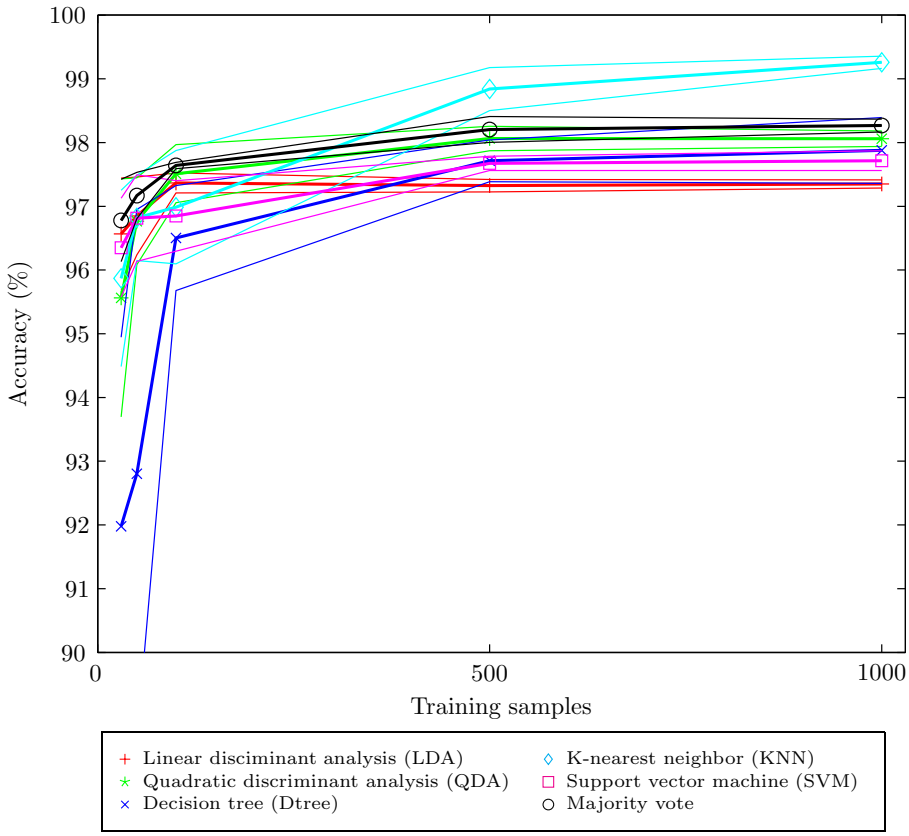

Figure 3.25: Learning curves. Data fusion method, day echoes and day features, strong/weak echoes, three classes.

Table 3.4: Learning curve results. Data fusion method, day echoes and day features, strong/weak echoes, two classes.

#	LDA		QDA		Dtree		KNN		SVM		Majority	
samples	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std
30	95.6	0.92	94.3	1.88	91.9	4.63	94.2	1.24	95.6	0.81	95.1	0.74
50	96.1	0.43	95.9	0.74	94.2	1.96	94.0	1.85	95.4	1.08	95.9	0.63
100	96.5	0.14	96.4	0.47	95.2	0.97	95.9	0.70	96.4	0.18	96.5	0.17
500	96.5	0.16	96.9	0.09	96.1	0.63	97.5	0.16	96.3	0.26	96.9	0.13
1000	96.5	0.19	97.0	0.09	96.8	0.44	98.1	0.22	96.4	0.20	97.2	0.22
2000	96.5	0.13	97.0	0.08	97.4	0.36	98.7	0.18	96.4	0.12	97.2	0.12
3000	96.5	0.08	97.0	0.04	97.9	0.12	98.8	0.11	96.4	0.08	97.2	0.10
4000	96.5	0.05	97.0	0.05	98.0	0.15	99.1	0.08	96.4	0.11	97.2	0.07

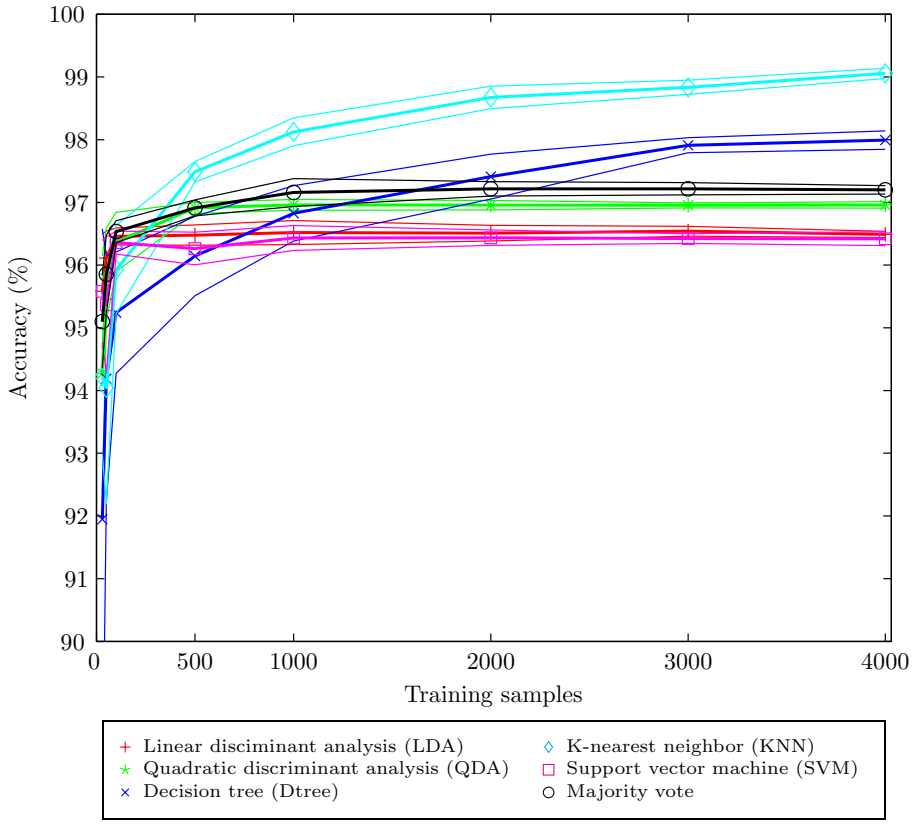

Figure 3.26: Learning curves. Data fusion method, day echoes and day features, strong/weak echoes, two classes.

3.11 Classifier evaluation

Upon building the classifiers their performance need to be accessed using independent test data to be able to know how well the classifiers are expected to perform when they are applied to new data.

3.11.1 Skill scores

Various skill scores [73] were used, namely the percent correct (PC), the false alarm ratio (FAR), and the Hanssen-Kuipers skill score (HKS). The skill scores are computed from the elements of the confusion matrix of the classification, as shown in Figure 3.27. The confusion matrix contains the number of correctly classified samples (hits and correct rejections) and the number of incorrectly classified samples (misses and false alarms). In this study of detection of clutter, clutter is the "target" of the classification, and hence clutter classified as clutter is counted as *hits*, whereas precipitation misclassified as clutter is a *false alarm*. Precipitation correctly classified as precipitation is a *correct rejection* and finally clutter misclassified as precipitation is a *miss*. Confusion matrices for the cases of three class classification were made by pooling sea and land clutter classes into a common clutter class.

		Classification	
		Clutter	**Precipitation**
Truth	**Clutter**	**Hit (A)**	**Miss (C)**
	Precipitation	**False alarm (B)**	**Correct rejection (D)**

$PC = \frac{A+D}{N} \times 100$	$FAR = \frac{B}{A+B}$	$HKS = \frac{A}{A+C} - \frac{B}{B+D}$
$POD = \frac{A}{A+C}$	$POFD = \frac{B}{B+D}$	$N = A + B + C + D$

Figure 3.27: Confusion matrix for two-class discrimination between clutter and precipitation (top). Skill scores: Percent Correct (PC), False Alarm Ratio (FAR), Hanssen-Kuipers Skill Score (HKS), Probability of Detection (POD), Probability of False Detection (PODF), and total number of observations (N).

Percent correct (PC), is similar to classification accuracy, and is the ratio between the sum of the hits and correct rejections and the total number of observations. The perfect classifier would have only hits and correct rejections

and no misses or false alarms. This would yield a Percent Correct score of 100 %. PC ranges between zero (worst) and 100 % (best).

False alarm ratio (FAR) is the ratio between the number of false alarms and the sum of hits and false alarms. It ranges between 0 (best) and 1 (worst) and in this study it is a measure of how much precipitation is misclassified out of the total number of samples classified as clutter.

Probability of detection (POD) is the ratio between the number of hits and the sum of hits and misses, i.e., how much clutter was classified correctly. Best score is 1 and worst score is 0.

Probability of false detection (POFD), is the ratio between number of false alarms and the sum of false alarms and correct rejections, i.e., how much precipitation was misclassified. Best score is 0, worst score is 1.

Hanssen-Kuipers skill score (HKS) is computed as the difference between the POD and the POFD. HKS thus takes values between -1 (worst) and 1 (best). HKS takes the value 1 when POD is 1 and POFD is 0, i.e., the perfect classification. A HKS value of 0 is obtained when the classification has no skill, i.e., it is no better than the random choice. HKS value smaller than 0 indicates worse skill than the random choice, and a value of -1 is obtained when the samples are consequently classified into their opposite class.

3.11.2 Skill score results

In this section the data fusion method is evaluated on the full set of training samples. The features chosen through feature selection were used to perform the ensemble classification. 4000 training samples per class were used for building the classifiers as chosen earlier as a sufficient number. The remaining samples from the complete training data were subsequently classified as test samples. The exact number of test samples can be computed from the total number of training samples listed in Figure 3.14 minus 4000 times the number of classes.

Upon classification of the test set, the three skill scores, Precent Correct (PC), False Alarm Ratio (FAR), and the Hanssen-Kuipers Skill (HKS) were computed. The results are shown for MSG data in Tables 3.5 to 3.8 and for the 'Precipitating Clouds' product in Tables 3.9 and 3.10. The scores are reported for each individual classifier (LDA and QDA, are the linear and quadratic discriminant analysis, Dtree is the decision tree, KNN is the k-nearest neighbor, and SVM is the support vector machine).

MSG classification

Firstly, for percent correct it is seen (in Table 3.5) that for day features, all echoes, the majority vote obtains a PC value of 95.64 %, i.e., 4.36 % of the test samples were misclassified. The value is better than the worst of the individual classifiers (The QDA with the score of 94.27 %). Looking at all the PC results (Tables 3.5 to 3.8) it is seen that the majority vote consistently outperforms the single worst classifiers. This is in accordance with the purpose of using an ensemble classifier. Improvements by using the majority classifier in comparison with the worst single classifier ranges from 0.78 % up to 12.1 %. The best one improves the classification accuracy from 85.50 to 97.58 % (LDA, night features, strong echoes, two classes, Table 3.7).

The best majority vote PC result is obtained for night features, strong echoes, and three classes with a score of 98.39 %. The worst score is 93.59 % for day features, weak echoes, two classes. For all echoes the best result is 95.82 % for day features and three classes, however, the other results differ by small figures only.

Evaluating the performance of the individual classifiers, it can be seen that the k-nearest neighbor (KNN) classifier almost consistently outperforms the other classifiers by approximately 1 %. The worst classifier is, also almost consistently, the linear discriminant analysis (LDA).

Looking at the false alarm ratio (FAR) for the majority vote classifier, values between a low 0.007 and 0.073 are seen. The best is obtained for night features and strong echoes, two classes, and the worst for day features, weak echoes, three classes. The FAR values indicate that between 0.7 and 7.3 % of the detected clutter were false alarms. For all echoes, night features, two classes shows the best performance (FAR of 0.013, Table 3.7). For day features the the best FAR value was 0.054 using two classes.

The Hanssen-Kuipers skill score which, as described earlier, takes into account both the probability of detection and the probability of false detection, scores from 0.870 (day features, weak echoes, three classes) to 0.948 for day features, strong echoes, three classes and a similar values for night features, strong echoes, two classes. The single best all echoes classifier for day features, is three class classification with a HKS of 0.919 and for night features 0.908 for two classes.

Echo types

Looking at the classification performance for the echo type categories, weak, and strong echoes, it can be seen that strong radar echoes are classified consistently better than weak echoes when considering all three skill scores. The classification

accuracy is between 1.48 % and 3.35 % better for strong echoes compared to weak echoes.

Number of classes

The effect of using two or three classes can also be assessed from the table. Looking at the results, it can be seen that a slightly better PC score is obtained for the majority vote classifier using three classes for day features (95.64 % to 95.82 %, 93.59 % to 93.93 % and 96.94 % to 97.05 % for all, weak and strong echoes respectively). The opposite effect is observed (i.e., lowered PC) for night features when using three classes.

For the FAR and HKS scores the same pattern is observed, however, the difference of using two or three classes in the classification is in general very small considering also that the standard deviation for the ensemble classifier was estimated to be around 0.1 % in Section 3.10.

Precipitating Clouds classification

In Tables 3.9 and 3.10, the performance of the 'Precipitating Clouds' classification method can be seen as applied to *day features* and *night features* respectively. Classification accuracies from 92.92 % to 96.95 % are observed in the PC scores with the best classification being strong echoes in the day and the worst being weak echoes also in the day. False alarm ratios from 0.030 to 0.078 are obtained for weak echoes in night and day, respectively. The Hanssen-Kuipers skill score highlights the day features for strong echoes as the best (0.939) and the strong echoes at night as the worst (0.701).

Table 3.5: Classification results. Data fusion method, day echoes and day features, two classes.

Echo	Percent correct (PC)						False alarm ratio (FAR)						Hanssen-Kuipers skill (HKS)					
type	LDA	QDA	Dtree	KNN	SVM	Majority	LDA	QDA	Dtree	KNN	SVM	Majority	LDA	QDA	Dtree	KNN	SVM	Majority
All	94.94	94.27	96.74	98.37	94.91	95.64	0.061	0.079	0.037	0.018	0.055	0.054	0.899	0.888	0.935	0.967	0.898	0.913
Weak	92.55	92.50	96.70	98.73	92.34	93.59	0.059	0.069	0.025	0.011	0.058	0.053	0.851	0.846	0.934	0.974	0.847	0.871
Strong	96.16	96.35	97.46	98.34	97.07	96.94	0.084	0.076	0.049	0.029	0.057	0.063	0.933	0.934	0.952	0.966	0.945	0.944

Table 3.6: Classification results. Data fusion method, day echoes and day features, three classes.

Echo	Percent correct (PC)						False alarm ratio (FAR)						Hanssen-Kuipers skill (HKS)					
type	LDA	QDA	Dtree	KNN	SVM	Majority	LDA	QDA	Dtree	KNN	SVM	Majority	LDA	QDA	Dtree	KNN	SVM	Majority
All	94.38	94.36	97.40	98.75	95.28	95.82	0.081	0.085	0.036	0.020	0.075	0.067	0.890	0.890	0.949	0.976	0.909	0.919
Weak	91.88	92.17	96.44	98.53	93.09	93.93	0.086	0.083	0.042	0.018	0.082	0.073	0.830	0.836	0.924	0.968	0.852	0.870
Strong	95.79	96.11	97.46	98.46	97.05	97.05	0.096	0.087	0.059	0.036	0.060	0.067	0.928	0.932	0.956	0.973	0.945	0.948

Table 3.7: Classification results. Data fusion method, night echoes and night features, two classes.

Echo	Percent correct (PC)						False alarm ratio (FAR)						Hanssen-Kuipers skill (HKS)					
type	LDA	QDA	Dtree	KNN	SVM	Majority	LDA	QDA	Dtree	KNN	SVM	Majority	LDA	QDA	Dtree	KNN	SVM	Majority
All	89.03	94.31	96.90	98.12	91.57	95.28	0.013	0.027	0.010	0.004	0.025	0.013	0.830	0.861	0.935	0.965	0.834	0.908
Weak	91.86	94.82	97.66	98.83	93.58	96.10	0.009	0.022	0.007	0.002	0.023	0.009	0.873	0.877	0.953	0.979	0.860	0.926
Strong	85.50	96.84	98.28	99.35	91.63	97.58	0.009	0.021	0.005	0.002	0.012	0.007	0.796	0.898	0.964	0.987	0.862	0.948

Table 3.8: Classification results. Data fusion method, night echoes and night features, three classes.

Echo	Percent correct (PC)						False alarm ratio (FAR)						Hanssen-Kuipers skill (HKS)					
type	LDA	QDA	Dtree	KNN	SVM	Majority	LDA	QDA	Dtree	KNN	SVM	Majority	LDA	QDA	Dtree	KNN	SVM	Majority
All	90.71	93.45	96.19	96.99	93.50	94.28	0.023	0.029	0.014	0.009	0.025	0.022	0.829	0.848	0.918	0.939	0.859	0.874
Weak	93.42	94.78	97.19	97.74	94.61	95.49	0.022	0.029	0.010	0.007	0.024	0.021	0.863	0.862	0.939	0.954	0.872	0.891
Strong	87.51	98.21	98.97	99.82	93.44	98.39	0.019	0.019	0.003	0.001	0.017	0.015	0.797	0.926	0.979	0.995	0.872	0.940

Table 3.9: Classification results and classification thresholds. Data fusion, nowcasting product method, day echoes.

Echo type	Threshold (Pct.)	Percent correct (PC)	False alarm ratio (FAR)	Hanssen-Kuipers skill (HKS)
All	0.0	94.81	0.069	0.898
Weak	0.0	92.92	0.078	0.849
Strong	0.0	96.95	0.053	0.939

Table 3.10: Classification results and classification thresholds. Data fusion, nowcasting product method, night echoes.

Echo type	Threshold (Pct.)	Percent correct (PC)	False alarm ratio (FAR)	Hanssen-Kuipers skill (HKS)
All	28.0	93.46	0.044	0.790
Weak	18.0	94.88	0.030	0.837
Strong	38.0	93.28	0.040	0.701

3.12 Reject option

Misclassification in the detection of clutter is, of course, always unwanted. However, for some applications it might be desirable to remove a large proportion of the clutter at the expense of removing small amounts of precipitation at the same time (or vice versa). For example, for qualitative use of the data, in its general use in operational weather forecasting or the display of the data for non-scientific users, this might be the case. The confusion caused by large amounts of clutter might be greater than the confusion from small amounts of missed precipitation. On the contrary, if the radar data is used quantitatively in a model to issue warnings of extreme thunderstorms then removing any precipitation might be unwanted and instead more unclassified clutter can be tolerated (it might not matter if a couple of false thunderstorm warnings are issued).

Using a *reject option* it is possible to adjust the expected amount of clutter classified as clutter (hits) at the expense of how much precipitation is mistakenly classified as clutter (false alarms).

The reject option makes it possible to avoid making classifications in the case of uncertainty on the class membership of an input vector. This is done by applying a threshold on the posterior probability of the classification. In Figure 3.28 this is illustrated: A threshold θ is used to reject input vectors inside the reject region. If θ is 1 then all inputs are rejected (i.e., nothing is classified as anything but precipitation) and if θ is 0.5 (in general 1 divided by the number of classes) then no inputs are rejected.

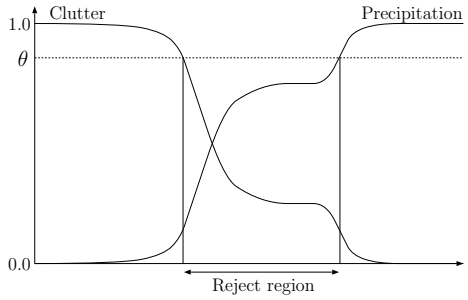

Figure 3.28: The reject option. Inputs to the classifier are rejected if it is in the reject region, i.e., if the largest of the posterior probabilities is smaller or equal to the threshold θ. After [17].

To illustrate how the reject option can be applied, reject thresholds from 0.5 to 0.99 were applied to the posterior probabilities from a quadratic discriminant classifier used to classify MSG channels 1, 5, 6, 9 10 (The day features, all echoes, two classes case).

If the largest of the two posterior probabilities for the classes, clutter and precipitation, was smaller than the threshold, then the input vector was assigned to the precipitation class. The POD and POFD scores were computed for each reject threshold and a receiver operating characteristic (ROC) curve was plotted (see Figure 3.28). On the plot the reject threshold is plotted along the ROC curve. It is seen how an increasing reject threshold lowers the probability of false detection, i.e., fewer false alarms are made (less precipitation misclassified as clutter). This is at the trade-off of a lower probability of detection of course (more misses, i.e., more clutter misclassified as precipitation). A good choice for many application could be the operating point of the ROC curve (where the tangent to the curve is 45 degrees).

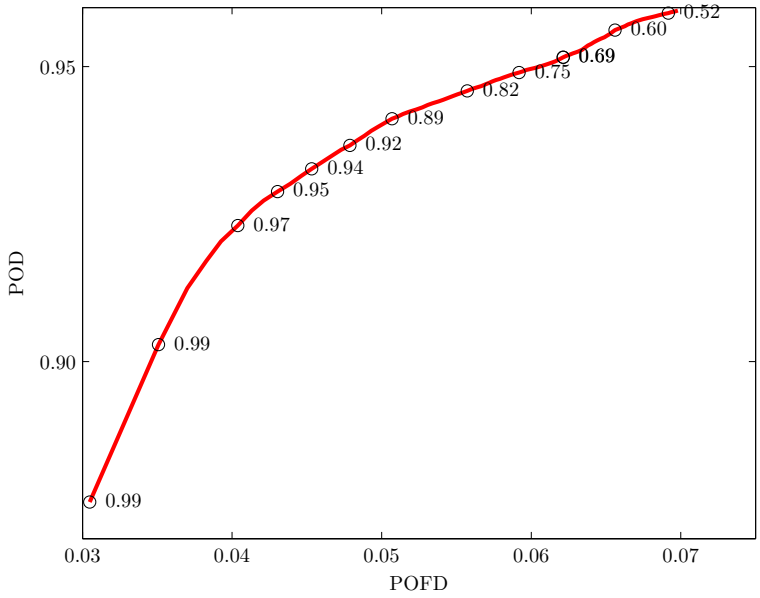

Figure 3.29: ROC curve for reject option. Day features, all classes, two class classification. Adjusting the reject threshold controls the trade-off between POD and POFD.

3.13 Results

After the design and evaluation of the classification model as described in the earlier sections, the classifier is ready to be applied to new data in operational applications. To illustrate this, the three case examples (presented in Section 3.3) were classified using the classification models developed and evaluated above. The results are shown in for the MSG classification method in Figure 3.30 and for the 'Precipitating Clouds' method in Figure 3.31.

For *Case I*, the case with almost only clutter echoes, the result for the MSG classification (Figure 3.30, first column) is very good with a classification accuracy of 99.3 % (The accuracy is shown in the bottom of the figure and was computed from the confusion matrix shown beneath it). Both the sea and land clutter is classified correctly which can be seen by comparison of the classified image in the middle row with the training areas (shown in Figure 3.12). The small error of 0.7 % originated from the small amount of precipitation in the lower North Sea of which two thirds were misclassified. In the lower row of the

figure, a “cleaned” radar image is seen. By this is meant that the echoes which were classified as clutter have been removed.

Now turning to the 'Precipitating Clouds' method for *Case I* (Figure 3.31, first column), the result is seen to be good also. The achieved classification accuracy is 97.1 %. The precipitation in the North Sea was not misclassified to such a great extent, but on the contrary a large patch of land clutter on the coast of Norway was misclassified as precipitation.

Classification results for *Case II* is shown in the middle column of the figures. For the MSG method 95.9 % of the echoes were classified correctly. The misclassifications are equally divided between missed clutter and false alarms (approximately 1000 echoes each). The 'Precipitating Clouds' method shows similar results with a accuracy of 95.0 % but a higher proportion of false alarms. The false alarms are seen to be located on the Southwestern side of the precipitation area in the middle of the image for both methods. The missed clutter is the land clutter close to the radar on Sindal. The location of the misclassification indicates that it is caused by misalignment. Either due to a slight temporal misalignment between the radar and satellite data (The precipitation was moving Northeast for that case example). The location of the misclassification could also hint at a contribution from the parallax displacement (which results in cloud and precipitating to be offset *away* from the nadir of the satellite to the north).

Case III is the all precipitation case and a high classification accuracy of 98.5 % and was achieved for the MSG method and a lower accuracy of 95.1 % for the 'Precipitating Clouds' method. Again, the false alarms are located on the Southwestern areas of the precipitation.

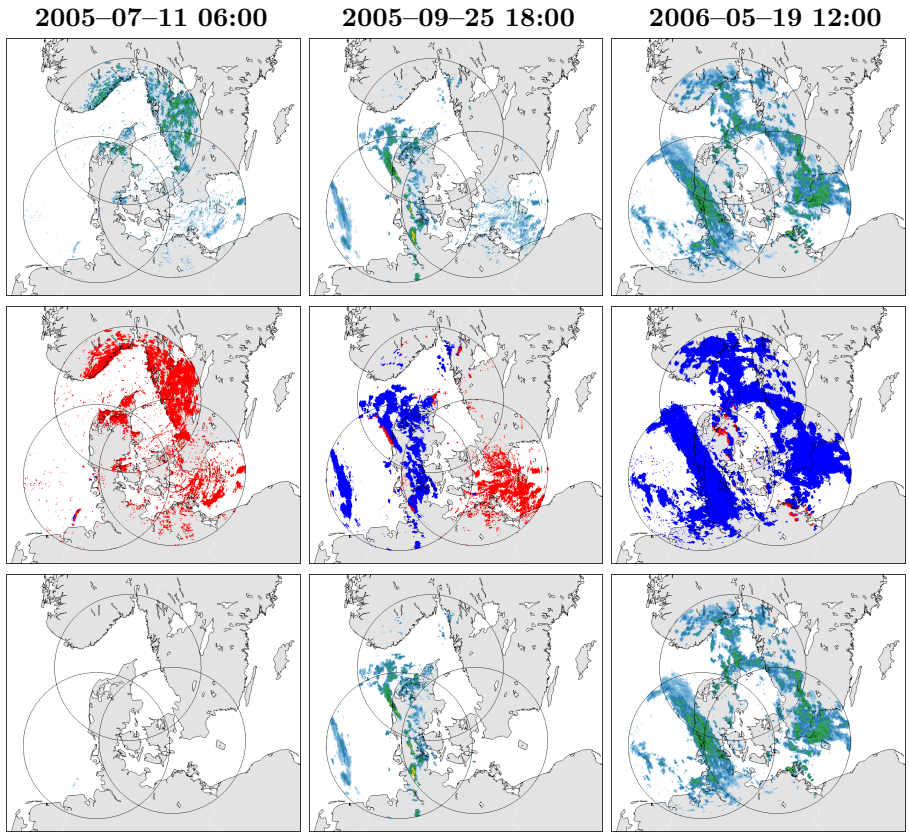

99.3 %

	C_c	P_c
C_t	39703	0
P_t	269	108

95.9 %

	C_c	P_c
C_t	15705	1065
P_t	1007	32299

98.5 %

	C_c	P_c
C_t	0	0
P_t	2074	134992

Figure 3.30: Classification results using the MSG data fusion method on the three case examples. Top row: Radar images, Middle row: Classified images (blue is precipitation and red is clutter), Bottom row: Filtered radar images. The classifications accuracy and confusion matrix are found below (Subscript t means *truth* and subscript c means *classification*, cf. Figure 3.27)

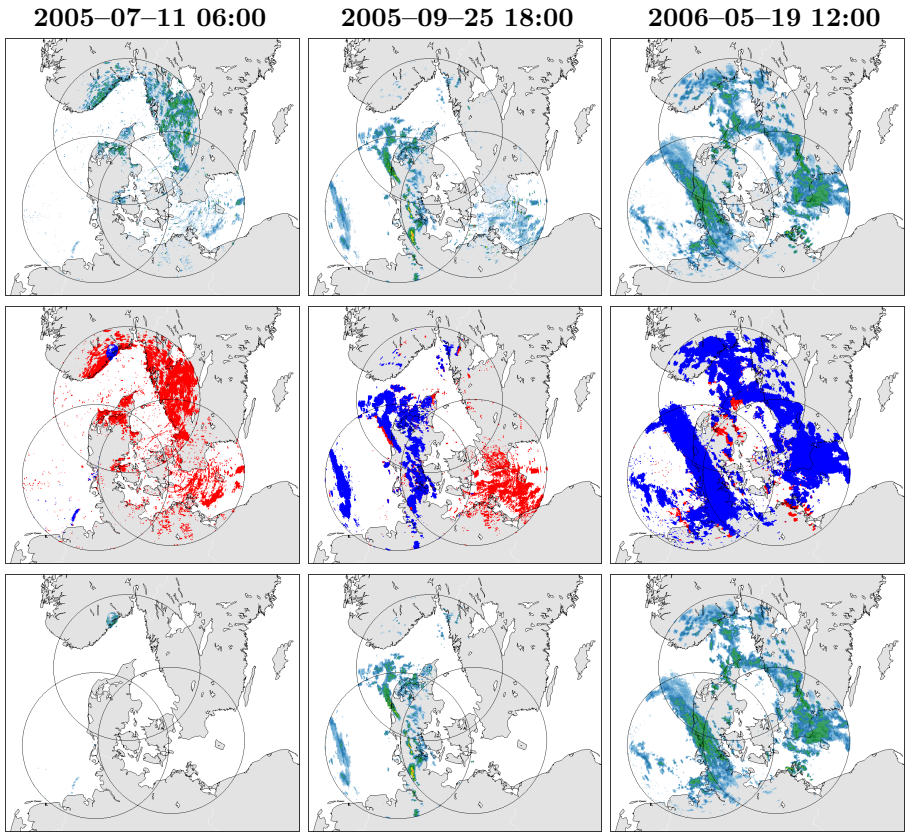

97.1 %

	C_c	P_c
C_t	38583	1120
P_t	23	354

95.0 %

	C_c	P_c
C_t	15747	1023
P_t	1460	31846

95.1 %

	C_c	P_c
C_t	0	0
P_t	6708	130358

Figure 3.31: Classification results using the 'Precipitating Clouds' data fusion method on the three case examples. Top row: Radar images, Middle row: Classified images (blue is precipitation and red is clutter), Bottom row: Filtered radar images. The classifications accuracy and confusion matrix are found below (Subscript t means *truth* and subscript c means *classification*, cf. Figure 3.27)

3.14 Summary and conclusion

In this chapter, a method using information from multispectral satellite images to detect clutter in weather radar data was developed, evaluated, and applied.

The data used as input to the method was 1) the 11 bands of from the SEVIRI sensor on board the Meteosat Second Generation platform and 2) an operational nowcasting product called 'Precipitating Clouds' which was derived from the 11 MSG bands. A different classification procedure was used for each of these.

Training data for the model was extracted from 60 pairs of co-located radar and satellite data sets from 15 days of varying meteorological conditions.

For the MSG bands, a procedure of feature selection was used to select a subset from the 11 bands which provided a high classification accuracy and at the same time was comprised of as few features as possible. The method was split into a day and a night part due to the lack of visual features at night. For the day features, MSG channels 1, 5, and 8, were consistently chosen as most significant in achieving a good classification accuracy. For the night features the channels 6, 7, 8, 10, and 11 were chosen.

Supervised classification was performed using five different classifiers. Linear and quadratic discriminant analysis, a decision tree, k-nearest neighbor classification, and a support vector machine. These were combined into an ensemble classifier using a majority vote between the results of each individual classifier. The amount of training data needed for building the classifiers was estimated from the construction of learning curves and 4000 samples per class were seen to be sufficient to achieve maximum classification accuracy.

The classification system was evaluated by training the models using 4000 samples per class and classifying the remaining of the training data as test data. The performance of the method was accessed using the skill scores percent correct (PC), false alarm ratio (FAR), and the Hanssen-Kuipers skill score (HKS).

The classification accuracy for day features and all echoes (weak and strong echoes combined) was estimated to be 95.82 % using the majority vote and three classes (land and sea clutter and precipitation). For night features the classification accuracy was 95.28 % using (two classes). The best single classifier on day features was the k-nearest neighbor classifier which scored 98.75 %. For night features and all echoes, a classification accuracy of 95.28 % was obtained for the majority vote and the single best was the k-nearest neighbor again with 98.12 % accuracy. The worst classifier was the linear discriminant function which performed up to 12.1 % worse than the majority vote classifier.

The use of an ensemble classifier was shown to improve the classification performance. The classification accuracy was improved by 0.78 % up to 12.1 % in comparison to the single worst classifier, however, using the ensemble classifier the classification accuracy was at the same time lowered in comparison to the single best classifier. Since it is unknown before classification which classifier will perform the best, an ensemble classifier can be an advantage.

For the 'Precipitating Clouds' method, a classification accuracy of 94.81 % was achieved for the day algorithm and 93.46 % for the night algorithm. This is 1.0 % and 1.8 % lower accuracy than for the MSG method. This difference might be caused by the use of the untuned version of the 'Precipitation Clouds' data set. If the 'Precipitation Clouds' product had been tuned to radar data, a higher performance would be expected.

No large difference in classification accuracy was observed from using two or three classes, i.e., to split up the clutter class into separate land and sea clutter classes, however, better clutter detection was obtained for strong precipitation echoes than for weak echoes. This indicates a higher correlation between strong precipitation and the multispectral features of the satellite data.

Finally, the method was applied to three case examples. Classification accuracies between 95.0 % and 98.5 % were obtained. Widespread clutter and precipitation was classified well, however, misclassification occurs on the edges of precipitation areas due to misalignment issues related to the use of multiple data sources from multiple sensors.

In the publications of the preliminary results of these methods, classification accuracies between 92.1 % and 95.5 % were seen for the MSG method [19] and 91.9 % and 100 % for the 'Precipitating Clouds' method [18]. These results were achieved on a limited size dataset, and it was concluded that the methods should be evaluated on longer time series of data, since it is expected that high performance of a method can be achieved easily using one or two images, but as the gamut of the feature spaces expands the performance decreases. Thus, in this study, a large dataset was used for training of the classifiers and for the evaluation of the classification. It can be concluded that the methods perform very well on a larger dataset as well and that the accuracies obtained (around 95 %) encourages the use of the methods in operational applications.

Chapter 4

Spatio-temporal method

4.1 Introduction

In qualitative operational use of weather radar images, a common way to visually distinguish ground clutter from precipitation is to inspect the evolution of radar echoes in a series of consecutive radar images. Typically, clutter echoes are static or display a random motion pattern (static clutter often originates from non-moving targets like mountains and buildings and less static clutter often originates from the sea surface). Precipitation, on the other hand, typically displays a more homogeneous motion pattern determined by the atmospheric wind field and other meteorological processes like precipitation growth and decay.

This is illustrated in Figure 4.1, where radar images at intervals of 10 and 60 minutes show how precipitation echoes move steadily and land clutter echoes remain in the same position.

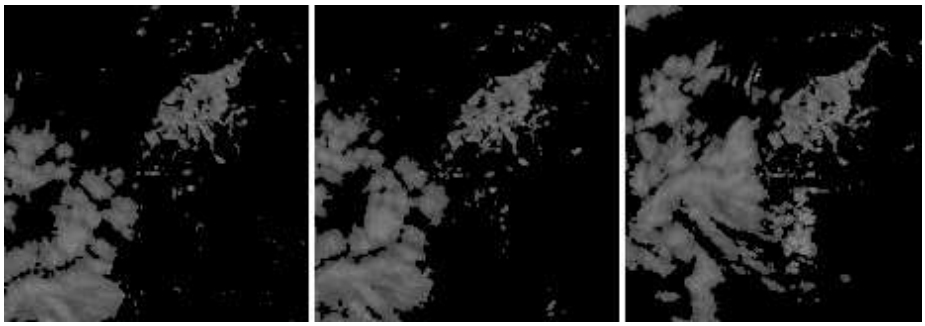

Figure 4.1: Radar images at times T (left), T+10 minutes (middle), and T+60 minutes (right). A precipitation area is moving in from the Southwest towards a stationary land clutter area. See also Figure 4.3.

The method proposed in this chapter is a heuristic, spatio-temporal method making use of the characteristics described above for classification of radar echoes into precipitation and clutter classes.

Firstly, the motion field between two consecutive radar images is identified using an optical flow technique. From the motion field, features are extracted which enable the discrimination of areas of homogeneous motion and areas of random (or no) motion. To make the classification more robust, additional external information on the motion of the precipitation and clouds is taken from a numerical weather prediction model.

The idea behind the method is similar to the methods that use the fact that the fluctuations and decorrelation of radar echoes is faster for clutter than for precipitation. The application of optical flow for clutter detection, however, is not widespread. One method [57], used the root-mean-square between five consecutive radar images to highlight areas of changes in motion and to detect areas of clutter and precipitation. The Doppler velocity of the radar echoes was used as an additional feature.

4.2 Data

Radar data

As input to the method the same radar images as used in the previous chapter were extracted from the database. The next consecutive images 10 minutes

after, i.e., at 01:10, 06:10, 12:10, and 18:10 UTC were also extracted to be able to detect the motion field.

Wind vectors from NWP

Prior information on the motion of precipitation areas is available from numerical weather models. HIRLAM (HIRLAM is short for: HIgh Resolution Limited Area Model) is an operational numerical weather prediction system used at several European meteorological institutes. HIRLAM provides predictions of a wide range of meteorological parameters (temperature, air pressure, humidity, wind, etc.) on a grid consisting of 40 vertical levels (up to 40 km into the atmosphere) and a horizontal grid of 0.05 ° (approximately 5 km). HIRLAM is described more in detail in Section 5.3.2 in Chapter 5. From HIRLAM, the wind vectors were extracted for the same times as the radar image pairs.

The atmospheric wind varies both in the vertical and horizontal direction. To relate the NWP wind predictions to the motion of clouds and precipitation the average wind velocity between 0.5 km and 3 km height (where the observed precipitation occurs) was computed from the 3D grid of wind vectors. Of course the motion of precipitation in weather radar images is not only determined by the atmospheric wind but also by other atmospheric processes related to cloud dynamics. The wind field is, however, probably the most significant contributor to the spatio-temporal evolution of radar echoes on short time scales such as the one used in this study.

4.3 Feature extraction

Optical flow

To estimate the motion of radar echoes in between two images, an optical flow technique is used. Several applications of optical flow techniques are found in meteorological image analysis. It is applied, for example, in the determination of atmospheric motion vectors (the wind field) by detection of the optical flow in the water vapor channels of multispectral satellite images [38]. To improve the temporal resolution in sequences of satellite images, optical flow has been used to enable temporal interpolation in-between images [50], and in the field of weather radar it is widely used in nowcasting of precipitation [37], [20].

Several methods for deriving the optical flow [39] in image series exists (cf. reviews in [8] and [50]). The main groups are *differential methods*, *frequency-based methods*, and *region-based matching methods*. The first use the gradients of the

image intensities to compute image motion and the second detects the flow in the frequency domain. The third method is based on detection of the flow direction and speed by matching of regions between images using for example the cross-correlation. Matching-based methods are recommended for cases where a small number of images exists (large motions and undersampled image sequences) in comparison to gradient-based methods [8], and hence a region-based matching method for computation of optical flow was chosen.

The widely used region-based matching method, *Anandan's method* [4], was applied. The method computes the optical flow using a pyramid approach. First, the large scale motions are derived on a coarse scale and these are subsequently propagated to finer scales. The matching was done using a 5-by-5 pixels moving window and the number of hierarchical pyramid levels were set to 4.

Experiments with preprocessing of the images by smoothing were carried out. Smoothing is often applied prior to optical flow computation to remove noise. However, in this study, if too much smoothing was applied then the features which make precise identification of the optical flow were removed as well and the high frequency motion pattern of clutter was dampened. The pre-smoothing provided no improvement to the results and was omitted. The reason why pre-smoothing had no positive effect could also be due to the coarse-to-fine flow detection method which automatically involves some noise reduction on the coarse levels.

From the detected flow in pixel coordinates the directions of the velocity field were converted into azimuth angles and the speeds were converted from pixels per 10 minutes into meters per second.

In Figure 4.2 the feature histograms are shown for the optical flow direction and speed for the clutter and precipitation classes. For the distribution of flow direction it is noticed how for the training data used in this study, the preferred flow direction of precipitation was around 240 degrees. This fits well with the fact that the prevailing wind directions in Denmark is West and Southwest (between 245 and 270 degrees). The clutter classes display nearly uniform feature histograms, as would be expected. The distribution of optical flow speed shows that the land clutter class has predominantly low speeds (peak value around 5 m/s. Of course, a value of 0 m/s would be expected for stationary land clutter, but when watching land clutter in a radar animation some fluctuations are observed, which are detected as low speed motion by the optical flow algorithm. The precipitation and sea clutter classes are quite similar with average speeds around 15 m/s. The overlap between all three classes is quite large and inspection of the conditional probabilities for the classes (Appendix C, Figures C.15 to C.20, top two plots) shows the little potential of using the detected speed alone to remove clutter.

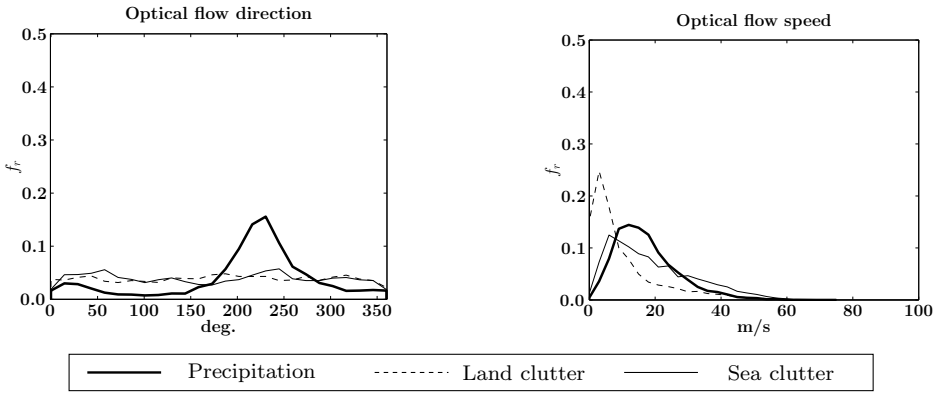

Figure 4.2: The distribution of optical flow direction (left) and speed (right) for clutter and precipitation classes. Total number of samples: 1.140.697.

To show the potential of the spatio-temporal method, an example is presented. For an area of precipitation and land clutter (Figure 4.3) the optical flow was computed (shown in Figure 4.4). It is seen that the detected optical flow of precipitation (in the left side of the image subset) has a homogeneous direction (to the Northeast), but that the clutter area shows a much more chaotic pattern. Comparison of the optical flow field with the predicted wind field from the numerical weather prediction model (Figures 4.5, 4.6, and 4.7) also highlights the difference between clutter and precipitation. The difference in the distribution of directions is clearly seen in rose diagrams of the two classes (Figure 4.8). The difference in speed is also very different (Figures 4.6 and 4.7). The precipitation area shows higher speeds than the clutter, which makes sense since the land clutter originates from stationary targets.

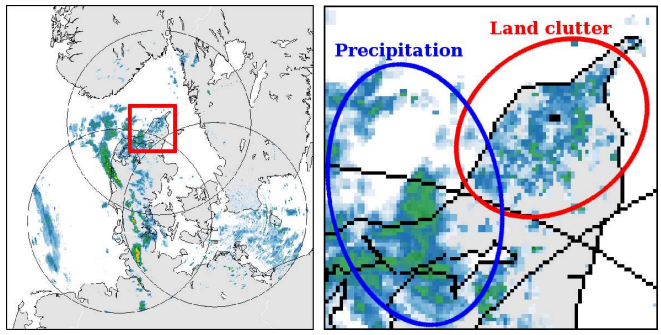

Figure 4.3: Study area for example of potential of optical flow features. 2005–09–25 18:00. The optical flow of the subset and its consecutive image is shown in Figure 4.4.

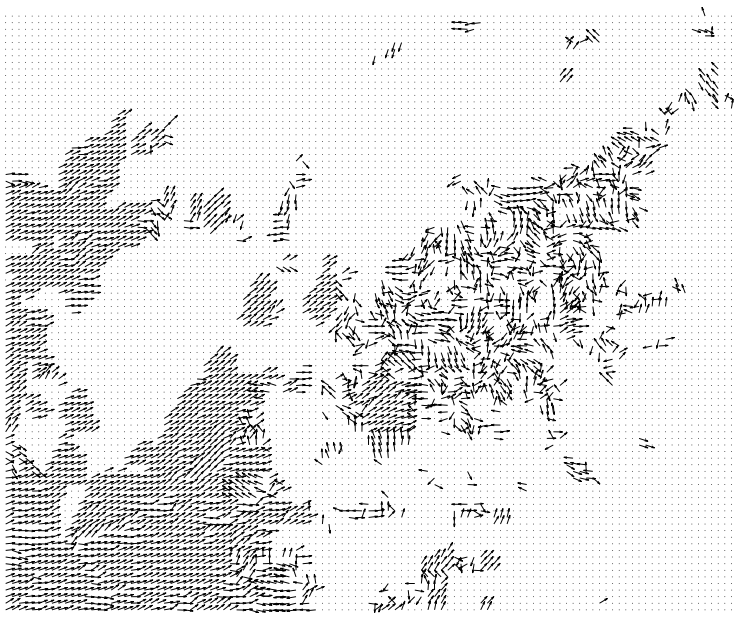

Figure 4.4: Optical flow. The direction of flow is homogeneous for precipitation but inhomogeneous for clutter. The speeds were normalized to 1 for clarity.

Figure 4.5: HIRLAM wind field. The direction of flow is homogeneous for the whole region. The speeds were normalized to 1 for clarity.

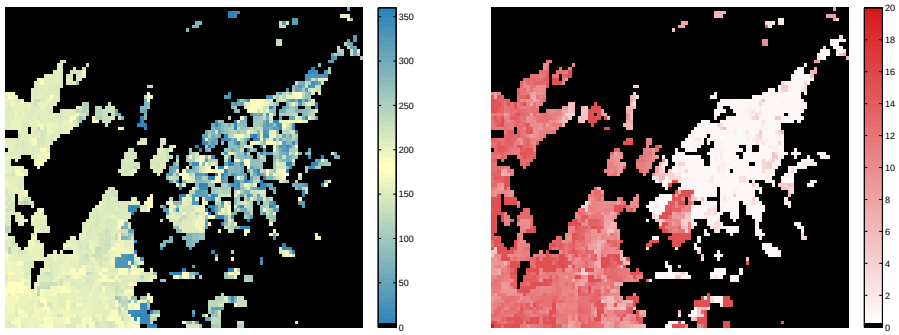

(a) Direction (deg.) **(b)** Speed (m/s)

Figure 4.6: Optical flow, direction and speed. 2005–09–25 18:00. Precipitation areas show homogeneous directions and high speeds. Clutter areas show fast changing directions and low speeds.

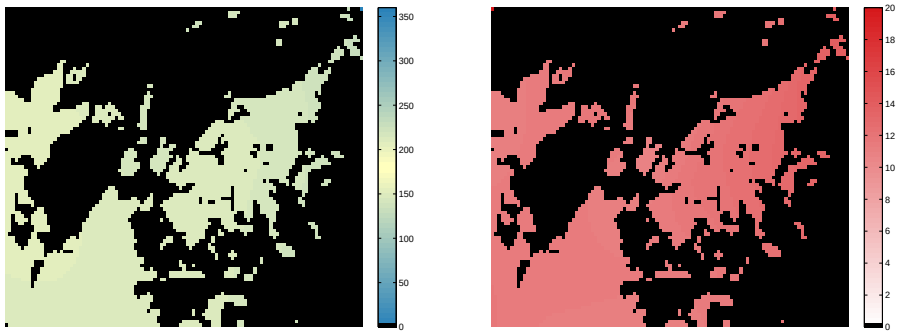

(a) Direction (deg) **(b)** Speed (m/s)

Figure 4.7: HIRLAM numerical weather prediction, direction and speed. 2005–09–25 18:00. The flow is homogeneous over the full region.

4.3.1 Entropy feature

To identify areas of clutter and precipitation, a measure is needed which can identify areas of slow and fast changes in the detected flow field. For this the *entropy* measure was chosen which can capture the diversity of the distribution of the values of the flow field in a region. The entropy of a population is a measure of the distribution the energy over its histogram. The entropy can also be described as a type of diversity index, i.e., how many different values are taken. The entropy $H(X)$ is defined as: $H(X) = -\sum_{i=1}^{n} p(x) \log_2 p(x)$, where

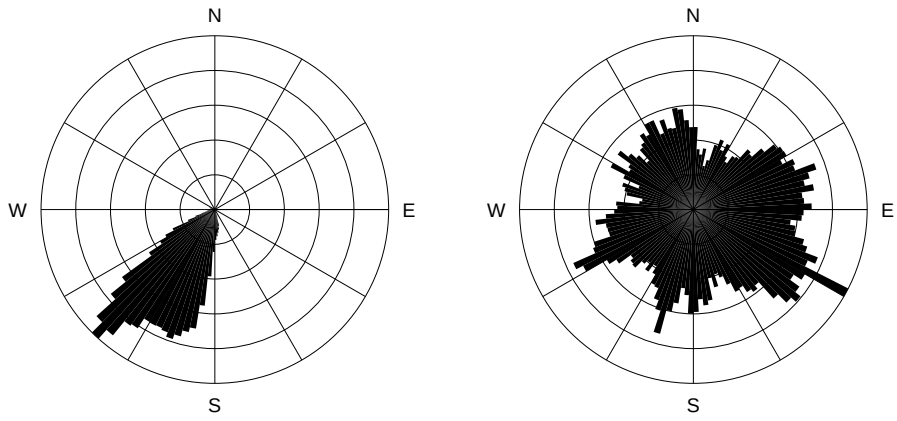

(a) Precipitation **(b)** Clutter

Figure 4.8: Rose diagrams for direction of the detected optical flow for precipitation (low entropy) and clutter (high entropy). Training data from 2005–09–25 18:00.

$p(x)$ is the values of the normalized histogram of the azimuth angles or speed values x using n bins.

The entropy was computed using a pixel-based and object-based approach.

Pixel-based entropy feature

For the pixel-based approach the entropy was computed using a moving window in form of a circular kernel of radius 31 pixels. From the direction or speed values inside this kernel the entropy was computed with the number of bins n set to 360. The entropy thus takes values between 0, if the values are gathered in one bin only, and $log_2(360) = 8.49$ if the values are distributed uniformly.

Thus, if the values of direction and speed are very similar inside the kernel, a low entropy is achieved but, on the other hand, if the values are distributed over many different values a high entropy value is taken. The entropy of the flow field is thought to take high values where the flow field is changing fast (clutter) and low values in areas of slow changing flow (precipitation).

The entropy was computed for all optical flow direction and speed images. As an example, the direction and speed and their corresponding entropy features are shown for *Case II* in Figures 4.11 and 4.12. Looking at the optical flow direction (Figures 4.11, middle) the precipitation area in the middle of the image shows quite homogeneous directions (yellow and light blue hues, corresponding to the

direction between 200 and 240). The sea clutter in the Baltic Sea, however, shows faster changing directions (and a very different direction overall than the precipitation in that the general color is blue). The land clutter close to the Sindal radar shows the same fast changing direction pattern, while the precipitation area in the North Sea shows a moderately changing direction pattern of blue hues.

These characteristics are reflected in the entropy images (Figure 4.12). The color maps for the entropy were chosen so that red hues correspond to entropy values which are more likely to be clutter and blue hues are more likely to be precipitation. The threshold value between red and blue was chosen to be 7.7 which was deduced from where the conditional probabilities for clutter and precipitation cross (See Appendix C, Figure C.15, top right plot, for example). As seen from the entropy of the optical flow direction (Figure 4.12, middle) the overall quality of that feature is not very impressive. It seems to capture some of the characteristics of the differences between clutter and precipitation, e.g., the precipitation area in the middle of the image has predominantly low entropy values and the sea clutter in the right hand side have predominantly high entropies. Also the land clutter close to the Sindal radar has all high entropy. The precipitation area in the left hand side has mostly high entropy, whereas it would be expected to have low entropy (precipitation). The changes in entropy are quite rapid in places which is caused by the somewhat noisy optical flow estimate.

The entropy of the optical flow speed (Figure 4.12, bottom) shows a surprising result: Low entropy of is found in the area the land clutter and the precipitation and sea clutter are of mostly of high entropy. As explained earlier, it was expected that the precipitation would be the class with low entropy also in the speed component of the flow. In Figure 4.14 the optical flow and entropy features are shown for all three case examples and similar conclusion can be drawn from the two other cases, *Case I* and *III*.

The entropy feature histograms (Figure 4.9) show that the observations above are general for the complete training data set. For the entropy of the optical flow direction it is noticed that precipitation generally has lower entropy measure than clutter. It is also seen that land clutter has the highest directional entropy of the two clutter classes. Sea clutter and precipitation are more overlapping.

The speed entropy, as mentioned, shows a different pattern: The precipitation and sea clutter are indiscernible with high entropy values while land clutter has low entropy, i.e., the local pace of changes in the speed is slower for land clutter than for precipitation and sea clutter.

The feature histograms and conditional probabilities for all echo strengths and

two and three class cases can be found in Appendix B, Figures B.15 to B.20 and Appendix C, Figures C.15 to C.20.

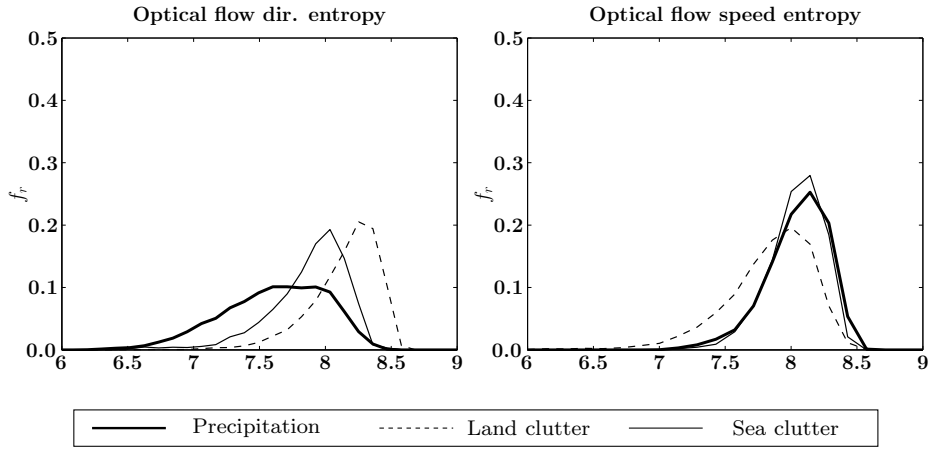

Figure 4.9: The feature histograms of the entropy of the optical flow direction (left) the speed (right) for clutter and precipitation classes. Total number of samples: 1.140.697. See Appendix C, Figure B.15 for the two class case.

From the feature histograms, showing great overlap between the classes, and the case examples, showing noisy entropy and difference patterns, is was realized that it would be difficult to perform good classification of clutter and precipitation using the pixel-based entropy features. Therefore it was investigated if evaluated on an object-basis the entropy features could provide better discrimination between clutter and precipitation.

Object-based entropy feature

Object-based image analysis is used increasingly the field of remote sensing [13] and the concept in object-based image analysis is to analyze groups of pixels belonging to real-world objects, i.e., a land use type, a forest type, a cloud type, in stead of analyzing on the pixel-level only. One of the corner stones of object-based image analysis is the successful *segmentation* of the image by which groups of pixels are labelled according to the object they belong to.

A simple method for segmentation was chosen for this study. A 4-connectivity connected component labelling [22] was applied which identifies groups of radar echoes which are connected to each other. It is assumed that objects of clutter and precipitation are separated by at least one pixel. This is a quite valid

assumption because precipitation and clutter are often separated in space, especially for the most common clutter type, anomalous propagation clutter, which occurs mostly in weather situations where precipitation is less likely.

For all objects[1] in the pixel-based entropy images (Figure 4.12) the mean of the entropy of the pixels within each object was computed. This resulting mean entropy was assigned to all pixels of that object. The example results can be seen in Figure 4.13. Computation of the entropy statistics for objects result in homogeneous regions and it is seen how the precipitation area now have mostly blue or magenta hues and the clutter areas have mostly red hues (corresponding to precipitation and clutter).

Inspection of the histograms of the object-based entropy features (Figure 4.10) reveals that the classes are better separated for the entropy of the direction. Land clutter is especially well separated from the precipitation while sea clutter has some overlap. For the entropy of the optical flow speed, only land clutter is well separated and the sea clutter and precipitation still overlap.

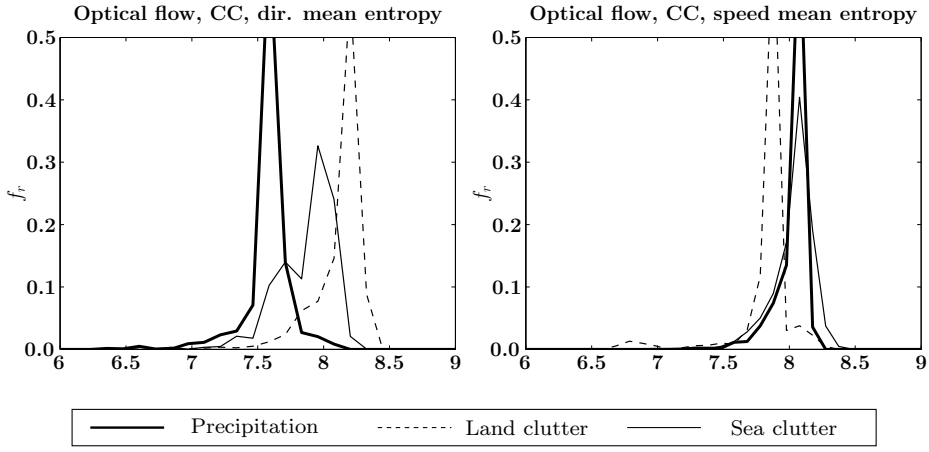

Figure 4.10: The feature histograms of the object-based entropy of the optical flow direction (left) and speed (right) for clutter and precipitation classes. Total number of samples: 1.140.697.

[1]The terms *object* and connected-component (abbreviated CC) are used synonymously in this thesis

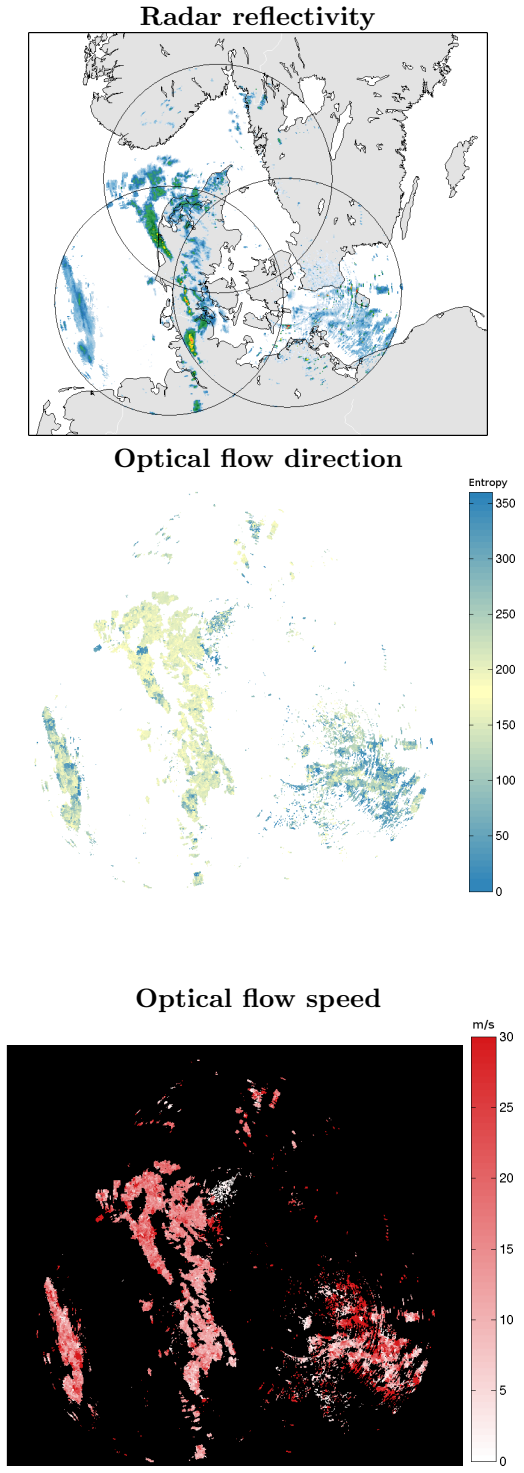

Figure 4.11: Optical flow features, *Case II*.

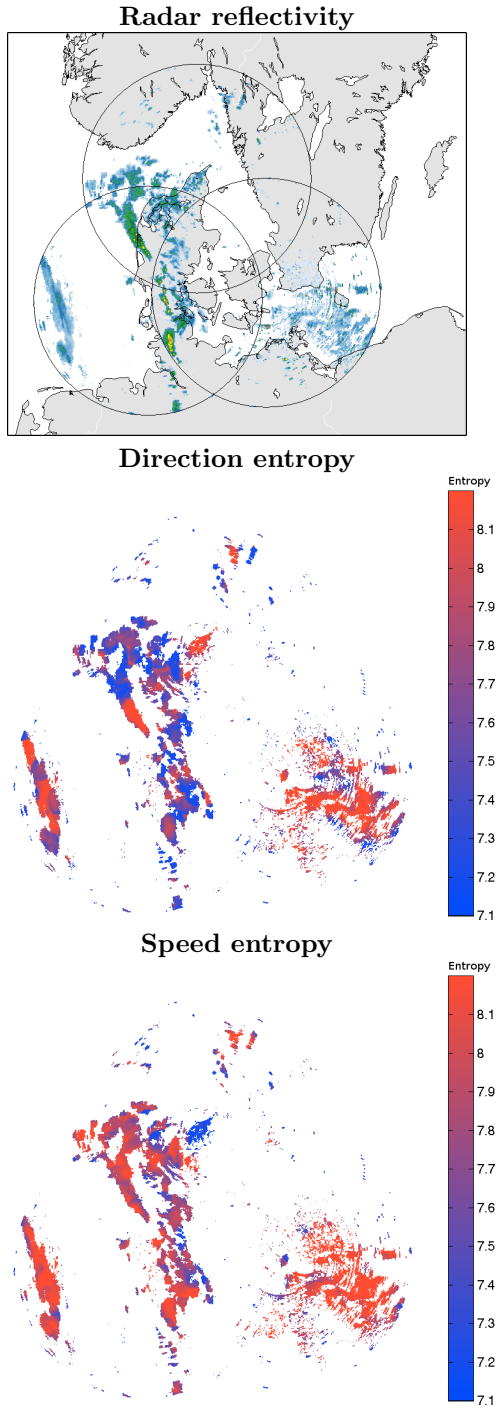

Figure 4.12: Entropy features, pixel-based, *Case II*.

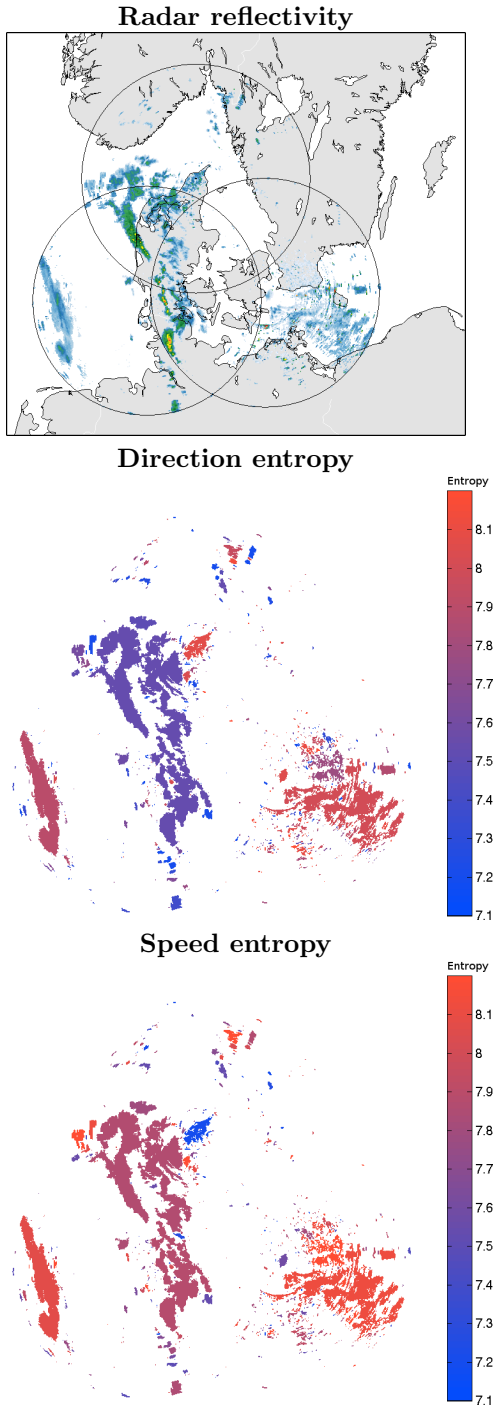

Figure 4.13: Entropy features, object-based, *Case II*. Compare with Figure 4.12.

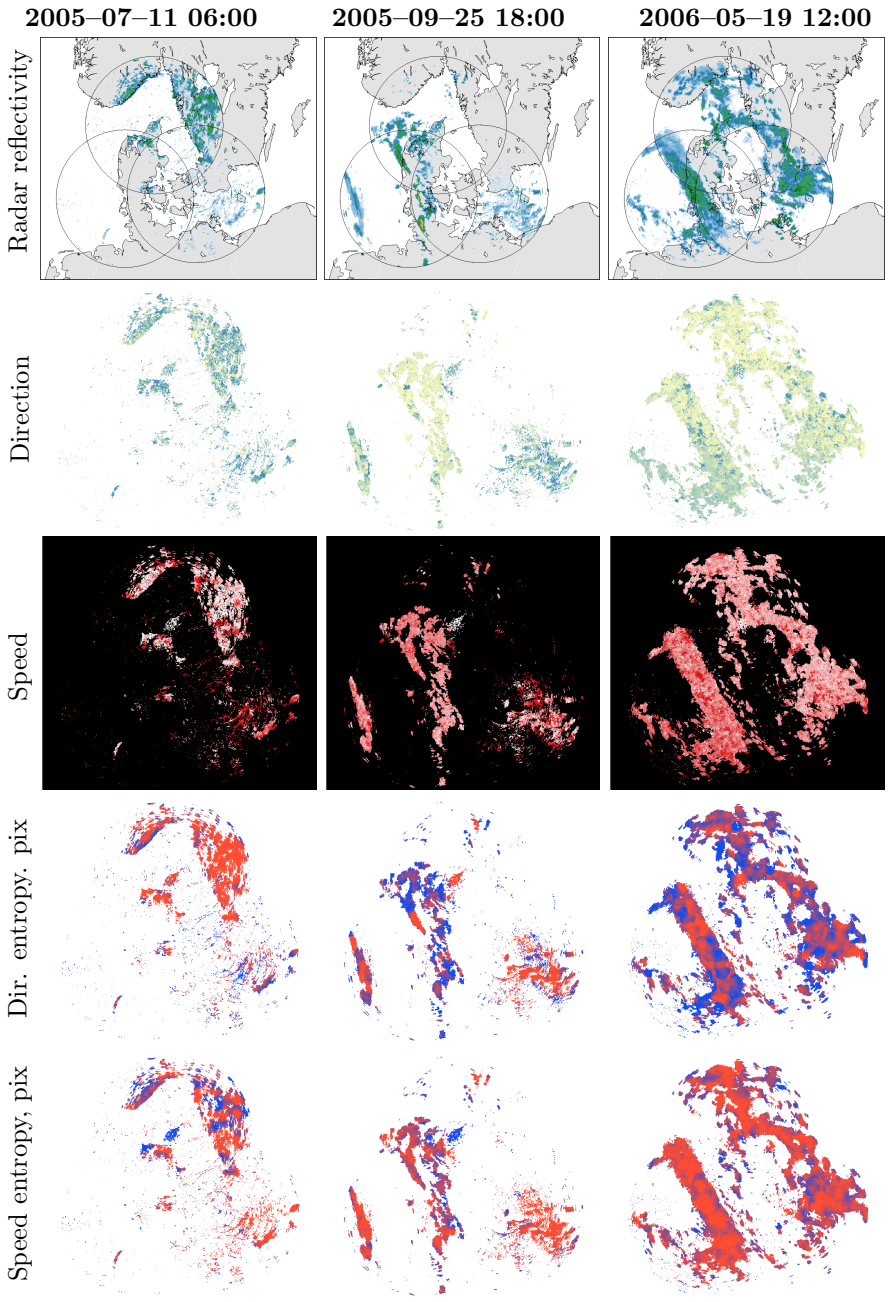

Figure 4.14: Optical flow features (direction and flow) and entropy features for the case examples.

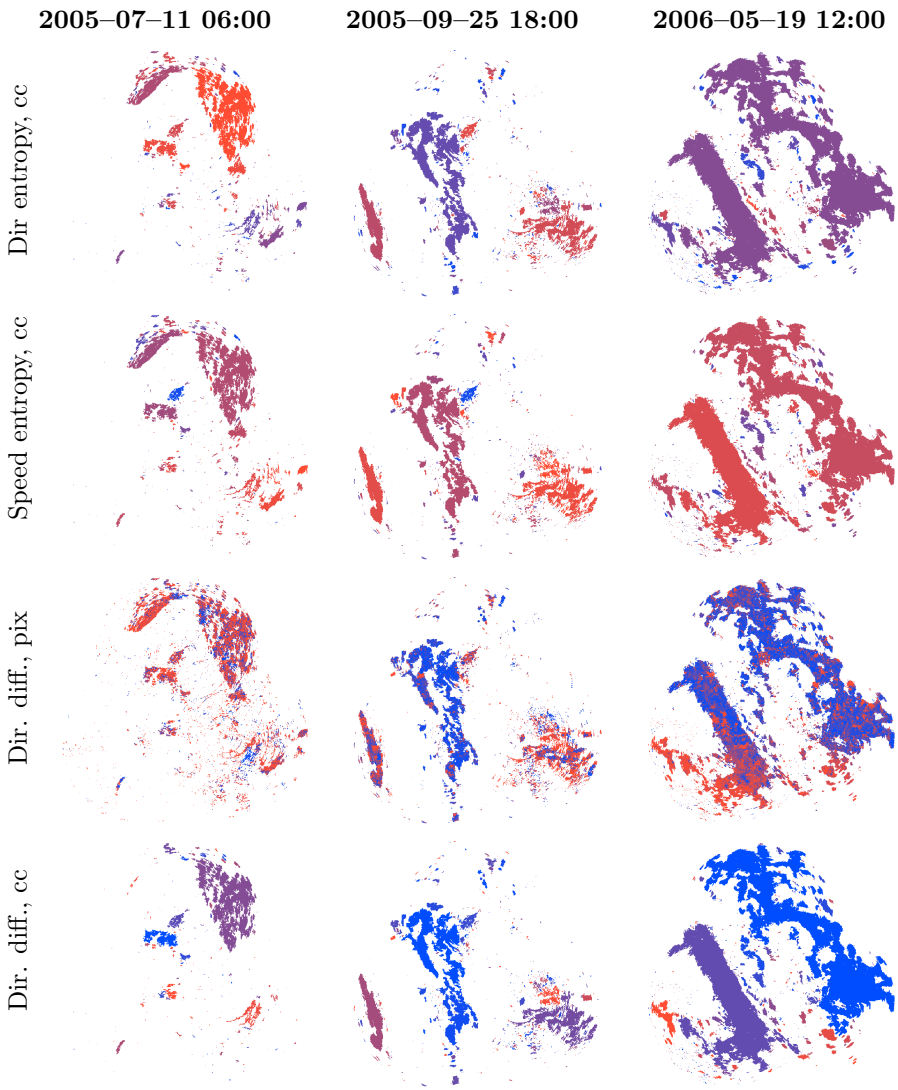

Figure 4.15: Optical flow features, object-based.

4.3.2 NWP wind difference feature

As seen in Figures 4.9 and 4.10, sea clutter shows quite similar entropy values as precipitation, which will lead to increased misclassification. Some of the time, however, it would be possible to eliminate some of the confusion between the classes by comparison with the wind field at that current time. I.e., when the detected motion of radar echoes is very different from the direction of the atmospheric wind, it is more likely to be clutter. Therefore the wind field was added as an additional source of information. By subtracting the wind field from the optical flow field, areas of similar or different velocity are highlighted.

The direction and speed difference can be seen for *Case II* in Figure 4.16. The color map for the figures was chosen so small differences are blue (corresponding to precipitation) and larger differences are red (corresponding to clutter). It is seen how in general the precipitation takes small values and clutter areas take larger values. The western precipitation area, however, shows some large direction differences, and the eastern most sea clutter shows some low differences in speed. Similar observations can be made from the difference features computed using the object-based method (Figure 4.17). In Figure 4.15, the difference features are seen for all three case examples.

The feature histogram for the pixel-based difference features (Figure 4.18) shows that precipitation in general has differences close to zero degrees as expected, but some spread around zero is seen. Land and sea clutter have much higher proportions of large and small differences. A preferences for negative and positive differences are seen for land clutter and sea clutter, respectively, which can be explained by the bias in the training data set for Southwesterly wind directions.

The distribution of speed differences shows that the distributions for land clutter and precipitation are overlapping and very similar, whereas sea clutter has a larger proportion of large positive differences. A bias in the speed difference is noticed in that the differences for precipitation are not centered on 0 m/s as would be expected, and that the distributions are skewed towards positive differences. This indicates that either the optical flow speeds have a positive bias or the NWP winds have a negative bias.

Looking at the object-based difference features (Figure 4.19) the same general trend is seen, but again with a better separation between the classes.

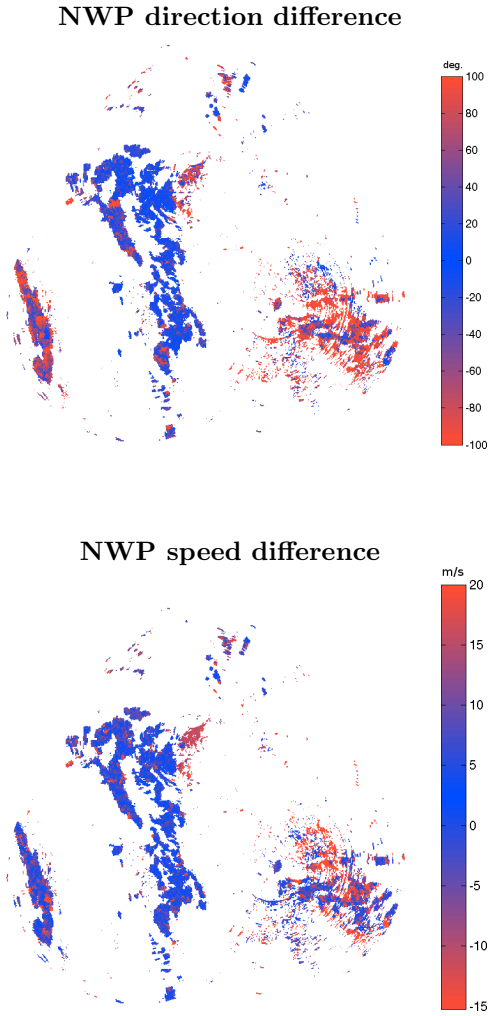

Figure 4.16: NWP difference features, pixel-based, Case II.

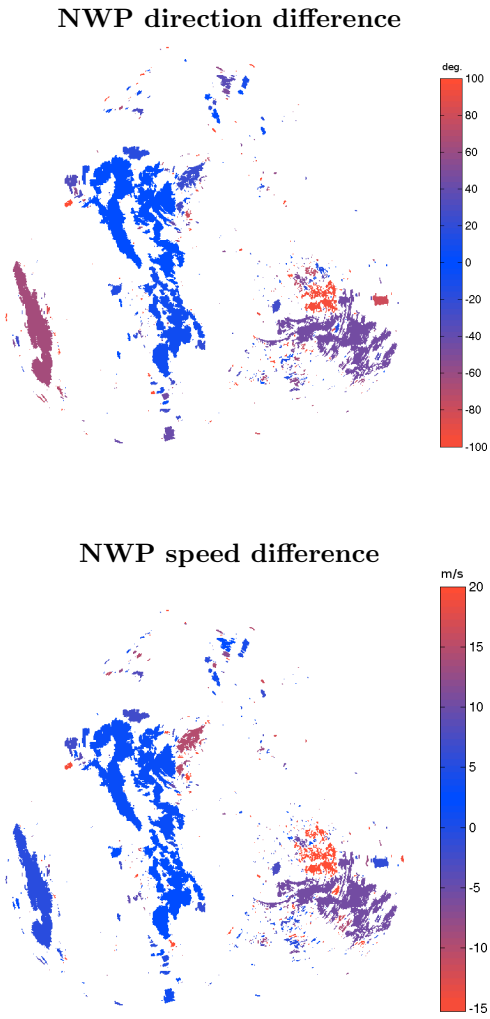

Figure 4.17: NWP difference features, object-based, Case II.

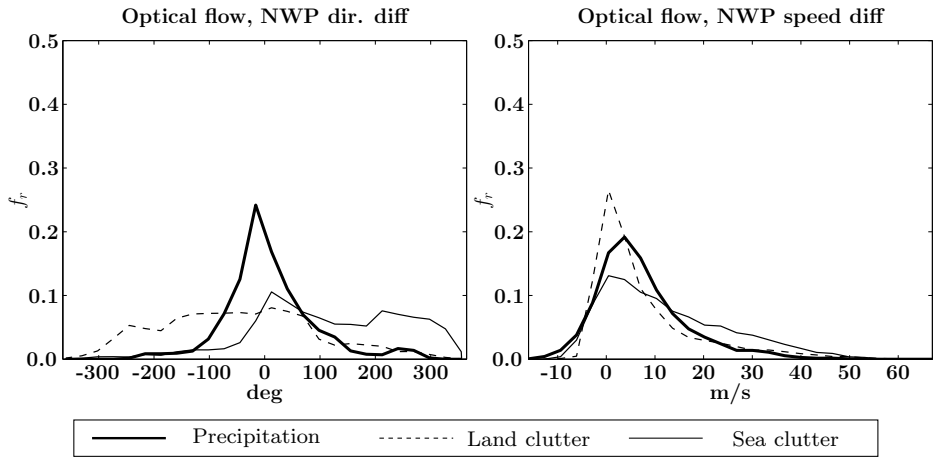

Figure 4.18: Feature histograms of the pixel-based difference between optical flow direction (left) and speed (right) and NWP direction and speed for clutter and precipitation classes. Total number of samples: 1.140.697.

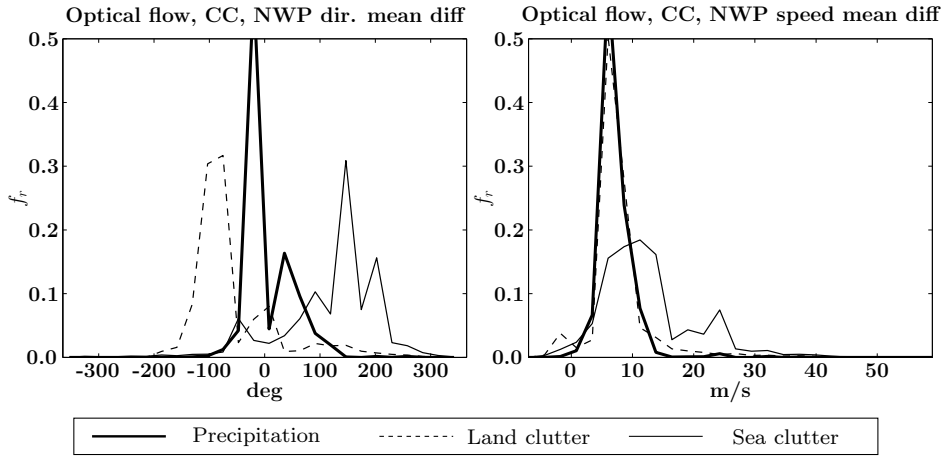

Figure 4.19: Feature histograms of the object-based difference between optical flow direction (left) and speed (right) NWP direction and speed for clutter and precipitation classes. Total number of samples: 1.140.697.

4.4 Feature selection

The entropy features and difference features were subjected to feature selection using the same procedure as described in detail in Section 3.8 in Chapter 3.

For the pixel-based method, the entropy of the direction and speed and the direction and speed NWP differences were used as features. For the object-based method, the corresponding object-based entropy and difference features were used and added to these were the variance of the direction and speed of each object.

The feature selection results are shown in Appendix D, Tables D.13 to D.18 for the pixel-based features and in Tables D.19 to D.24 for the object-based features.

The results of the feature selection for the pixel-based features show that the best single feature is the entropy of the optical flow direction. For using two features the direction entropy is still consistently chosen, together with either the direction difference or the speed difference. The chosen number of features are either three or four, i.e., all the features.

For the object-based features, the single best feature was either the direction entropy or the direction variance, which was also consistently chosen for two features together with the velocity difference. The chosen sets consists of both the entropy and difference features which shows that both features are significant in achieving a high classification accuracy.

4.5 Classification

Classification was carried out using the same approach as described in the previous chapter, in which an ensemble of five classifiers were applied to each chosen feature subset and combined using a majority vote between the individual classifiers.

4.6 Classifier evaluation

Using the methodology in Section 4.6, Chapter 3, the classifiers were evaluated by classification using independent test samples and several skill scores were computed.

4.6.1 Skill score results

The classification evaluation results are found in Tables 4.1 to 4.4.

Pixel-based method

For the pixel-based method (Tables 4.1 and 4.2) it is seen that the majority classifier achieves percent correct scores from 76.20 % (weak echoes, two classes) to 84.84 % (strong echoes, two classes). For all echoes, the best PC score is 79.36 % using two classes. The single best classifier is the quadratic discriminant analysis, while the k-nearest neighbor showed the worst performance. The majority classifier again scores consistently better than the worst single classifier, and quite close to the single best one in some cases. In one case the majority vote classifier achieved a higher PC value than the single best (84.84 % for the majority classifier while the QDA scored only 84.52 % (for strong echoes and two classes.

The FAR scores for the majority classifier range between 0.161 and 0.289 indicating a high number of false alarms in the detection of clutter. The best score for all echoes was 0.176 using two classes. The Hanssen-Kuipers skill score also shows mediocre values for all echo types. The best score for all echoes is a mere 0.582 for two classes.

While the pixel-based method definitely has skill at removal of clutter it is at a mediocre accuracy, a high false alarm rate and mediocre HKS scores.

Object-based method

The picture is very different for classification using the object-based features. PC scores range between 94.16 % and 98.03 % for the majority classifier, with the highest PC for all echoes being 96.22 % for two class classification. The best PC value was for strong echoes, three classes (98.03 %).

Looking at the individual classifiers the decision tree (Dtree) scores on average more than a 2 % better than the other classifiers.

The false alarm ratio ranges between the low values of 0.019 and 0.045 (strong and weak echoes, three and two classes), and for all echoes the best FAR was 0.031. The HKS scored high values between 0.887 and 0.961 with the best all echoes value of 0.924 for two classes.

Echo types

As for the data fusion method, the best discrimination between clutter and precipitation is for strong radar echoes with PC values several percent better than for weak echoes.

Number of classes

No improvement is seen by using three class classification in comparison with two class classification. Both PC, FAR, and HKS scores generally worse for the three class case. These results are the opposite of the expected, because the motion fields of land and sea clutter and precipitation are different both in the feature histograms and when judged visually, which were thought to improve classification using all three classes separately.

Table 4.1: Classification results. Spatio-temporal method, pixel-based, two classes.

Echo	Percent correct (PC)						False alarm ratio (FAR)						Hanssen-Kuipers skill (HKS)					
type	LDA	QDA	Dtree	KNN	SVM	Majority	LDA	QDA	Dtree	KNN	SVM	Majority	LDA	QDA	Dtree	KNN	SVM	Majority
All	77.22	79.51	75.28	61.94	77.15	79.36	0.198	0.166	0.201	0.315	0.189	0.176	0.538	0.589	0.505	0.242	0.541	0.582
Weak	72.85	77.48	70.43	61.16	72.68	76.20	0.217	0.190	0.208	0.275	0.215	0.198	0.422	0.509	0.402	0.225	0.422	0.484
Strong	83.09	84.53	81.36	65.23	83.47	84.84	0.183	0.166	0.190	0.355	0.162	0.161	0.662	0.691	0.627	0.305	0.669	0.697

Table 4.2: Classification results. Spatio-temporal method, pixel-based, three classes.

Echo	Percent correct (PC)						False alarm ratio (FAR)						Hanssen-Kuipers skill (HKS)					
type	LDA	QDA	Dtree	KNN	SVM	Majority	LDA	QDA	Dtree	KNN	SVM	Majority	LDA	QDA	Dtree	KNN	SVM	Majority
All	76.36	79.28	74.71	62.38	76.25	78.13	0.250	0.228	0.258	0.356	0.250	0.243	0.495	0.554	0.464	0.210	0.493	0.526
Weak	75.71	76.23	70.72	68.80	75.57	76.23	0.220	0.230	0.256	0.266	0.218	0.222	0.455	0.448	0.350	0.315	0.456	0.459
Strong	77.09	79.78	77.67	62.54	77.41	78.63	0.299	0.276	0.286	0.410	0.295	0.289	0.549	0.603	0.560	0.257	0.555	0.580

Table 4.3: Classification results. Spatio-temporal method, object-based, two classes.

Echo	Percent correct (PC)						False alarm ratio (FAR)						Hanssen-Kuipers skill (HKS)					
type	LDA	QDA	Dtree	KNN	SVM	Majority	LDA	QDA	Dtree	KNN	SVM	Majority	LDA	QDA	Dtree	KNN	SVM	Majority
All	93.88	95.86	97.75	94.47	94.91	96.22	0.036	0.043	0.014	0.058	0.038	0.031	0.881	0.914	0.956	0.884	0.899	0.924
Weak	91.18	93.97	96.84	94.31	92.35	94.52	0.049	0.047	0.018	0.044	0.047	0.039	0.828	0.872	0.938	0.880	0.847	0.887
Strong	95.46	96.61	98.77	95.11	96.69	97.80	0.016	0.025	0.012	0.081	0.030	0.019	0.908	0.932	0.975	0.903	0.934	0.956

Table 4.4: Classification results. Spatio-temporal method, object-based, three classes.

Echo	Percent correct (PC)						False alarm ratio (FAR)						Hanssen-Kuipers skill (HKS)					
type	LDA	QDA	Dtree	KNN	SVM	Majority	LDA	QDA	Dtree	KNN	SVM	Majority	LDA	QDA	Dtree	KNN	SVM	Majority
All	93.48	94.73	97.19	94.94	93.74	96.08	0.049	0.057	0.029	0.064	0.054	0.041	0.870	0.889	0.942	0.890	0.873	0.918
Weak	91.35	93.02	96.93	94.63	91.11	94.16	0.057	0.059	0.024	0.053	0.066	0.045	0.825	0.849	0.935	0.879	0.815	0.877
Strong	96.95	96.89	98.81	94.95	97.45	98.03	0.029	0.052	0.015	0.086	0.034	0.031	0.939	0.939	0.976	0.901	0.949	0.961

4.7 Results

In this section the method is evaluated on the three case examples again to show the performance of application of the method on single images.

The classified images and "cleaned" (filtered) radar images can be found for the pixel-based method in Figure 4.20 and for the object-based in Figure 4.21.

The classification for the pixel-based method shows mediocre accuracies (as expected) between 73.7 % and 83.5 %. Many misses and false alarms are observed and the misclassification are distributed all over the radar images. The misses and false alarms can also be seen in the confusion matrices.

The object-based method, on the other hand, shows results that are very good with values between 96.2 % and 99.6 %. The best is *Case I* where almost all the clutter is correctly classified and the small precipitation area in the lower North Sea is also classified correctly as well. For *Case II* and *III* there are a number of smaller pixel objects misclassified which is probably due to insufficient flow statistics for objects with a small number of pixels.

4.8 Summary and conclusion

In this chapter, a method for clutter detection using spatio-temporal features in radar data combined with wind field information from a numerical weather model was presented.

The method firstly identifies the optical flow between two consecutive radar images and from the flow field the entropy of the direction and speed is computed. It was found that precipitation has lower entropy than clutter of the detected flow *directions* (The change in the direction of the motion field is slower for precipitation than for clutter). Land clutter showed a higher directional entropy in comparison with sea clutter. For the entropy of *speed*, low values were surprisingly observed for land clutter, while sea clutter and precipitation took higher values.

Wind field information from a numerical weather prediction model was used to derive a difference feature in the form of the difference between the optical flow field and the NWP field. High differences were found for areas of clutter while low differences were found for precipitation.

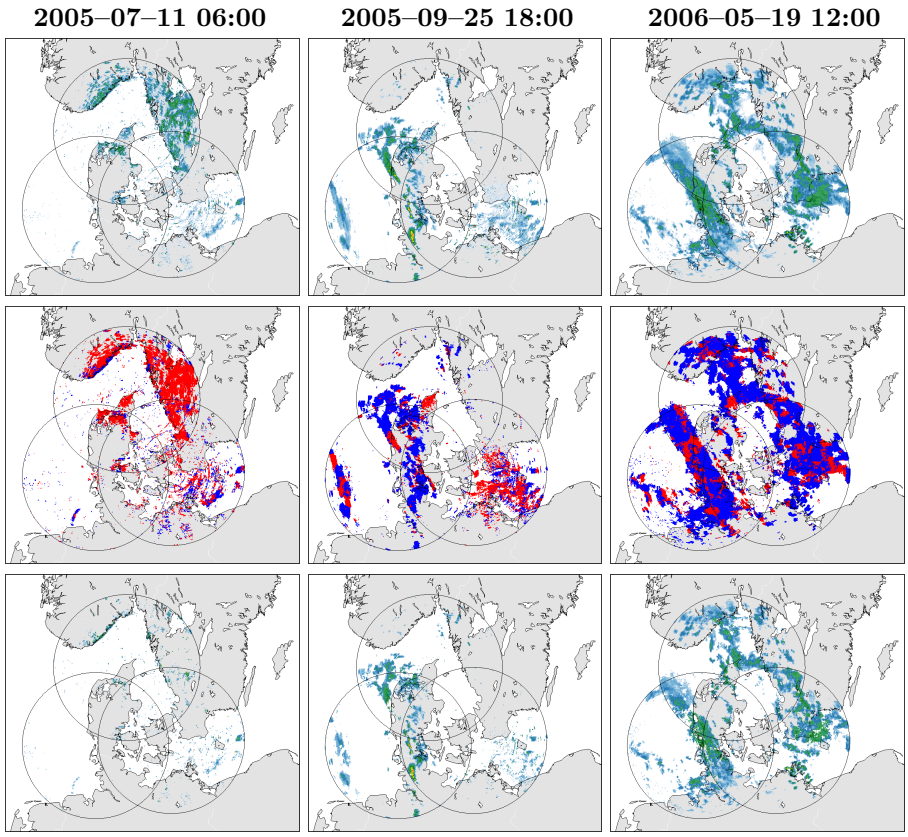

83.5 %

	C_c	P_c
C_t	33169	6534
P_t	74	303

73.7 %

	C_c	P_c
C_t	11217	5553
P_t	7626	25680

75.0 %

	C_c	P_c
C_t	0	0
P_t	34231	102835

Figure 4.20: Classification results using the spatio-temporal pixel-based method on the three case examples. Top row: Radar images, Middle row: Classified images (blue is precipitation and red is clutter), Bottom row: Filtered radar images. The classifications accuracy and confusion matrix are found below (Subscript t means *truth* and subscript c means *classification*, cf. Figure 3.27)

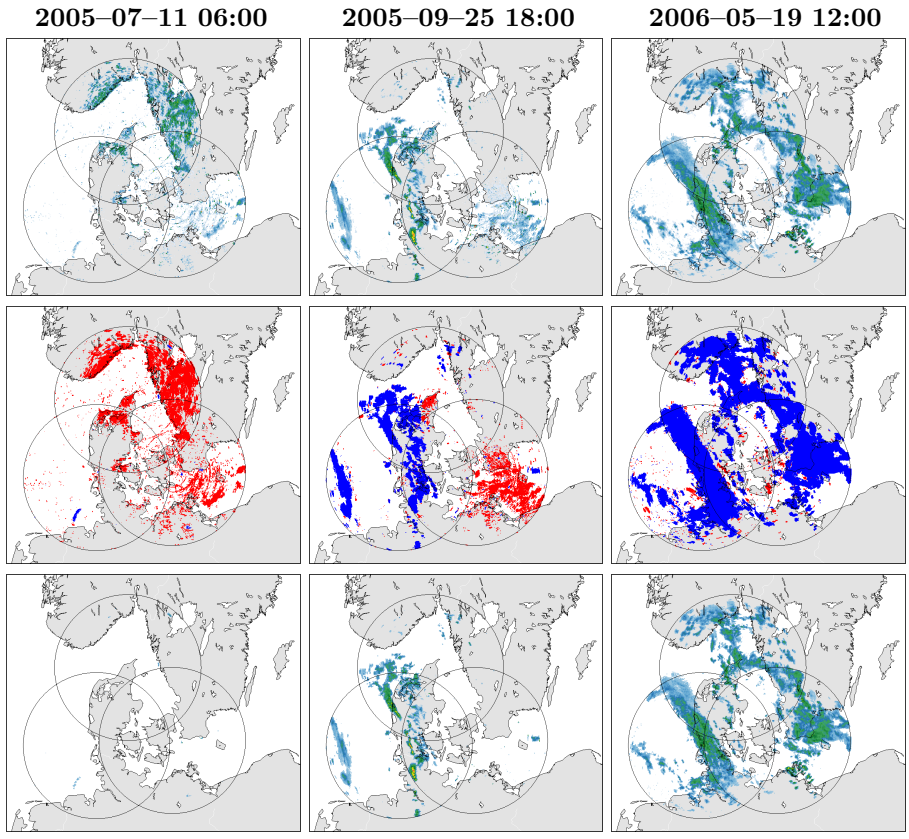

99.6 %

	C_c	P_c
C_t	39574	129
P_t	44	333

96.9 %

	C_c	P_c
C_t	16042	728
P_t	812	32494

96.2 %

	C_c	P_c
C_t	0	0
P_t	5180	131886

Figure 4.21: Classification results using the spatio-temporal object-based method on the three case examples. Top row: Radar images, Middle row: Classified images (blue is precipitation and red is clutter), Bottom row: Filtered radar images. The classifications accuracy and confusion matrix are found below (Subscript t means *truth* and subscript c means *classification*, cf. Figure 3.27)

The features were computed on a pixel-basis and an object-basis. The object-based method assumed clutter and precipitation to be separated in space and provided more robust statistics of the entropy and difference measures.

Classification of the input features were carried out using five different classifiers, which were combined through a majority vote. The pixel-based features were not very successful, and low skill scores were obtained. The classification accuracy for the pixel-based features was estimated to be 79.36 % for all echoes using two class classification. The obtained FAR was 0.176 and the HKS 0.582. These scores are not encouraging the use in any operational applications. The object-based features, on the contrary, showed good to very good performance with high accuracy and low false alarm ratios. The accuracy for all echoes was 96.22 %, the false alarm ratio was 0.031, and the Hanssen-Kuipers skill score 0.924, all for two class classification.

Using one classifier only, the quadratic discriminant analysis performed the best for the pixel-based method, while the decision tree classifier was consistently the single best classifier for the object-based method.

The case examples were classified with mediocre results for the pixel-based method: a PC value of up to 83.5 % only and a large proportion of misclassified radar echoes. The object-based method, however, showed results that were very good. PC values up to 99.6 % were seen with low FAR and high HKS scores.

The pixel-based method could probably be improved upon by using the flow information from more than two images. In fact, when a radar user applies visual clutter detection, more than two images are typically used (maybe in the range of four or five images are used). The proposed method uses only two and future work on the method could include extension of the method to use statistics from the flow of image series of more than two images.

Another issue that could be improved upon is the time interval between radar images. The sometimes large confusion between precipitation and clutter in the optical flow field indicates that the motion of precipitation is undersampled at the 10 minute interval between images. Radar images with a higher temporal resolution is likely to improve the detection method. An increase of the temporal resolution of radar data acquisition is under implementation for the weather radar network of the Danish Meteorological Institute at the time of writing.

Finally, the parameters of the optical flow algorithm were chosen from a limited number of empirical tests of various combinations of size of the matching kernel and number of pyramid levels. Optimization of these parameters to enhance discrimination of clutter and precipitation motion patterns could also be part of future work.

Chapter 5

Propagation method

5.1 Introduction

As described in the review of methods for clutter detection, a group of clutter detection methods aim at predicting the meteorological conditions which cause anomalous propagation and potential anomalous propagation clutter.

The path of propagation of electromagnetic waves in the atmosphere is determined by the atmospheric refractive index and hence using observations or predictions of the geophysical parameters (temperature, pressure and humidity) which determine the refractive index, it is possible to forecast anomalous propagation.

Prediction of the propagation can be done using analyses of the vertical gradients of the refractive index or by actual simulation of the propagation path of the weather radar's electromagnetic energy through a two or three dimensional refractivity field. Propagation methods use ray tracing or wave propagation techniques (or combinations of these). Sources of data for determining the refractive index of the atmosphere are either observations (typically from radiosondes) or predictions from numerical weather prediction models.

In this chapter, two propagation methods for clutter detection are presented. The first method uses a wave propagator to simulate areas of potential ground clutter and the second method uses an azimuth-dependent propagation index which is derived from vertical refractivity gradient measures.

In the next two sections, radar propagation and the used data are described. Then in the next sections (Sections 5.4 and 5.5) the two methods are described and evaluated.

5.2 Radar propagation

The theory of propagation of microwaves in the atmosphere is well established [43], and the key to determining the propagation path of radio waves is to estimate the variation of the refractive index.

For air in the atmosphere, the refractive index n is close to 1 (the refractive index of vacuum) with typical values around 1.0003 at sea-level. Since the changes in the refractive index of the atmospheric air are usually very small, the refractivity N, one million times n-1, is used instead. The refractivity at sea-level is thus normally 300. The refractivity N for a point in the atmosphere can be computed from the temperature, air pressure, and water vapor pressure in that point using:

$$N = (n-1) \times 10^6 = \frac{77.6}{T}(p + 4810\frac{e}{T}), \tag{5.1}$$

where T is the temperature (K), p is the air pressure (hPa), and e is the water vapor pressure (hPa) [52].

For a standard atmosphere, the refractive index decreases with altitude which will result in radio waves getting bent toward the surface of the earth. This is illustrated in Figure 5.1 by successive application of Snell's law for four layers with decreasing values of n [52].

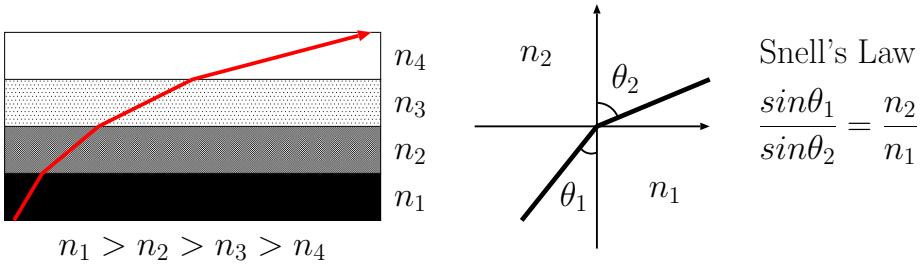

Figure 5.1: For decreasing refractive indices n_1 to n_4 with height the radio wave (red) is bent downwards, cf. Snell's Law.

The degree of bending of the radar waves is determined by the vertical gradient of the refractivity, and hence the refractivity gradient is a tool for predicting the propagation of radar waves. The International Telecommunication Union (ITU) proposes the measure VRG (Vertical Refractivity Gradient) [41] defined as the gradient between the refractivity at surface level, $N_{surface}$ and in 1 km height, N_1:

$$VRG = N_{surface} - N_1 \tag{5.2}$$

From the VRG, propagation conditions can be classified into three categories: Standard propagation, subrefraction and superrefraction (See Figure 5.2). The vertical refractivity gradient takes values around −40 N-units/km for a standard atmosphere (standard propagation takes place in the interval 0 and −79 N-units/km). Values larger than 0 N-units/km results in subrefraction (the wave is bent more upwards than normal, which does not cause increased clutter), while refractivity gradients below −79 N-units/km causes superrefraction and the waves are bent downwards more than normal. At values below −157 N-units/km ducting occurs, which causes the radar waves to be trapped in a layer close to the ground which acts like a wave guide.

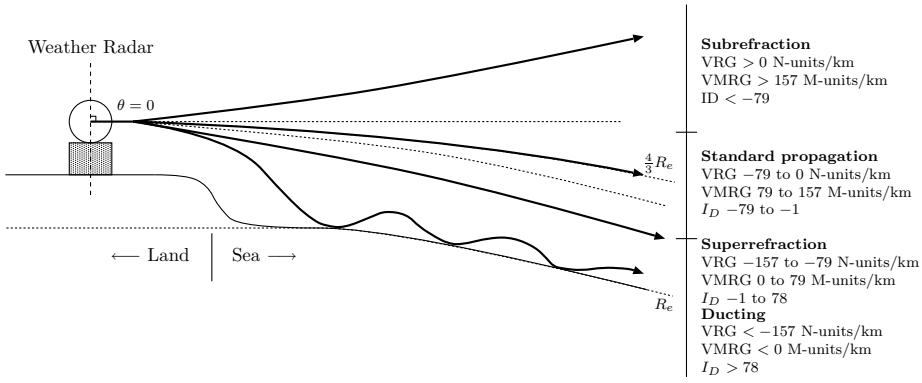

Figure 5.2: Different types of propagation in the atmosphere in relation to the vertical refractivity gradient (VRG), vertical modified refractivity gradient (VMRG), and the ducting index (I_D). Here shown for elevation angle equal to 0. The earth's local radius is R_e, and for standard propagation, the propagation follows a curve along a sphere with radius of $4/3R_e$.

Another commonly used propagation value is the modified refractivity M which can be computed from the refractivity N:

$$M = N + \frac{z}{R} \times 10^6 = N + 0.0157z, \tag{5.3}$$

where N is the refractivity, z is the height above mean sea level (m), and R is the Earth radius (m). The values for the vertical modified refractivity gradient (VMRG) is also shown in Figure 5.2.

From the modified refractivity the Ducting Index, I_D can be derived. The Ducting Index as described and used in [10] is defined as:

$$I_D = 78\delta z - \delta M, \tag{5.4}$$

where δz is the height between the points that the gradient is computed over, and δM is the modified refractivity gradient between these. The I_D is computed over the available refractivity data below 3 km and the maximum value is chosen. The Ducting Index expresses how much the modified refractivity gradient differs from the superrefraction threshold value of 78 M-units/km. The thresholds for the different propagation types can likewise be seen in figure Figure 5.2.

5.3 Data

5.3.1 Radar data

All weather radars experience radar clutter from time to time, and of the four weather radars of the Danish Meteorological Institute (DMI) especially the radar at Stevns (55.326 °N, 12.449 °E) is prone to sea clutter caused by superrefraction. The radar is situated at a height of 53 m above sea level, very close to the coast of the Baltic Sea (see Figure 5.3) and therefore a large part of the electromagnetic field from the radar is radiated over the sea. Anomalous propagation is more pronounced in the months of summer and fall, where ducting of the radar waves are frequently observed.

In the following section on wave propagation modelling, two case examples from the Stevns radar are used. In the final section on the use of a azimuth-dependent propagation index, the full database of radar data is used.

Figure 5.3: The Stevns radar seen from the sea side (Photo: Søren Overgaard).

5.3.2 Numerical weather prediction model output

For the simulation of anomalous propagation, detailed information about the refractive index and its variation in the atmosphere is needed. As described, the refractive index in a point can be determined from the temperature, pressure, and humidity and these parameters are standard outputs of numerical weather prediction models.

For this work, data was extracted from the HIRLAM numerical weather prediction model. HIRLAM (short for: HIgh Resolution Limited Area Model) is an operational NWP system used at several European meteorological institutes and it is focussed on providing high-resolution short-range weather forecasts [67]. HIRLAM consists of several nested models with different spatial resolutions for different regions (Figure 5.4), and in the DMI implementation of HIRLAM (DMI-HIRLAM) the highest resolution model is currently the S05 model. The specifications of DMI-HIRLAM is listed in Table 5.1. A long list of geophysical parameters are predicted on a grid consisting of 40 vertical levels and a horizontal grid of 0.05 ° (approximately 5 km). The spatial resolution of the HIRLAM grid is 5 km in the horizontal plane and from 30 m to 7 km in the vertical plane (highest resolution at ground level).

Output from the NWP model is extracted from 6 hour forecasts at time intervals of 10 minutes to ensure a good correspondence between the NWP and radar data.

The HIRLAM data were resampled to the same stereographic 1000 m grid as the radar data, and the refractive index, refractivity, modified refractivity, were computed for all grid nodes.

Table 5.1: Parameters of the DMI-HIRLAM models. After [67].

Model Identification	T15	S05
Grid points (mlon)	610	496
Grid points (mlat)	568	372
No. of vertical levels	40	40
Horizontal resolution	0.15°	0.05°
Hor. res. (assimilation)	0.45	–
Time step (dynamics)	360 s	120 s
Time step (physics)	360 s	120 s
Boundary age (in forec.)	6h	0h
Boundary age (in ass.)	0 h-6 h	0 h
Host model	ECMWF	T15
Boundary frequency	1/(3 h)	1/(1 h)
Data assimilation cycle	3 h	3 h
Forecast length (long)	60 h	54 h
Long forecasts per day	4	4

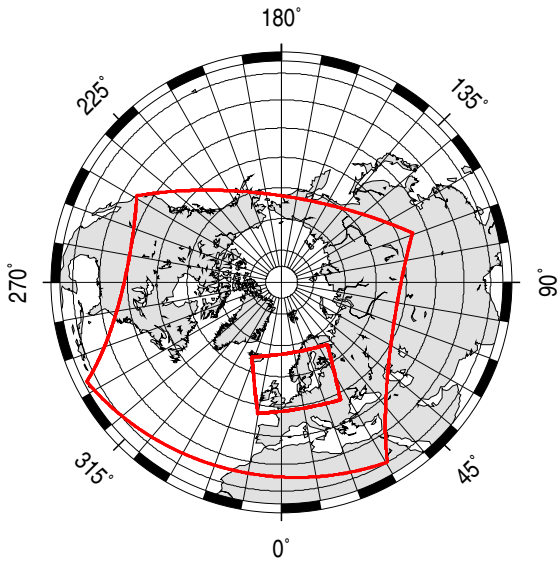

Figure 5.4: Extent of the HIRLAM models: DMI-HIRLAM-T15 (large box covering part of the Northern Hemisphere) and DMI-HIRLAM-S05 (smaller box covering Northern Europe).

5.4 Method I: Wave propagation method

In the next sections, a method for simulation and prediction of anomalous propagation clutter using a wave propagator is presented. The wave propagator uses the three dimensional refractivity field derived from the HIRLAM NWP model. The wave propagator is based on the parabolic equation approximation to the electromagnetic wave equation. The radar clutter prediction technique is applied to two cases of superrefraction propagation conditions and the potentials for detection and removal of clutter echoes are illustrated.

Other work in this field includes [2] who presented a method for simulation of anomalous wave propagation through a 2D inhomogeneous refractivity field using observed refractivity data and [5] who showed prediction of radar propagation in ducting conditions using NWP data with 5 km horizontal grid resolution as input. This approach was also used by [9] who showed the feasibility and potentials for modelling clutter using a fast hybrid parabolic equation method on 8 km resolution NWP data. Methods for propagation modelling range from ray tracing to numerical solutions to the wave equation. The ray tracer is generally the fastest method compared to wave propagators, but the ray tracer cannot model atmospheric multipath effects. A number of authors, e.g. [9], have investigated the use of hybrid methods where a ray tracer is combined with a wave propagator based on a solution of the parabolic equation. For this study, however, the method is solely based on the solution of the parabolic equation.

Electromagnetic wave propagation in the atmosphere can be described by the wave equation and this equation can be approximated by the parabolic equation (PE) in a rectangle containing the neutral atmosphere [27], [28], [52]. The parabolic equation in the simulator is solved using a split-step sine transformation [15], [47], [7].

To illustrate the wave propagator two case examples are used.

5.4.1 Case example A: 2006-05-05 12:00 UTC

Figure 5.5 shows an example of a radar image from the Stevns radar. Southeast from the radar site (in the centre of the image), an extended area of sea clutter over the Baltic Sea is seen. No precipitation was present within the radar's coverage at that time, which was confirmed by inspection of multispectral satellite images from Meteosat-8, showing cloud free conditions.

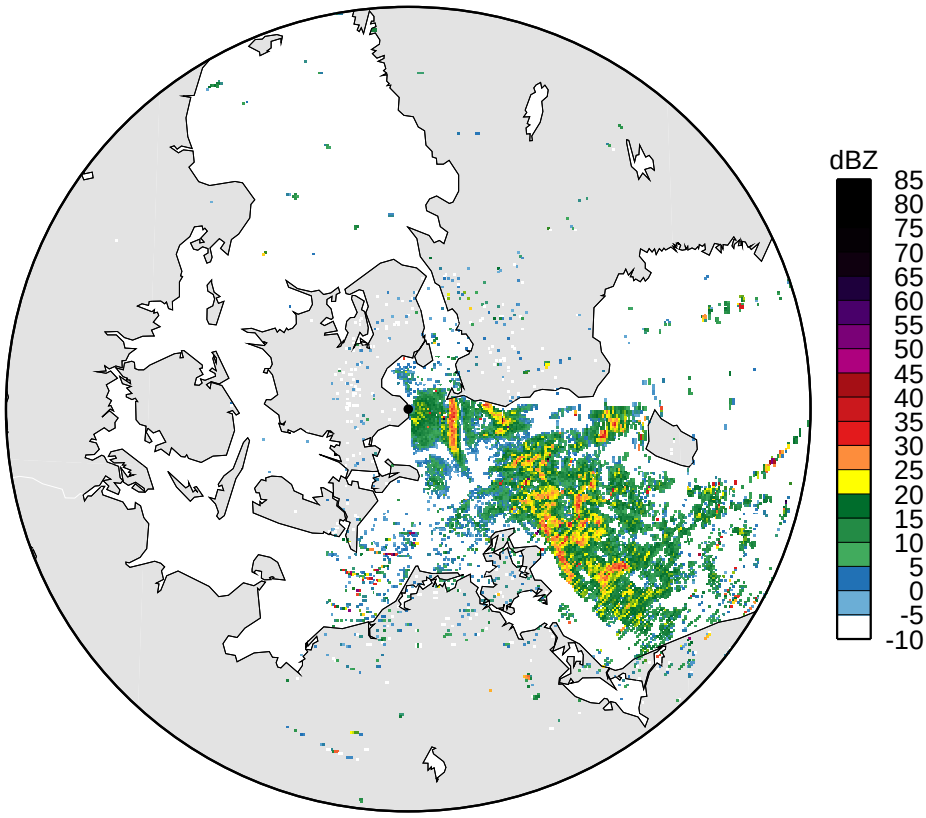

Figure 5.5: Radar reflectivity factor, Stevns radar at May 5, 2006 12:00 UTC. The maximum range of the radar is 240 km. Radar elevation angle is 0.62 degrees. All the radar echoes are sea clutter caused by superrefraction.

5.4.2 Propagation model set-up

The modified index of refraction from the NWP data is extracted along rectangles. These rectangles (of dimensions 250 km by 2 km) are placed in a circle on the surface of the earth, centered on the radar site. The angle between the rectangles is 1 degree, corresponding to a total of 360 rectangles. The 3D coordinate system that is used is positioned so that the x-axis is pointing in the west-east direction while the y-axis is pointing in the south-north direction. The azimuth angle is measured from the positive y-axis in the clockwise direction. The z-axis is perpendicular to the plane given by the x and y-axis and this axis corresponds to the height above ground level.

The wave propagator was applied to NWP data from the same time as the very pronounced sea clutter seen in Figure 5.5. Figures 5.7 and 5.8 below show the amplitude of the electromagnetic field on two of the rectangles placed at azimuth angles of 10 degrees and 130 degrees respectively. The distance between the calculated field points is 0.2 m along the z-axis and 50 m along the x-axis. The modified refractive index profiles on the rectangles are found by linear interpolation of the geophysical values in the NWP model.

Normal propagation of the simulated radar wave is seen in Figure 5.7, while in Figure 5.8, simulated superrefraction, due to atmospheric ducting and reflections in the sea, is seen. It can be seen that the wave travels along the earth surface in a 100 m high surface duct to a distance of 180 km before it escapes the duct and follows the standard behaviour. The precise wave propagation of the wave field is very depended on the azimuth angle. It is seen that the radar wave propagates a longer distance along the earth surface in Figure 5.8 compared to the normal atmospheric conditions found in Figure 5.7.

From the NWP output of temperature, pressure, and humidity for the time of the observed clutter in the radar image, it can be seen that non-standard propagation conditions indeed were present over large parts of the Baltic Sea. In Figures 5.9 and 5.10, vertical profiles of temperature, pressure, and humidity are shown. Additionally, profiles of the corresponding refractivity and the refractivity gradients are shown. The profiles are made along rectangles at azimuth angles of 10 and 130 degrees, respectively, and the five curves in each profile correspond to positions at distances of 10, 50, 100, 150, and 200 km from the radar (see Figure 5.6 for the location of the profiles). It is seen from the plots that the refractivity gradients have much larger negative values (down to almost -400 N-units/km) at the rectangle positioned at an azimuth angle of 130 degrees than for the rectangle with an azimuth angle of 10 degrees (around -40 N-units/km). This is due to a temperature inversion over the Baltic Sea where dry warm air overlays cool humid air over the sea. Normally, the temperature decreases with altitude (as seen in Figures 5.9), however, a temperature inversion is the phenomenon seen in Figure 5.10, where the temperature increases with altitude (here up to a height of 100 m). The temperature inversion causes high gradients in the refractivity profile.

The high gradients seen at the azimuth angle of 130 degrees cause superrefraction which was also seen in the simulated field in Figure 5.8. To illustrate the spatial distribution of sub and superrefracted areas the vertical refractivity gradient between the two lowest NWP model levels is shown in different colours in Figure 5.6. Superrefracted areas are shown in red hues and subrefracted areas are shown in blue hues. The area east of the radar site is dominated by superrefraction while land areas show subrefraction conditions.

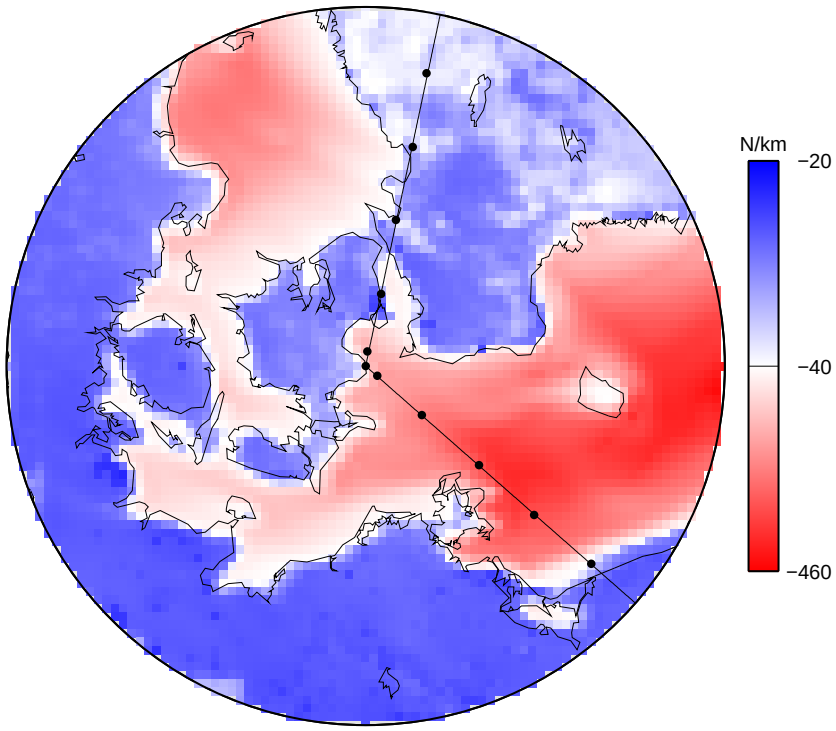

Figure 5.6: The refractivity gradient between the lowest model levels calculated from the NWP model. The rectangles positioned at an azimuth angle of 10 degrees and 130 degrees is also indicated. The black dots indicate the positions of the profiles in Figure 5.9 and Figure 5.10.

From the 2D propagation computation a 3D view of the propagation of the radar wave was made, see Figure 5.11. The field is here shown on three rectangles together with a projection of the average of the lower 50 m of the field onto the earth surface. The projection of the field is also shown in Figure 5.12, where it can be seen that the amplitude of the field has large values in the lower right corner of the plot, corresponding to areas of sea clutter.

Comparing the simulated near-surface field (Figure 5.12) with the actual measured radar data (Figure 5.5), a very good agreement is seen. Clutter is both simulated and observed in an extended area in the Baltic Sea. Local variations in propagation are also captured in the model. This can be seen by the change from non-standard to standard propagation when going from sea to land; also both the simulated field and the radar image have as anticipated low field values over the island of Bornholm found on the right side of the plots.

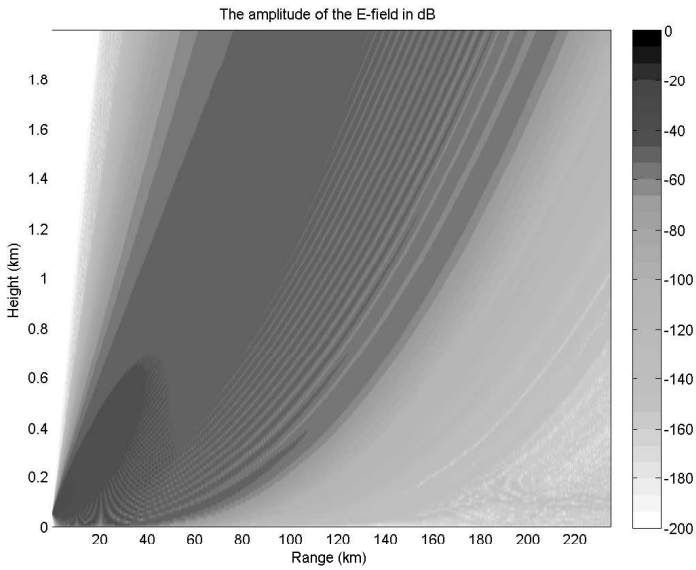

Figure 5.7: The amplitude of the electromagnetic field for the rectangle placed at an azimuth angle equal to 10 degrees.

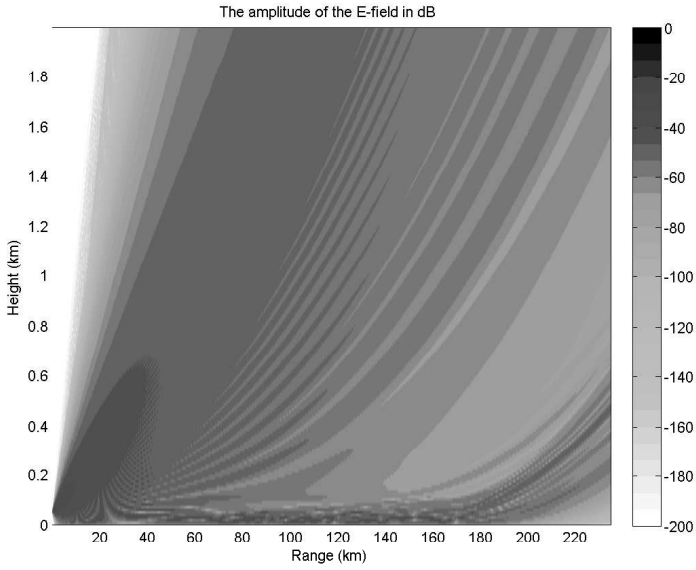

Figure 5.8: The amplitude of the electromagnetic field for the rectangle placed at an azimuth angle equal to 130 degrees.

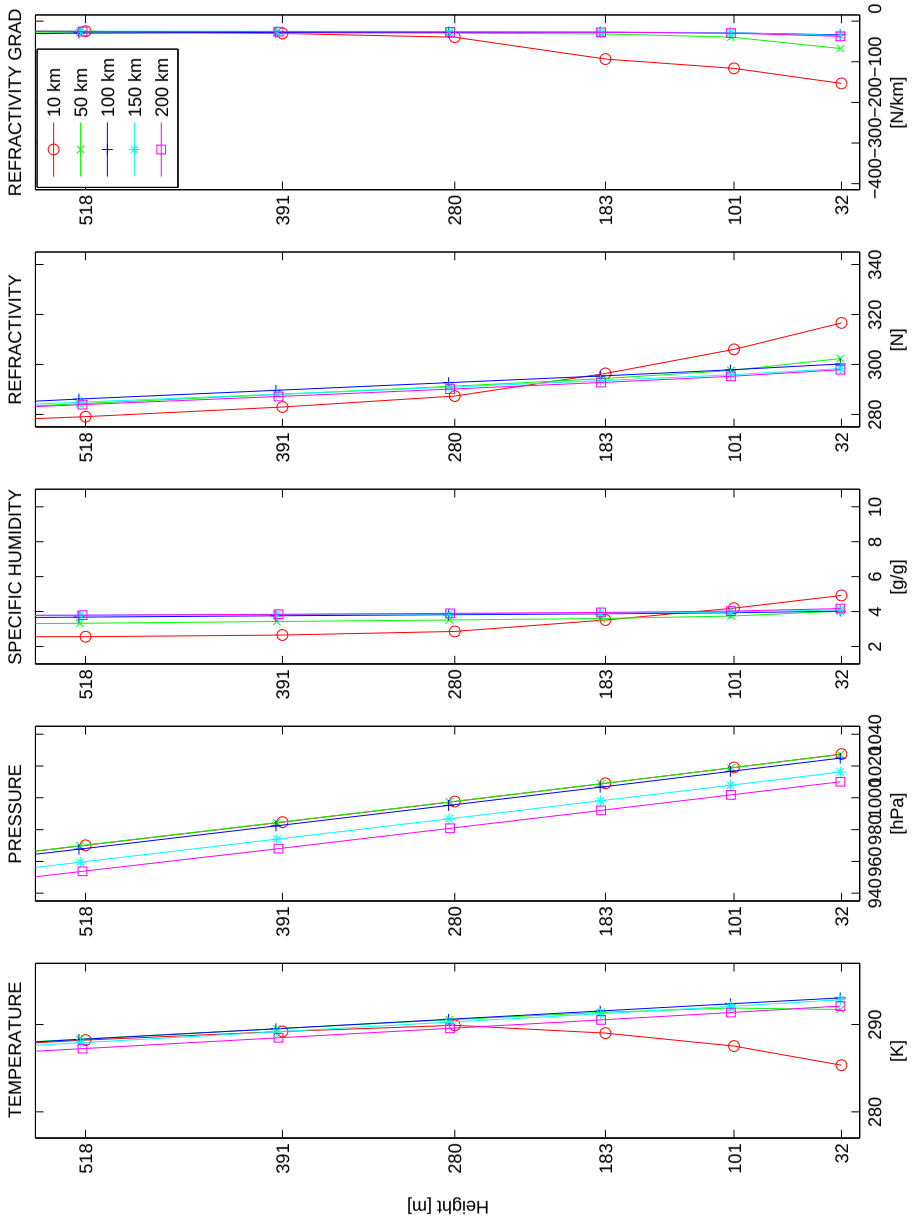

Figure 5.9: Vertical profiles of temperature, pressure, humidity, refractivity and refractivity gradient at five positions along the earth surface (as seen in Figure 5.6). The values correspond to the rectangle placed at an azimuth angle of 10 degrees.

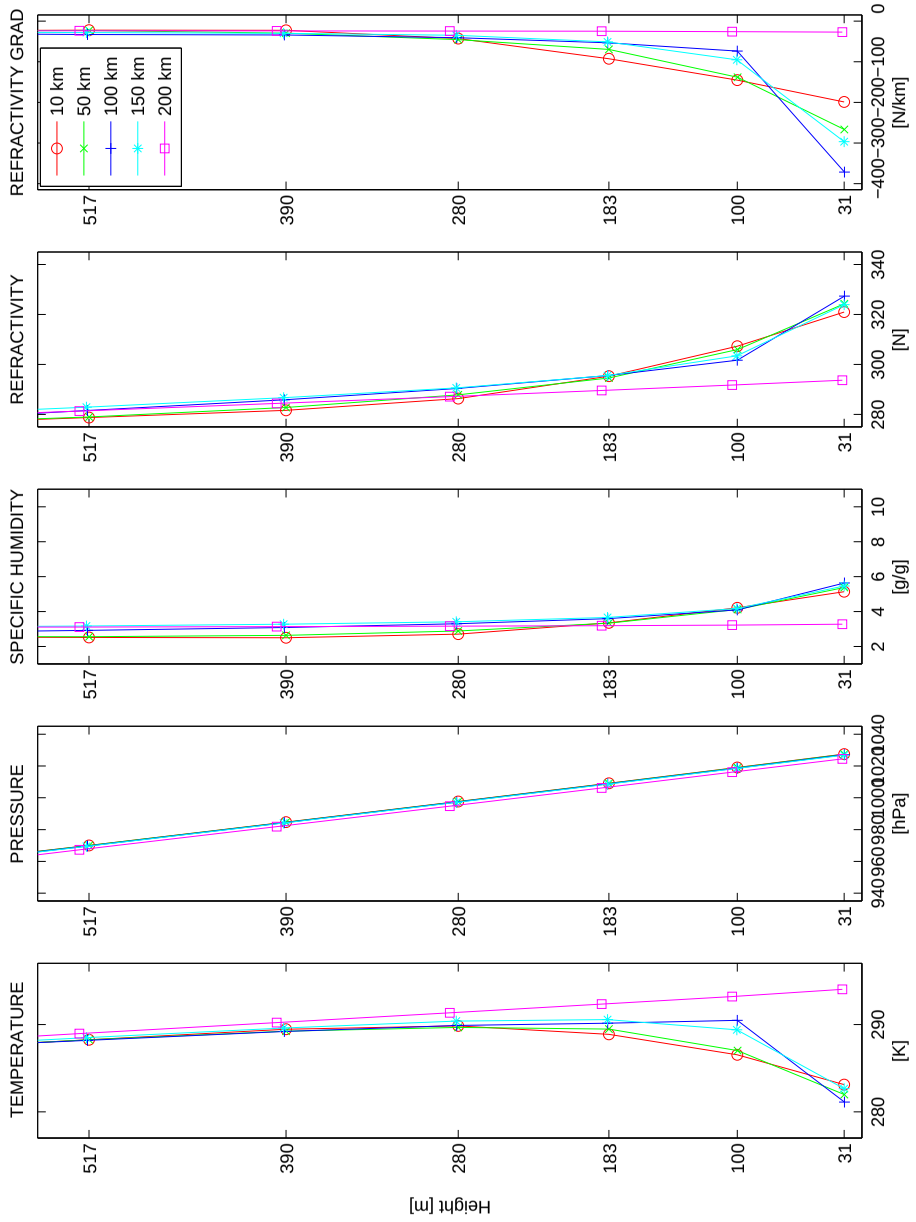

Figure 5.10: Vertical profiles of temperature, pressure, humidity, refractivity and refractivity gradient at five positions along the earth surface (as seen in Figure 5.6). The values correspond to the rectangle placed at an azimuth angle of 130 degrees.

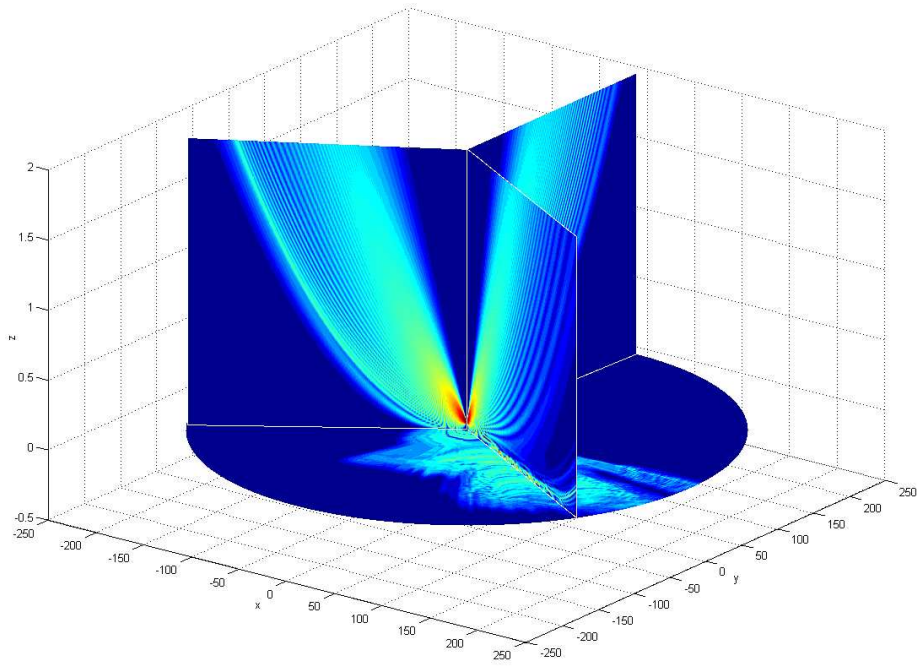

Figure 5.11: The amplitude of the electromagnetic field in a 3D view.

It should be noted that the wave propagation simulation is based on predicted values of the refractive index. It is, however, anticipated that there will always be some difference between the predicted and the true values of the modified refractive index, due to the very dynamic behaviour of the anomalous propagation phenomenon. When comparing Figure 5.12 with Figure 5.5 it can be noted that the wave propagation simulation predicts a higher level of surface reflections close to the antenna position than what is seen in the radar image. The reason to this lies in the fact that the ground clutter filter in the radar's signal processor removes ground clutter from stationary targets which have zero Doppler velocities (unlike sea clutter).

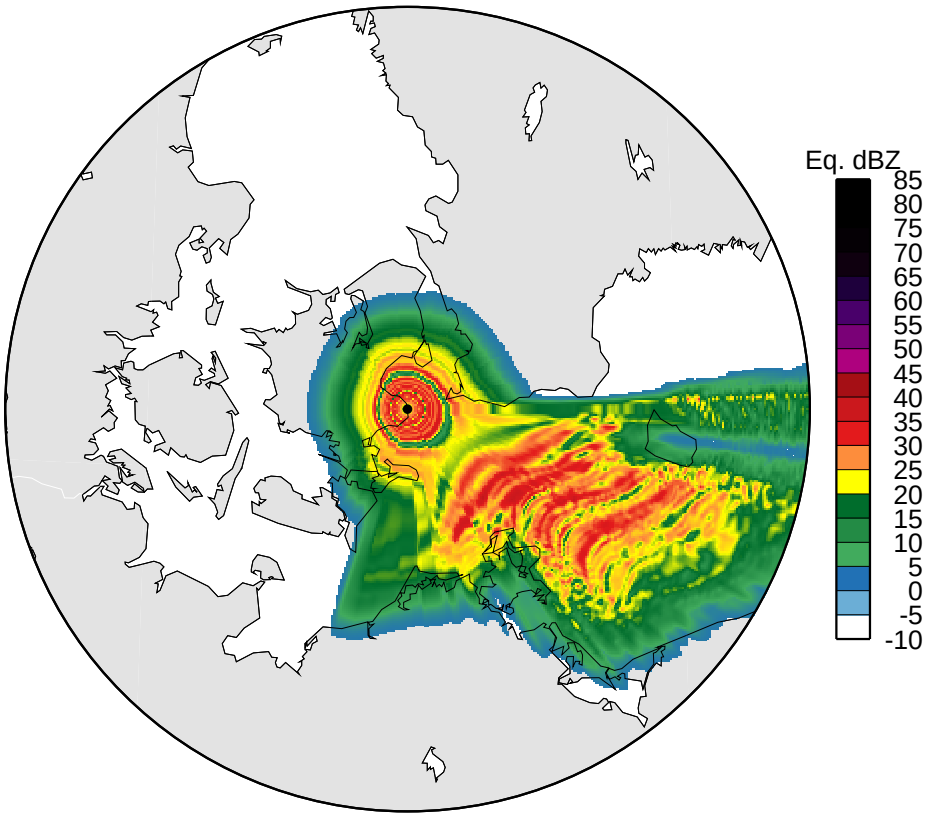

Figure 5.12: The amplitude of the electromagnetic field (converted to equivalent dBZ values) along the earth surface in a 2D view. The radar is positioned in the centre of this plot.

5.4.3 Case example B: 2005-09-25 20:00 UTC

The radar simulations were also been performed on a day when clutter as well as precipitation was present in the observed radar data. Figure 5.13 shows the radar image from September 25, 2005 at 20:00 UTC. A precipitation front can be seen in the left side of the image while moderate to strong sea clutter is present in the right side of the image. Multispectral satellite images from Meteosat-8 confirmed again that the coherent area to the left of the radar image is precipitation while the area to the right is radar clutter.

The result of running the wave propagator through the predicted refractivity is seen in Figure 5.14. The ducting phenomenon to the east of the radar in the Baltic Sea as well as the normal propagation area to the west of the radar are well captured.

The correspondence is, however, not perfect. The spatial distribution of areas of strong clutter and areas of high predicted clutter differs to some extent. This is due to the fact that the radar wave propagation is based on inputs of refractivity fields from a numerical weather model and it can not be expected that predictions from these models will be able to precisely model the very dynamic phenomena of atmospheric refractivity.

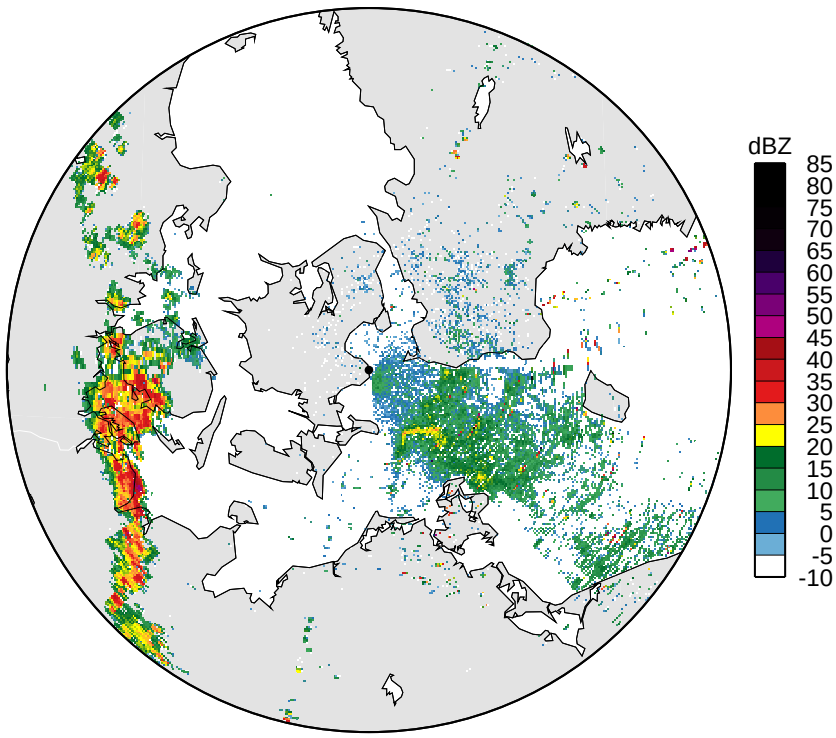

Figure 5.13: Stevns radar image from a scan performed September 25, 2005 at 20 UTC time. In the western part of the image a precipitation front is seen, whereas the backscatter in the eastern sector is solely clutter. Note that clutter from interfering radio transmitters at azimuth angles of approximately 25, 80, and 190 degrees is also present in the image.

5.4.4 Detection and removal of radar clutter

One of the immediate applications of the radar propagation simulation is to identify areas of clutter and facilitate the removal of clutter echoes.

A simple classification of radar echoes into clutter and precipitation classes was done by using a threshold value on the simulated near-surface EM energy and use this to mask out radar pixels in areas of simulated clutter.

The result of applying this method can be seen in Figure 5.15 which shows the radar image from Figure 5.13, but with the area where the wave propagation simulation predicts radar clutter masked out. It can be seen that the radar image contains a significantly lower amount of clutter (there is still some land clutter present over land areas) and therefore the data quality has been increased considerably.

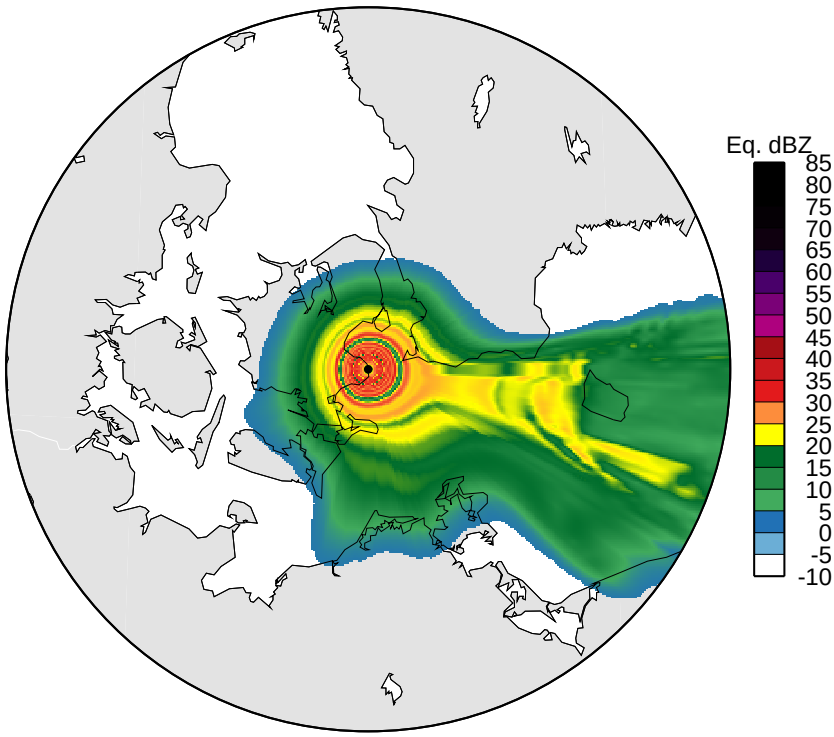

Figure 5.14: The amplitude of the electromagnetic field along the earth surface in a 2D view. The radar is positioned in the centre of this plot. This simulation corresponds to the radar scan performed September 25, 2005 at 20 UTC time.

The classification accuracy was computed to be 94.5 % with a false alarm ratio of 0 (i.e. no precipitation was misclassified as clutter). The probability of detection was 0.92 which indicates that 8 % of the clutter was not detected in the classification. For comparison, in Case Example A the classification was carried out with 94.7 % accuracy, a false alarm ratio of 0, and a probability of detection of 0.95.

The wave propagator technique will of course be challenged in cases of areas of mixed clutter and precipitation. The radar echoes from the ground and from the air will be mixed. However, as is described in the next section, the propagation model allows for simulation of propagation of increasing elevation angles, which makes it possible to determine which elevations are expected to be affected with ground clutter. It should however be noted that precipitation and clutter tend to be spatially separated which make the method more robust.

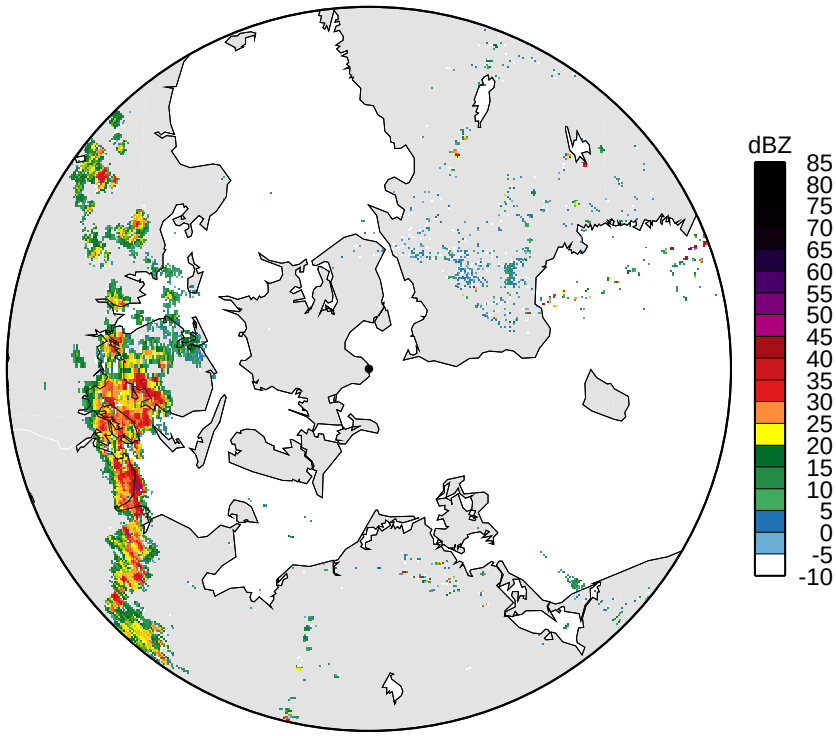

Figure 5.15: Radar image with predicted clutter echoes removed.

5.4.5 Elevation angle dependence on clutter

Superrefraction of radar waves occur only at small grazing angles (small elevation angles). This is due to the fact that the rapid change in refractivity happens in the vertical plane and not in the horizontal plane. Realistic meteorological values of refraction and successive application of Snell's law show that ducting is limited to elevation angles below one degree. This is supported by the radar observations which show that clutter is not present at higher elevation angles. At lower elevations, superrefraction is more frequent and at increasing elevation angles superrefraction decreases. Furthermore, it is the atmospheric conditions close to the radar site which determine whether non-standard propagation occurs. If AP conditions are only present at a far range from the radar, the radar waves have already escaped the near-surface altitudes where non-standard propagation and ducting occurs.

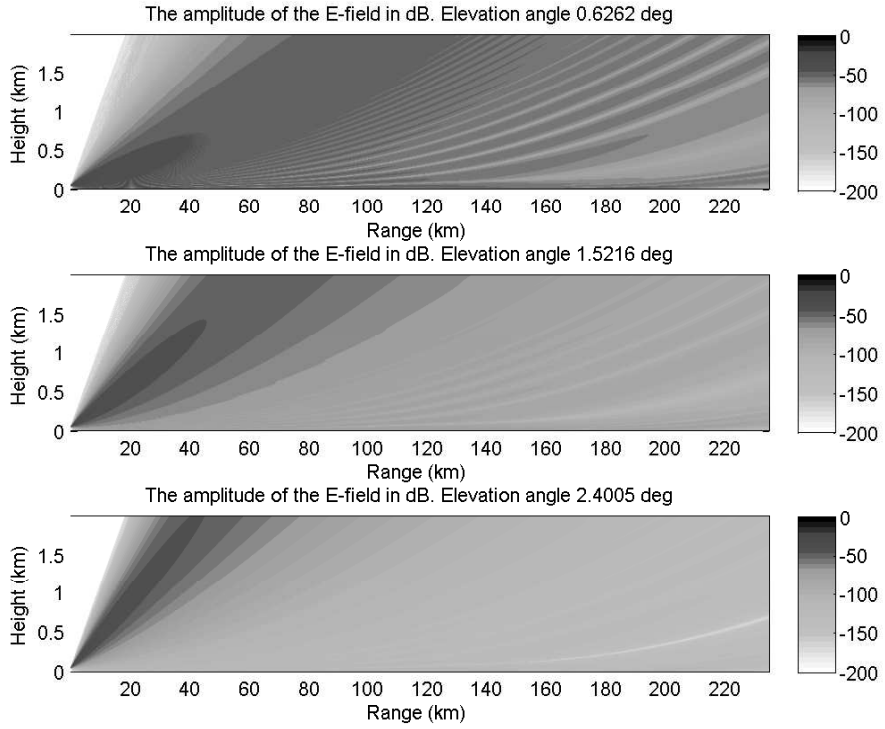

Figure 5.16: The amplitude of the radar field for three different elevation angles. The azimuth angle is 130 degrees.

To show the potential of the wave propagator, the simulation was run for increasing elevation angles 0.63, 1.52 and 2.40 degrees (Figure 5.16). These elevation angles correspond to the first, fourth, and fifth elevation in the operational scanning strategy of the DMI radars. The figure shows propagation through the strong duct at azimuth angle 130 degrees as described earlier.

It can be seen that for this case example, superrefraction and clutter is predicted to occur only in the lowest two elevation angles than in the third elevation, where normal propagation is expected. Using this information, identification of precipitation and radar clutter echoes can be carried out on an elevation basis and clutter can be removed from the lowest elevation angles while potential precipitation in the higher elevations can be preserved. This is at the expense of a potentially limited detection of precipitation at far ranges due to overshooting of precipitation at higher elevation angles.

5.4.6 Summary and conclusions on wave propagation method

A method for simulation and prediction of weather radar clutter caused by anomalous propagation using a wave propagator has been presented. The method uses the parabolic equation approximation to the wave equation to propagate the radar waves through a 3D refractivity field from a NWP model. It has been shown in two case examples that radar clutter caused by temperature inversion in the atmosphere and reflection of the radar wave from the sea can be predicted with the use of a wave propagation simulation. Very good agreement is seen when the simulated near-surface field is compared to the actual observed radar data. Clutter was in the two examples both simulated and observed in an extended area in the Baltic Sea.

The case examples show that wave propagation simulations can be used to identify areas of potential ground clutter. This facilitates the improvement of the quality of radar images by removal of radar echoes in areas under suspicion of anomalous propagation. Likewise, it was show how precipitation echoes in areas of standard propagation can be preserved, while clutter was removed. From inspection of the model results and the computed classification performance (accuracies of 94.5 % and 94.7 %, false alarm ratios of 0, and probabilities of detections of 5 % and 8 %) it can be seen that the sea clutter phenomenon is reasonably well captured, although not with the same level of details as the observed radar clutter.

Improvements of the propagation model include increased horizontal and especially vertical resolution of the NWP grid as well as more accurate input data, i.e., water vapor fields.

The method for radar clutter prediction has the potential of being used in operational meteorological services; however, it would however mean that simulation time should be decreased from the present level of hours on a standard office PC to minutes. The performed wave propagation calculations can be performed in parallel and it is therefore expected that the calculation time can be reduced to minutes on a multiprocessor supercomputer.

5.5 Method II: Azimuth-dependent propagation index

The main limitation on the wave propagation method in its current form is the computation time. Therefore the possibility to predict anomalous propagation from the refractivity gradient alone was investigated. This lead to the development of the azimuth-dependent propagation index (API) which for each azimuth angle expresses the degree of expected anomalous propagation in that direction.

5.5.1 Azimuth-dependent propagation index

Data

As input to the azimuth-dependent propagation index can be used any measure of propagation conditions. Here the vertical refractivity gradient and the ducting index are used. As described earlier the VRG is the refractivity gradient between surface level and 1 km height. As this measure does not take into account the potentially high gradients lower than 1 km, also the mean and minimum VRG values between surface level and 1 km were computed. These are denoted VRGmean1000 and VRGmin1000 in the following, and the VRG is denoted VRG1000.

The propagation features were computed for all times in the radar database for use in computation of the API. In Figures 5.17 to 5.20 are shown the VRG1000, VRGmean1000, VRGmin1000, and I_D for the three case examples used in Chapter 3 and 4. The color map for the figures were chosen to highlight areas of superrefraction conditions (red hues) and standard or subrefraction (blue hues).

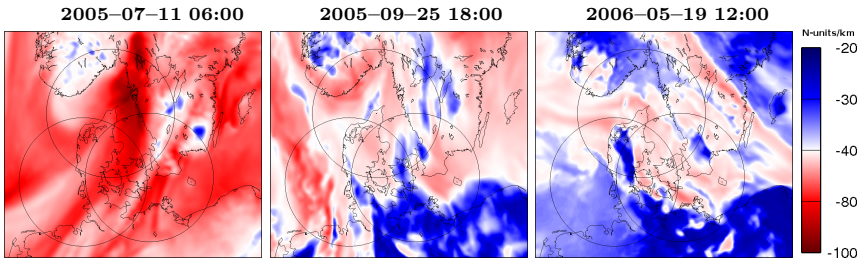

Figure 5.17: VRG1000 for the three cases.

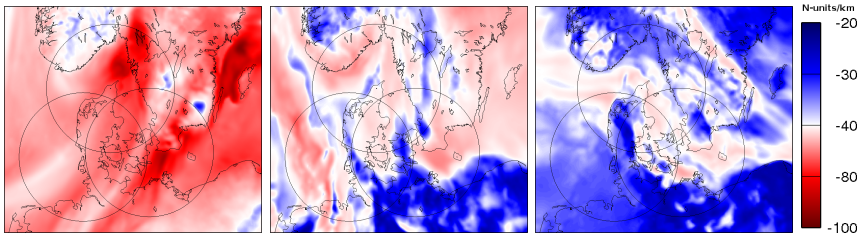

Figure 5.18: VRGmean1000 for the three cases.

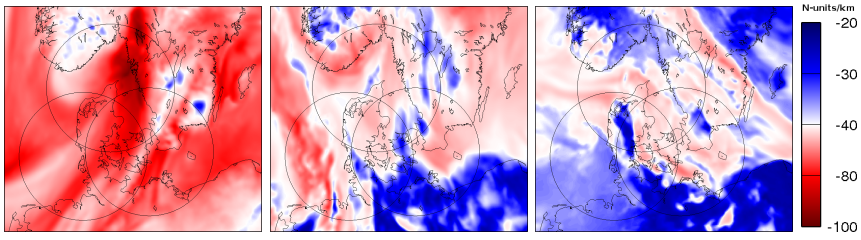

Figure 5.19: VRGmin1000 for the three cases.

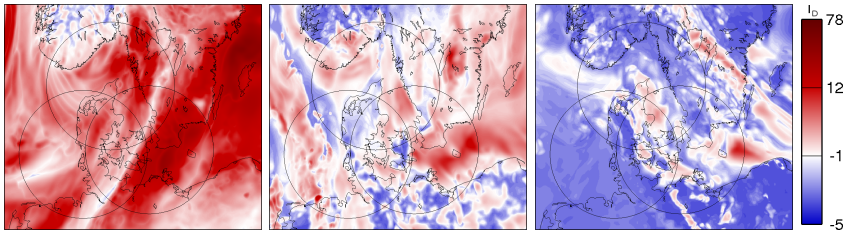

Figure 5.20: I_D for the three cases.

Azimuth-dependent propagation index

The azimuth-dependent propagation index is the distance weighted average of the propagation measure along each azimuth angle from the radar site (See Figure 5.21), here derived with the Ducting Index as input parameter:

$$API_{I_D} = \frac{\sum \frac{I_D}{r}}{R}, \tag{5.5}$$

where I_D is the Ducting Index value, r is the distance from the radar site, and R is the maximum range. In other words, as the distance to the radar increases the weight of the propagation measure (here I_D) decreases. Using a distance weighted average was chosen because it is the propagation conditions close to the radar site that has the highest influence on the risk of anomalous propagation. For example, if close to the radar site, standard propagation is present, then the radar waves will escape the part of the atmosphere where anomalous propagation can occur. On the contrary, if anomalous propagation conditions are present near the radar site, then ducting can occur which will affect the rest of the echoes along that azimuth angle. The distance weighting is done by linear weighting but other weight functions can be applied.

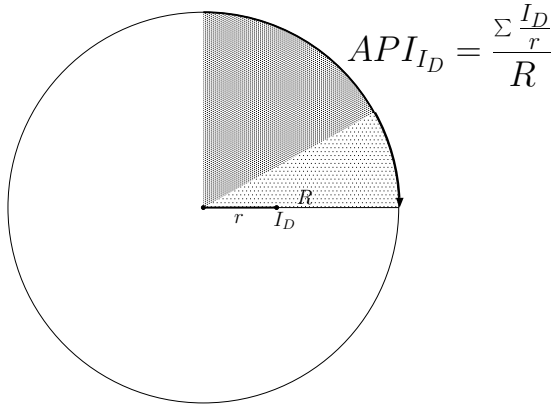

Figure 5.21: The API (azimuth-dependent propagation index) is computed as the distance weighted average of the propagation measure (here I_D) along each azimuth angle (here for azimuth 90 degrees)

The API was computed for each radar site and then mosaiced together. The maximum value is chosen for the API for the Ducting Index, while the minimum value is chosen for the VRG features. This of course can lead to clutter being removed from the other radar's observations, and ideally, the API should be

applied to the radar data before it is mosaiced.

In Figure 5.22 the resulting API feature is shown for the Ducting Index. The color map is again selected to highlight areas of potential clutter (red) and precipitation (blue) The values for the color map was selected by inspection of the feature histograms for the API, which are described in the next section. It can be seen that for *Case I* there is a risk of anomalous propagation for most of the scene, except for a sector West of the Rømø radar. Comparison with Figure 5.20, left image, reveals that the ducting index is close to -1 West of the Rømø radar, i.e., close to standard propagation. *Case II* predicts well the anomalous propagation conditions from the Stevns radar which produced sea clutter in the Baltic Sea and Figure 5.20, middle image shows that the Ducting Index predicts quite strong ducting in the Baltic Sea at that time. For the rest of the image the API_{I_D} correctly predicts standard propagation. Finally, *Case II* shows standard propagation conditions for the whole scene.

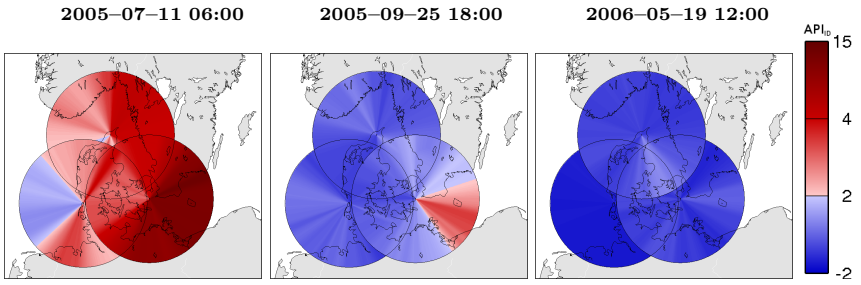

Figure 5.22: The azimuth-dependent propagation index, API_{I_D} for the Ducting Index I_D.

5.5.2 Feature histograms

The feature histograms of the API_{I_D} (Figure 5.23) show that precipitation and clutter are well separated. The clutter class has low values with a mean value of approximately 0 and the precipitation class has higher values with a mean of about 4. As described above, the Ducting Index indicates anomalous propagation for values above -1 and standard propagation at values below -1. These values are not in line with the distribution of the classes in the feature histogram which could indicate that the Ducting Index computed from the HIRLAM NWP is underestimated.

Looking at the distribution of the API_{I_D} for three classes, a good separation of all three classes, land and sea clutter, and precipitation, is also seen. Land clutter takes API values between precipitation and sea clutter.

In Appendix B, Figures B.21 and B.22, and Appendix C, Figures C.21 and C.22 the feature histograms and conditional probabilities can be seen for the API_{ID}, $\text{API}_{VRG1000}$, and $\text{API}_{VRGmean1000}$, for all echo strengths and two and three classes.

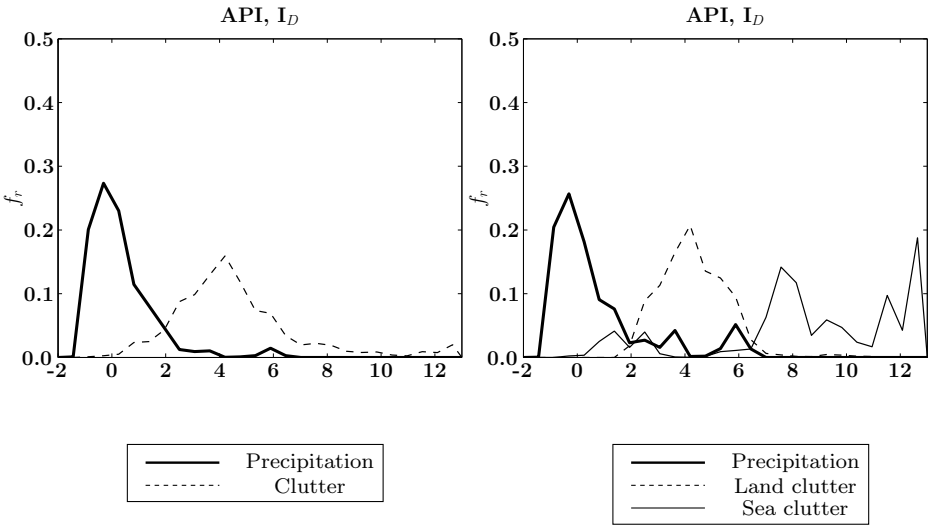

Figure 5.23: Feature histrograms for the API_{I_D} for clutter and precipitation classes. Two classes (left) and three classes (right).

5.5.3 Feature selection

Even if the API measures, API_{ID}, $\text{API}_{VRG1000}$, $\text{API}_{VRGmean1000}$, $\text{API}_{VRGmin1000}$ are all derived from the same refraction data, they were subjected to feature selection to investigate if they contain complimentary information.

The feature selection results are shown in Appendix D, Tables D.25 to D.30.

From the tables it can be concluded that under the assumption of a quadratic classifier, the $\text{API}_{VRGmean1000}$ feature is the best feature, if only one was to be chosen. The best two features are the $\text{API}_{VRGmean1000}$ in combination with the $\text{API}_{VRG1000}$. For three features, the API_{I_D} is added. The finally chosen subsets can be seen in the appendix.

5.5.4 Classification evaluation

The results from evaluation of the ensemble classification of the full training data set can be seen in Tables 5.2 and 5.2.

In general, good classification accuracy is achieved with the best majority vote classifier for all echoes being for three class classification with a PC value of 93.95 % (Table 5.3). The majority classifier consistently scores better than the single worst classifier (the LDA for two classes, and the QDA for three classes) and it scores consistently second to the best classifier (the k-nearest neighbor).

For the false alarm ratio, the best value for all echoes classified by the majority classifier is 0.021 (two classes). The best HKS value for all echoes is 0.886 for three classes.

Echo types

Again, it is seen how classification of strong echoes in general scores better than weak echoes and all echoes. For example, the best PC for the majority vote classifier using three classes is achieved for strong radar echoes with 95.14 % vs. 92.88 % for weak echoes and 93.95 % for all echoes.

Number of classes

Contrary to the data fusion method and the spatio-temporal method, here there is a significant improvement when employing three class classification over using two class classification (improvements up to 1.5 % PC (for all echoes, using the majority vote classifier)). This indicates that there is information in the propagation features which is dependent on clutter type, i.e., whether it is land or sea clutter.

Table 5.2: Classification results. Propagation method, two classes.

Echo	Percent correct (PC)						False alarm ratio (FAR)						Hanssen-Kuipers skill (HKS)					
type	LDA	QDA	Dtree	KNN	SVM	Majority	LDA	QDA	Dtree	KNN	SVM	Majority	LDA	QDA	Dtree	KNN	SVM	Majority
All	87.35	90.33	92.28	94.12	89.58	92.49	0.017	0.046	0.055	0.041	0.043	0.021	0.773	0.816	0.847	0.884	0.804	0.862
Weak	89.76	90.84	94.80	96.12	91.19	92.55	0.026	0.033	0.033	0.018	0.034	0.022	0.821	0.833	0.895	0.927	0.838	0.867
Strong	89.22	93.02	94.62	96.03	92.32	94.58	0.005	0.036	0.043	0.028	0.031	0.005	0.781	0.859	0.892	0.920	0.845	0.890

Table 5.3: Classification results. Propagation method, three classes.

Echo	Percent correct (PC)						False alarm ratio (FAR)						Hanssen-Kuipers skill (HKS)					
type	LDA	QDA	Dtree	KNN	SVM	Majority	LDA	QDA	Dtree	KNN	SVM	Majority	LDA	QDA	Dtree	KNN	SVM	Majority
All	93.09	91.92	94.14	96.00	91.99	93.95	0.021	0.046	0.057	0.029	0.045	0.025	0.872	0.844	0.879	0.921	0.845	0.886
Weak	91.41	91.35	94.13	95.95	91.47	92.88	0.029	0.042	0.049	0.026	0.037	0.033	0.845	0.836	0.874	0.918	0.840	0.866
Strong	92.89	92.79	93.53	95.20	92.79	95.14	0.017	0.055	0.075	0.052	0.045	0.010	0.855	0.855	0.871	0.904	0.854	0.901

5.5.5 Results

The azimuth-dependent propagation index was applied to the three case examples to see the performance on the case of extensive land and sea clutter (*Case I*), the case of mixed clutter and precipitation (*Case II*), and the case of precipitation only (*Case III*).

For *Case I* the results are very good with a percent correct score of 99.1 %. From the API_{ID} feature (Figure 5.22, left) it can be seen that all radars were operating in superrefraction conditions leading to anomalous propagation and API clutter. It is also seen how standard propagation is predicted over the North Sea which almost made it possible to classify the small precipitation as precipitation. It was, however, misclassified.

In *Case II* the method classifies part of the sea clutter in the Baltic Sea correctly. Again, comparing with Figure 5.22, middle image, it is seen how a sector of the coverage of the Stevns radar is subject to superrefraction. Only part of the sea clutter is classified as clutter, and thus a large proportion of missed clutter results in a classification accuracy of 91.4 %. This is a around 4 % worse than the performance of the wave propagator using data a little later on the same day, but for a single radar's data only.

The final case, *Case III*, is classified with the excellent classification accuracy of 100.0 %. The whole region takes API values indicating standard propagation (Figure 5.22, right) and no risk of anomalous propagation clutter.

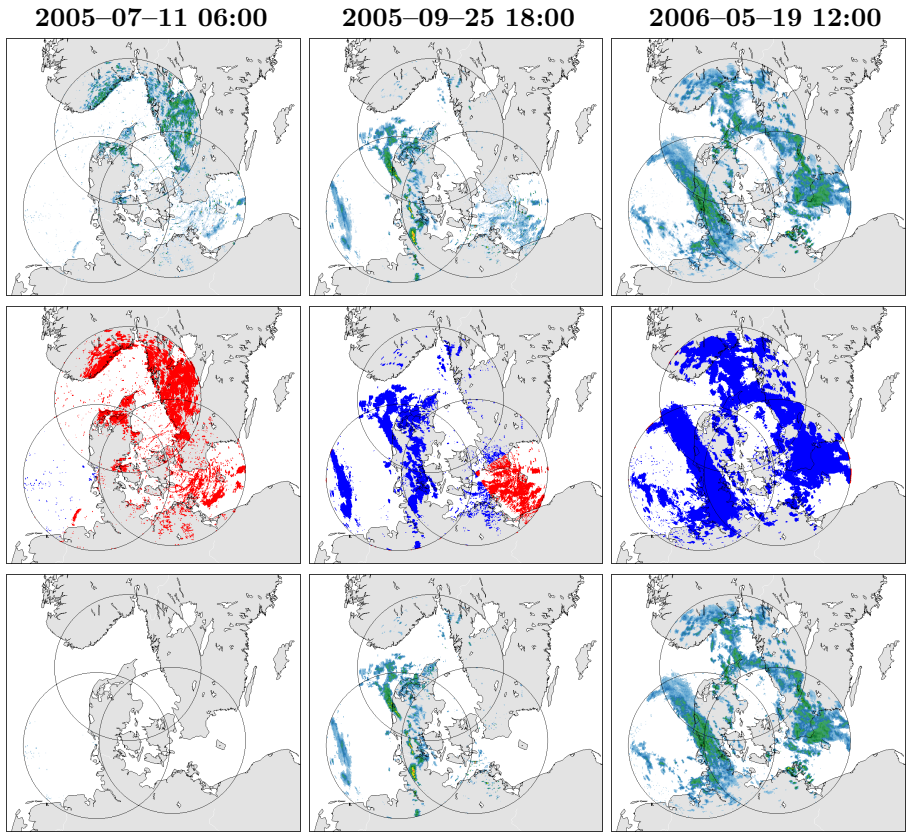

99.1 %

	C_c	P_c
C_t	39703	0
P_t	377	0

91.4 %

	C_c	P_c
C_t	12439	4331
P_t	0	33306

100.0 %

	C_c	P_c
C_t	0	0
P_t	22	137044

Figure 5.24: Classification results using the azimuth-dependent propagation index on the three case examples. Top row: Radar images, Middle row: Classified images (blue is precipitation and red is clutter), Bottom row: Filtered radar images. The classifications accuracy and confusion matrix are found below (Subscript t means *truth* and subscript c means *classification*, cf. Figure 3.27)

5.5.6 Summary and conclusion on the API method

A method for detecting clutter using information about radar propagation from a numerical weather prediction model was presented.

From the temperature, pressure, and humidity extracted from DMI-HIRLAM-S05 the refractivity was extracted. From this, various propagation measures were computed, amongst other the vertical refractivity gradient (VRG) and the ducting index (I_D). These measures were used as input to a novel azimuth-dependent propagation index, which is the distance weighted average of the VRG or the I_D on an azimuth basis. Thus the API facilitates clutter detection on an azimuth angle basis.

Good classification performance was obtained using an ensemble classifier combining the results of five individual classifiers. The skill score percent correct of 93.95 % was obtained for all echos using three class classification. The best false alarm rate achieved was 0.021 (two classes) and the best HKS value for all echoes was 0.886 for three classes.

The API method is simple to implement and computationally inexpensive, which encourages the use as an operational clutter detection technique.

The performance of the method can probably be improved on by optimization of the API weighting. In the current version it uses a linearly decaying weight with distance. Future work, could include investigations of the impact of using different weight functions.

Another improvement to the method would be to apply it on a per radar basis, i.e., before making the mosaic of radar images. This would most likely lower the number of false alarms in comparison to the method used in this study (to apply the API classification on the mosaiced images).

5.6 General summary and conclusions

In this chapter, two methods were presented which use predictions of the radar propagation to classify radar echoes into clutter and precipitation classes.

Both methods used information on the refractive index of the atmosphere from the HIRLAM numerical weather prediction model.

The first method used a wave propagator to simulate the propagation of electromagnetic energy through a varying refractivity field. The method was able to identify areas of anomalous propagation leading to sea clutter. At the same time the method also predicted areas of standard propagation well, which made it possible to remove a great amount of the clutter in two case examples. Classification accuracies between 94.5 % and 94.7 % were achieved.

The second method used as input the gradient of the refractivity which can be used to detect different propagation conditions. A simple azimuth-dependent propagation index was proposed which expresses for each azimuth angle the degree of anomalous propagation. An accuracy of 93.95 % was obtained.

The two methods, used quite different methods for estimating the propagation conditions, but achieved results that were quite similar. In the operational use of propagation methods for clutter detection, a hybrid methodology could be employed. If the azimuth-dependent propagation index predicts anomalous propagation in a given sector of a radar's coverage, then the wave propagator could subsequently be used for generating a more detailed simulation of the propagation. This would reduce the amount of computations and make the method faster.

Chapter 6

General summary and conclusion

6.1 Summary

In this thesis, three methods for detection of clutter in weather radar data were proposed. One method used data fusion with multispectral satellite data, another used the spatio-temporal information in the optical flow in radar images combined with wind information from a weather model, and the final method used the refractivity information from a weather model to predict radar propagation and anomalous propagation clutter using a wave propagator and an azimuth-dependent propagation index.

The methods, thus, approached the problem from different angles, but parts of all methods used the same supervised classification framework to detect clutter in the data. 60 radar images from 15 days of typical events of clutter and precipitation were annotated and a large training dataset was established, consisting of more than one million samples for the three classes, land clutter, sea clutter, and precipitation.

For each method, the best sets of features were selected using subset selection. Supervised classification was performed using an ensemble of five different classifiers: Linear and quadratic discriminant analysis, a decision tree, k-nearest

neighbor classification, and a support vector machine. These were combined into the ensemble classifier using the majority vote between the results of each individual classifier. The amount of sufficient training data for achieving good accuracy of classification was estimated to be 4000 samples per class which was concluded from the construction of learning curves.

Table 6.1: Classification evaluation of the clutter detection methods. Percent correct (PC), false alarm ratio (FAR), and Hanssen-Kuipers Skill (KHS) for all echoes.

Method	PC	FAR	HKS
Data fusion, MSG, day	95.82	0.054	0.919
Data fusion, MSG, night	95.28	0.013	0.908
Data fusion, Precip. Clouds, day	94.81	0.069	0.898
Data fusion, Precip. Clouds, night	93.46	0.044	0.790
Spatio-temporal, pixel-based	79.36	0.176	0.582
Spatio-temporal, object-based	96.22	0.031	0.924
Azimuth-dep. propagation index	93.95	0.021	0.886

The classifiers were evaluated by training the models using 4000 samples per class and using the remaining of the training data as independent test data. The performance of the methods in classification of clutter was assessed using the skill scores, percent correct (PC), the false alarm ratio (FAR), and the Hanssen-Kuipers skill score (HKS). In Table 6.1, the overall performances of the developed methods are summarized. The skill scores are shown for classification of all echoes (i.e., both weak and strong echoes) and using two class classification (except for the azimuth-dependent propagation index for which the best performance was for using three classes in the classification). Note that different classification methods were used for the 'Precipitating Clouds' data and the wave propagation method, as described in the respective chapters.

In the following, the main results are summarized for each of the three methods.

The **data fusion method** was split into two parts. One using the 11 Meteosat Second Generation (MSG) bands as input, the other using the 'Precipitating Clouds' product which is derived from the MSG data to express the probability of precipitation. The classification accuracies for both methods were good with a percent correct value for the use of MSG data of around 95 % and of around 94 % for the use of the 'Precipitating Clouds' data. The false alarm ratios were higher for application of the methods using day features (visual/near-infrared, and infrared bands), 5.4 % and 6.9 %, in comparison with night features (only

infrared bands) with FAR values of 1.3 % and 4.4 %. The Hanssen-Kuipers skill score highlights the best result as the use of MSG day features with a value of 0.919.

Application of the data fusion method on three case examples showed good classification accuracies between 95.0 % and 98.5 % with good detection of clutter. The error rates of up to 5 % were due to misclassification on the edges of precipitation areas which were thought to be caused by misalignment issues related to the use of multiple data sources from multiple sensors.

The **spatio-temporal method** used the optical flow between two radar images to identify areas with fast or slow changes in the direction and the speed of the motion of radar echoes. The entropy was used as a measure of the local diversity of the flow. For the entropy of the flow direction, it was found that clutter takes high entropy values while precipitation takes lower values. This corresponds to fast and slow changing directions, respectively. For the entropy of the flow speed, land clutter takes low values while precipitation and sea clutter both take high values. To strengthen the classification, the wind field information from a numerical weather prediction model was used to compute a difference feature in the form of the difference between the optical flow field and the NWP field. As would be expected, high differences were found for areas of clutter and low differences for precipitation areas.

The entropy and difference features were computed on a pixel-basis and an object-basis. Pixel-based features were computed using a moving window, while the object-based features were computed by averaging the values for each connected component in the image. The evaluation of the spatio-temporal method showed mediocre results for the pixel-based method with a PC value of 79.36 %, a high FAR of 17.6 % and a mediocre HKS of 0.582. The use of object-based features, however, provided a significant improvement to all skill scores. The PC value achieved was 96.22 % (the highest of all the methods), a false alarm ratio of 3.1 % and a HKS of value 0.924. The same improvement using object-based features was seen when the method was applied to the case examples. The maximum achieved PC value was increased from 83.5 % to 99.6 %.

The **propagation approach** to detection of clutter was divided into two different parts. Both, however, used refractivity data extracted for the DMI-HIRLAM-S05 numerical weather prediction model. The first method showed the potential of using a wave propagator for simulation of anomalous propagation clutter on the DMI radar at Stevns. The wave propagator successfully identified areas of sea clutter in the Baltic Sea by simulation of the superrefraction conditions over the ocean while preserving precipitation echoes in areas of standard propagation. The wave propagator was applied to two case examples only and the classification results were very good with percent correct values

between 94.5 % and 94.7 %.

The second propagation method used as input the vertical gradient of the refractivity which can be used to identify different propagation conditions. An azimuth-dependent propagation index (API) was proposed, which, for each azimuth angle of each radar site, contains the distance weighted average of the vertical refractivity gradient or the ducting index. The API was found to capture well the propagation conditions on an azimuthal basis and a classification accuracy (percent correct value) of 93.95 % was obtained. The false alarm ratio was low (2.1 %) and the Hanssen-Kuipsers skill score at a value of 0.886. The three case examples were classified with accuracies up from 91.4 % to 100 %.

Ensemble classifier

The use of an ensemble classifier was found to improve the performance of classification of radar echoes. Consistently, the ensemble classifier performed better than the single worst classifier of the ensemble. For example, for the data fusion method, the percent correct skill was improved by 0.78 % up to 12.1 % in comparison to the single worst classifier. Using the ensemble classifier, the classification accuracy was, in most cases, at the same time lowered in comparison to the single best classifier, but since it is unknown in advance which classifier would perform the best, an ensemble classifier was found to be an advantage. Outside the improvement of the classification accuracy, using an ensemble classifier made it possible to compare the classifiers and find the best one for each method.

Number of classes

Classification was carried out for all three methods using two configurations of classes. One configuration used three class classification using the classes: land and sea clutter and precipitation. The other one, used two class classification with the two clutter classes pooled into a common clutter class. Only the propagation method showed a better performance by using three classes. The data fusion and spatio-temporal methods both showed best detection using only one clutter class for the classification.

Echo strengths

Summarizing the performance of clutter detection methods for strong and weak radar echoes, it was found that better classification performance, in general, was obtained for strong echoes in comparison to weak echoes.

Reject option

The use of a reject option to control the amount of false alarms and misses in the classification was showcased on the data fusion method. By omitting classification of input vectors for which the maximum posterior probability of the classes is lower than a certain threshold, the number of false alarms can be decreased (at the expense of a higher number of misses). The reject option, thus, is a knob that can be used to balance the number of misclassified clutter (misses) against the number of misclassified precipitation (false alarms).

6.2 Conclusion

The aim of this work was to improve on the data quality of weather radar data by developing new methods for clutter detection. Three methods were developed, tested, and applied, and the results show that all three methods are capable of detecting clutter with high accuracy. An average classification accuracy of 93 % was achieved, and if the worst method is left out (the pixel-based spatio-temporal method) an average accuracy of 95 % is obtained.

Of course, classification errors are always unwanted—and even if being rather low, an average error rate of 5 % can indeed cause great uncertainty in the subsequent application of "clutter filtered" weather radar data. However, it can be argued that in many applications of weather radar data, the detection of a great majority of the clutter at the expense of some remaining clutter and some mistakenly removed precipitation, would result in an overall increase of the data quality.

The contributions of this work is a step forward in the continued work towards improved methods for clutter detection. The three proposed methods were developed with the aim of being applicable as practical methods in the operational use of radar data. A common framework was used for all three methods, which allowed for the intercomparison of the methods. The same database of radar data and training data were used, and the same procedures for classification and evaluation of the methods were applied. Furthermore, the sample size for the training and testing sets is thought sufficiently large to make it possible to be able to make general statements about the expected performance of the method when applied to new data.

Regarding the data fusion method, this contribution provides new results of using the high-resolution data from Meteosat Second Generation and the derived 'Precipitation Clouds' product on an large training and test set. The results of

the thesis underlined the preliminary results achieved in earlier versions of the method.

The novel method of using an optical flow technique for clutter detection is also part of this work's original contribution. The method provided disappointing results if used on a pixel-basis but very good results were obtained when used in an object-based approach.

The propagation approach firstly showed the potential of using a wave propagator for clutter detection and secondly the evaluation of a azimuth-dependent propagation index showed that high accuracy clutter detection is possible using a simple measure of the propagation conditions of the weather radar.

6.3 Future work

Several improvements to the presented methods can be made. Starting with the data fusion method, further research into mitigation of the misalignment of multi-source datasets could provide an increased performance. The correction for parallax displacement of the satellite data is one aspect and the other is the extrapolation of satellite data to match the radar data. The further exploration of a scale-space approach to mitigate border effects between datasets is also of interest. Finally, the tuning of the parameters of the 'Precipitating Clouds' product with radar data is thought to provide an immediate improvement of detection. Ideally, the tuning should be made with radar data from the radar network that the method will be applied to.

The spatio-temporal method could be improved on by optimization of the optical flow features. Especially, work towards improving the skill of the pixel-based features is needed and would also benefit the object-based features.

Regarding the propagation methods, the wave propagator should be applied to a larger test set before operational use. Furthermore, the computation of the wave propagation needs optimization for computational speed before real-time use is realistic. The propagation index method could be improved on by applying it to the data on a per radar basis rather than applying it to a radar mosaic image.

In general, one can never get too much training and test data for a supervised classification method, so naturally, the collection of more training data from a wider range of precipitation events (snow and hail) and more cases of clutter is part of the future operationalization of the methods.

Finally, a study of ways to combine the above clutter detection methods into a single integrated clutter detection system could be carried out. It is expected that the use of several different methods in parallel would improve the detection of clutter since each method has its own individual strengths and weaknesses.

Appendix A

Radar image database

Radar images

In this appendix are shown the radar images used for development and testing of the methods for clutter detection. Images from all the days described in Table 2.3 are shown.

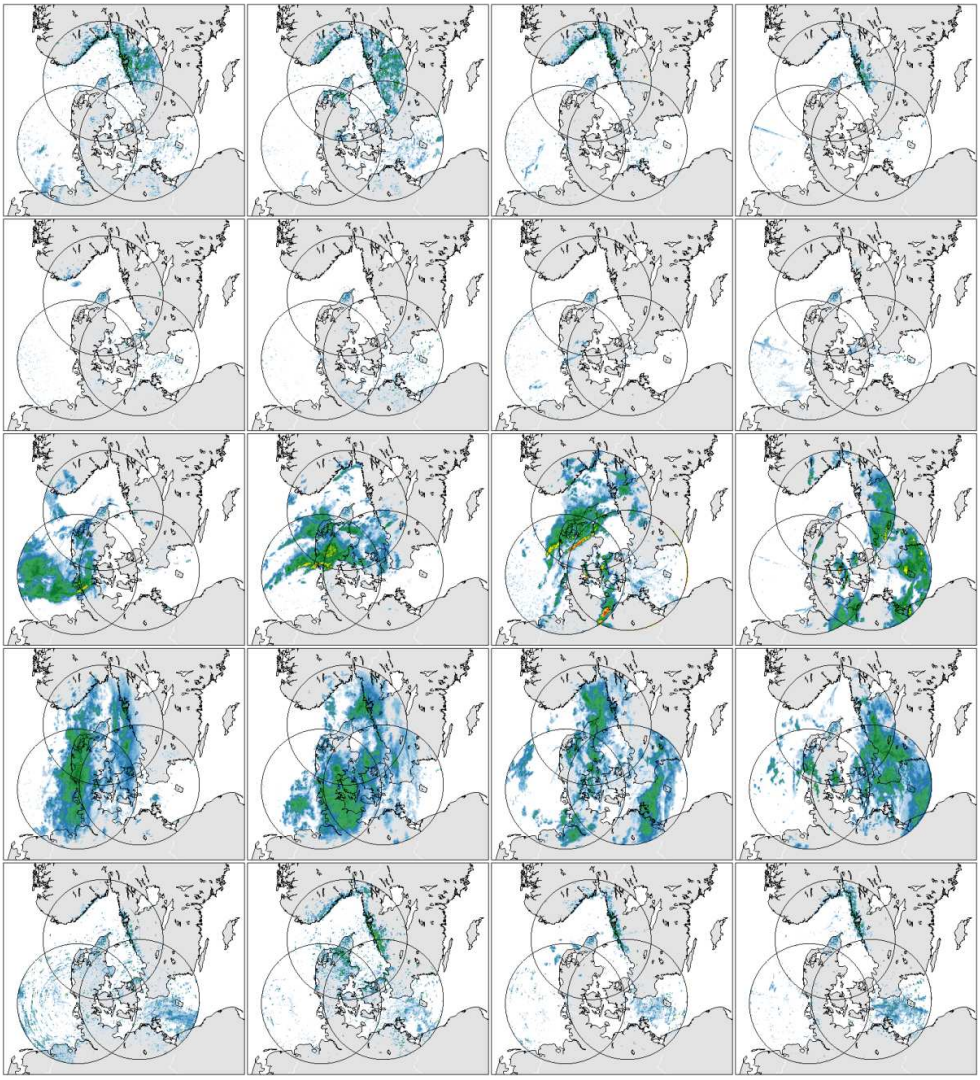

Figure A.1: Radar image database. Columns: 01:00, 06:00, 12:00, 18:00 UTC. Rows: 2005–07–11, 2005–07–13, 2005–07–15, 2005–08–25, 2005–09–06.

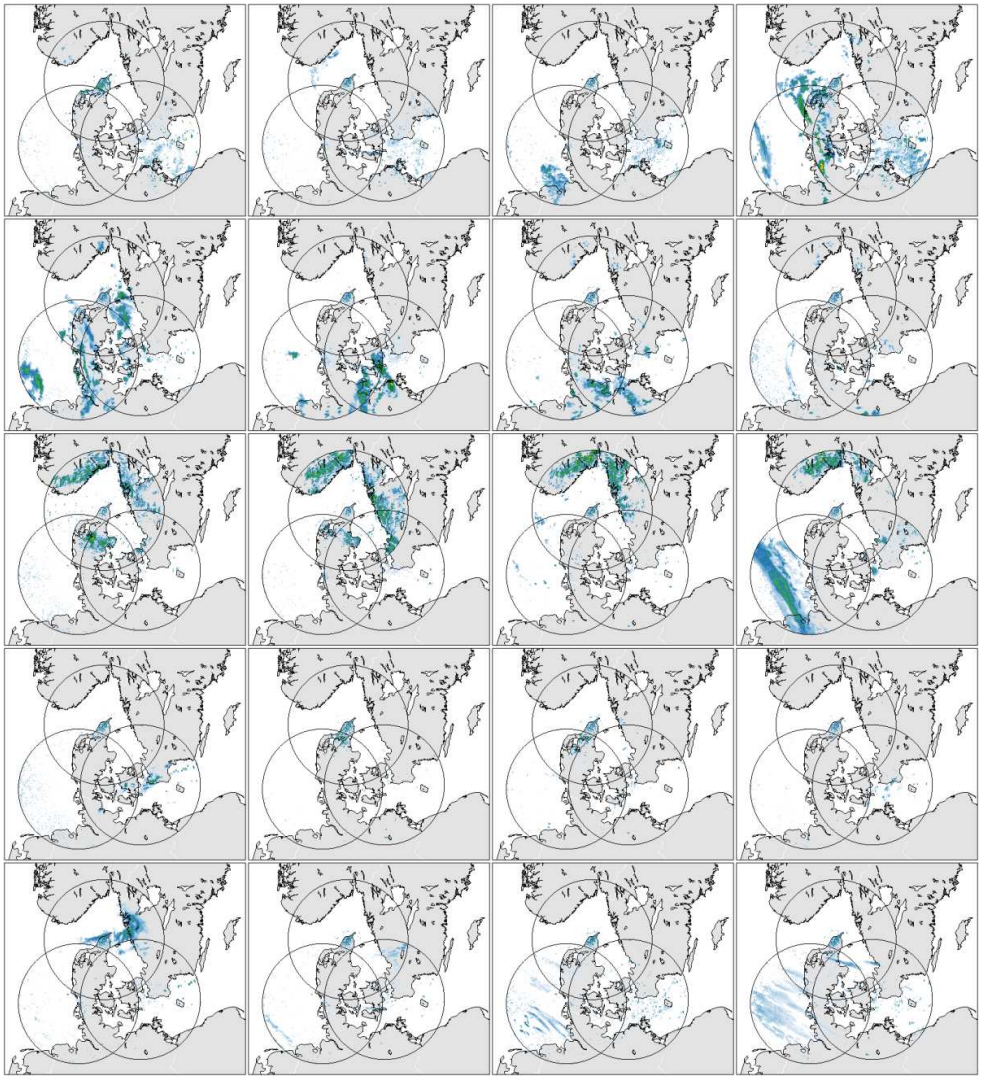

Figure A.2: Radar image database. Columns: 01:00, 06:00, 12:00, 18:00 UTC. Rows: 2005–09–25, 2005–09–26, 2005–10–19, 2006–01–02, 2006–03–24.

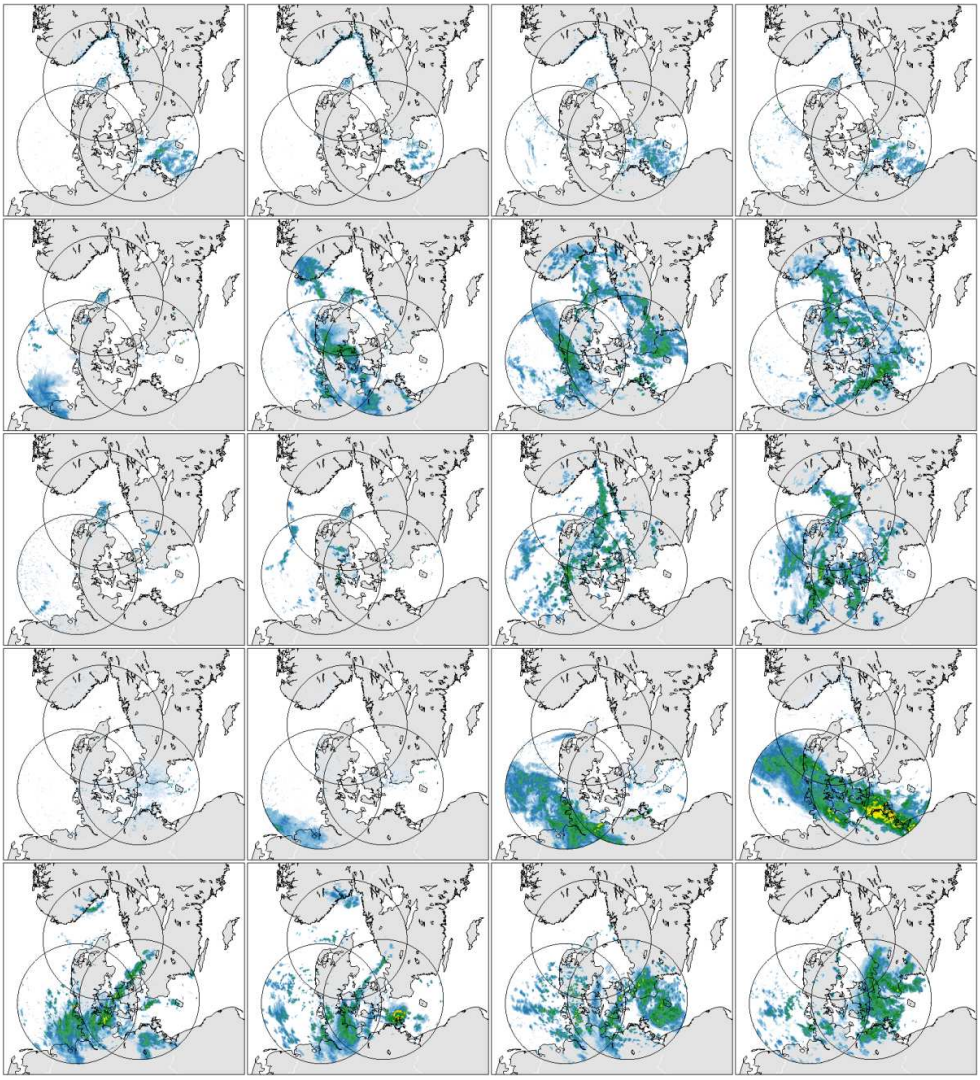

Figure A.3: Radar image database. Columns: 01:00, 06:00, 12:00, 18:00 UTC. Rows: 2006–05–05, 2006–05–19, 2006–05–24, 2006–08–18, 2006–08–29.

Appendix B

Feature histograms

Feature histograms

This appendix contains features histograms (relative frequency polygons) for the features which are used in the methods.

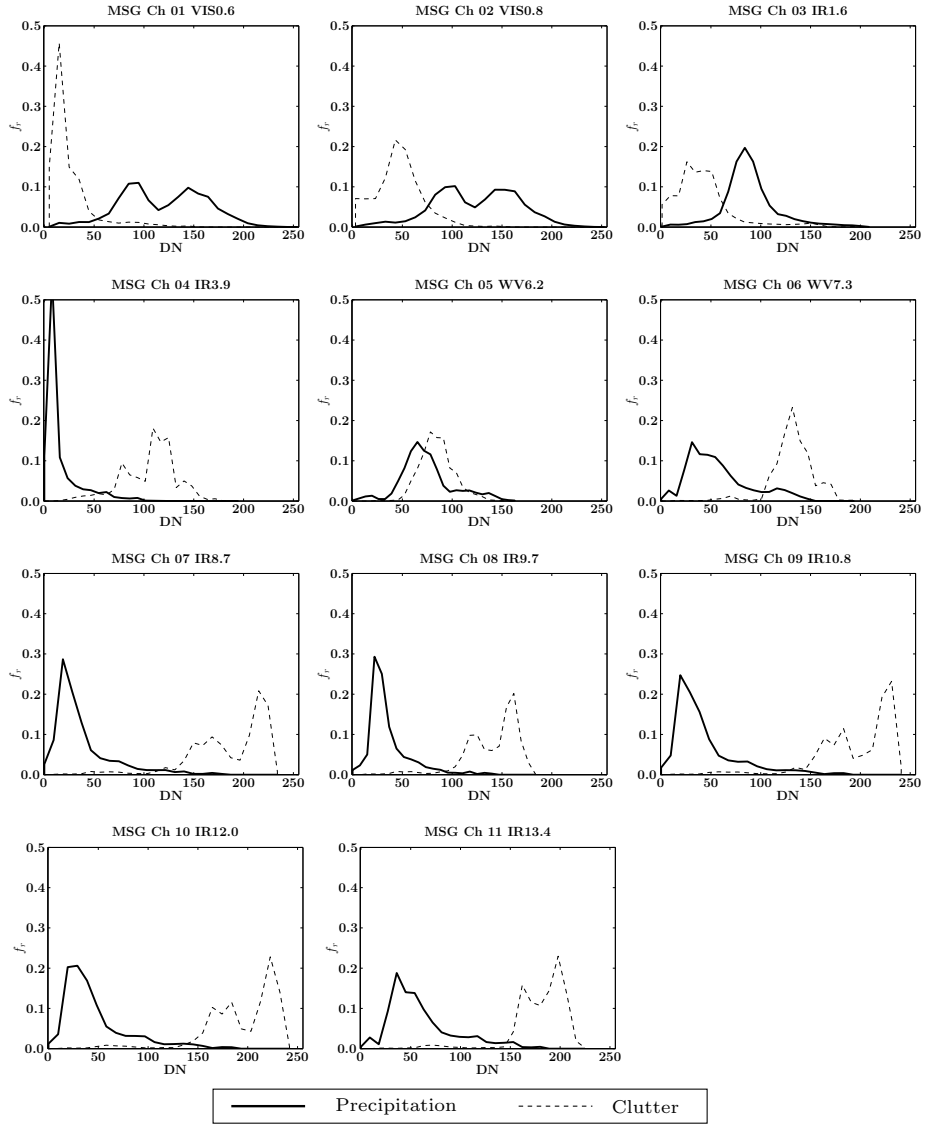

Figure B.1: Data Fusion features. Day features and all echoes. Two classes.

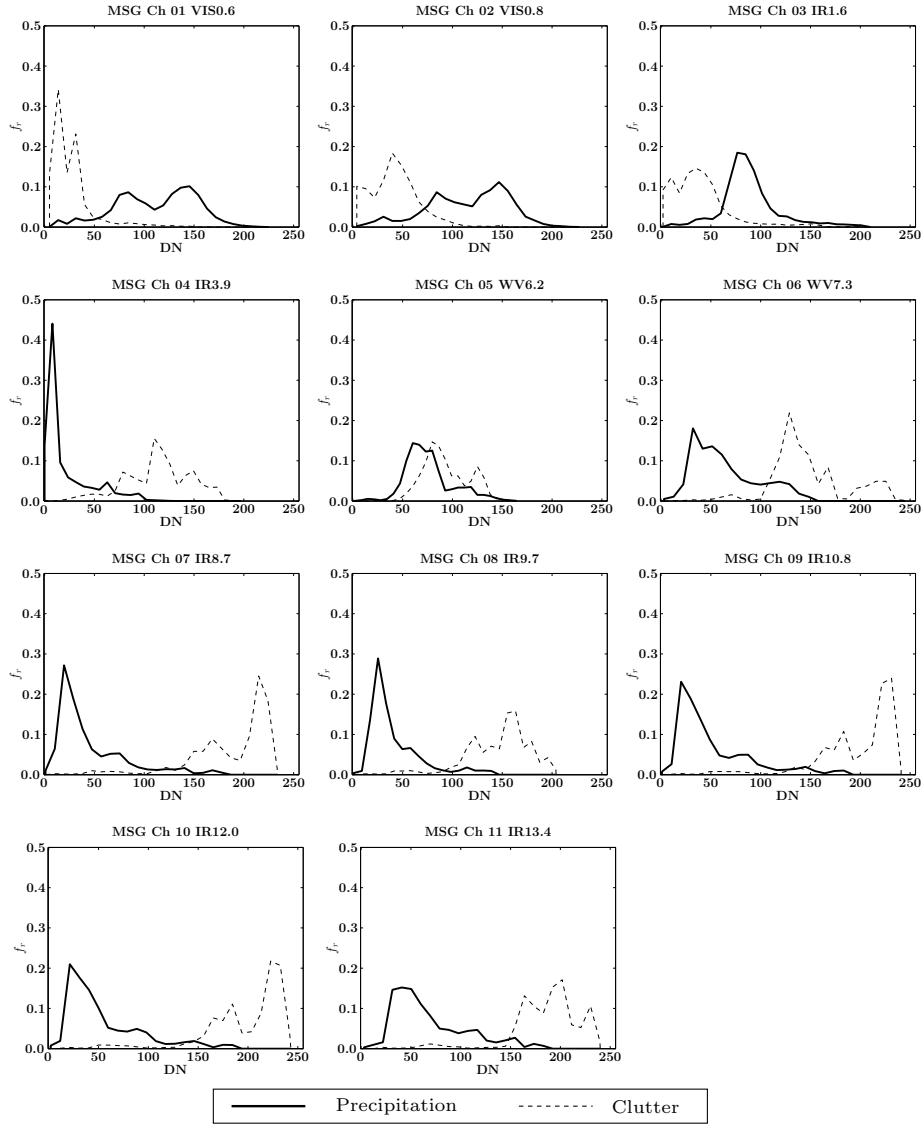

Figure B.2: Data Fusion features. Day features and weak echoes. Two classes.

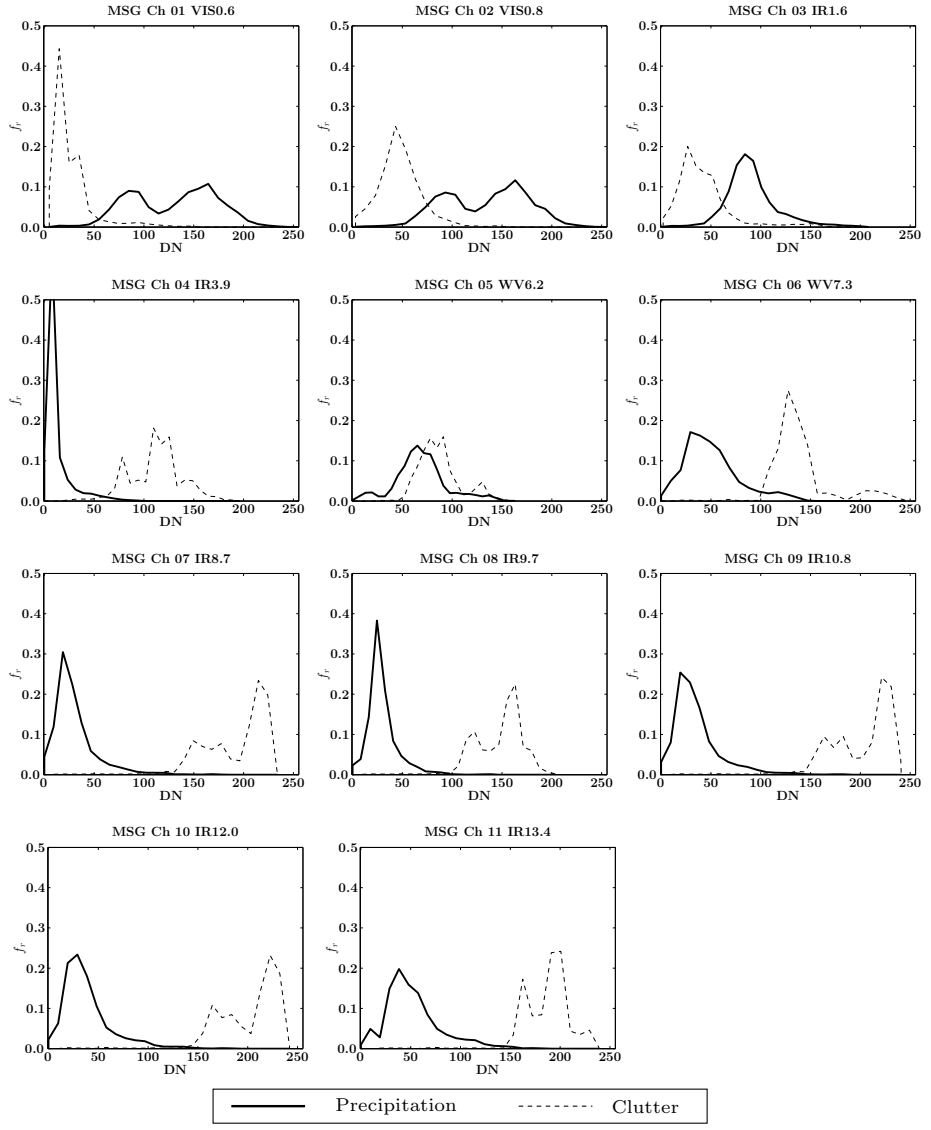

Figure B.3: Data Fusion features. Day features and strong echoes. Two classes.

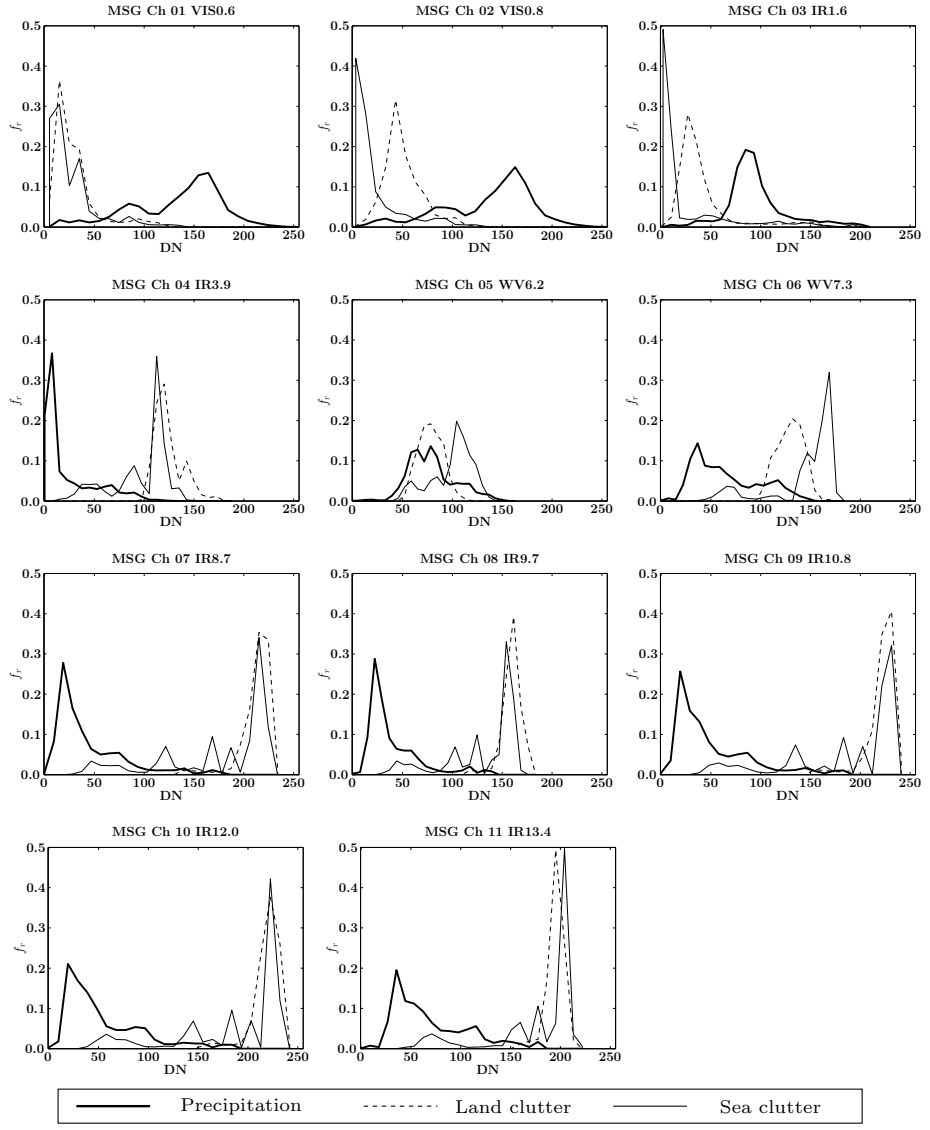

Figure B.4: Data Fusion features. Day features and all echoes. Three classes.

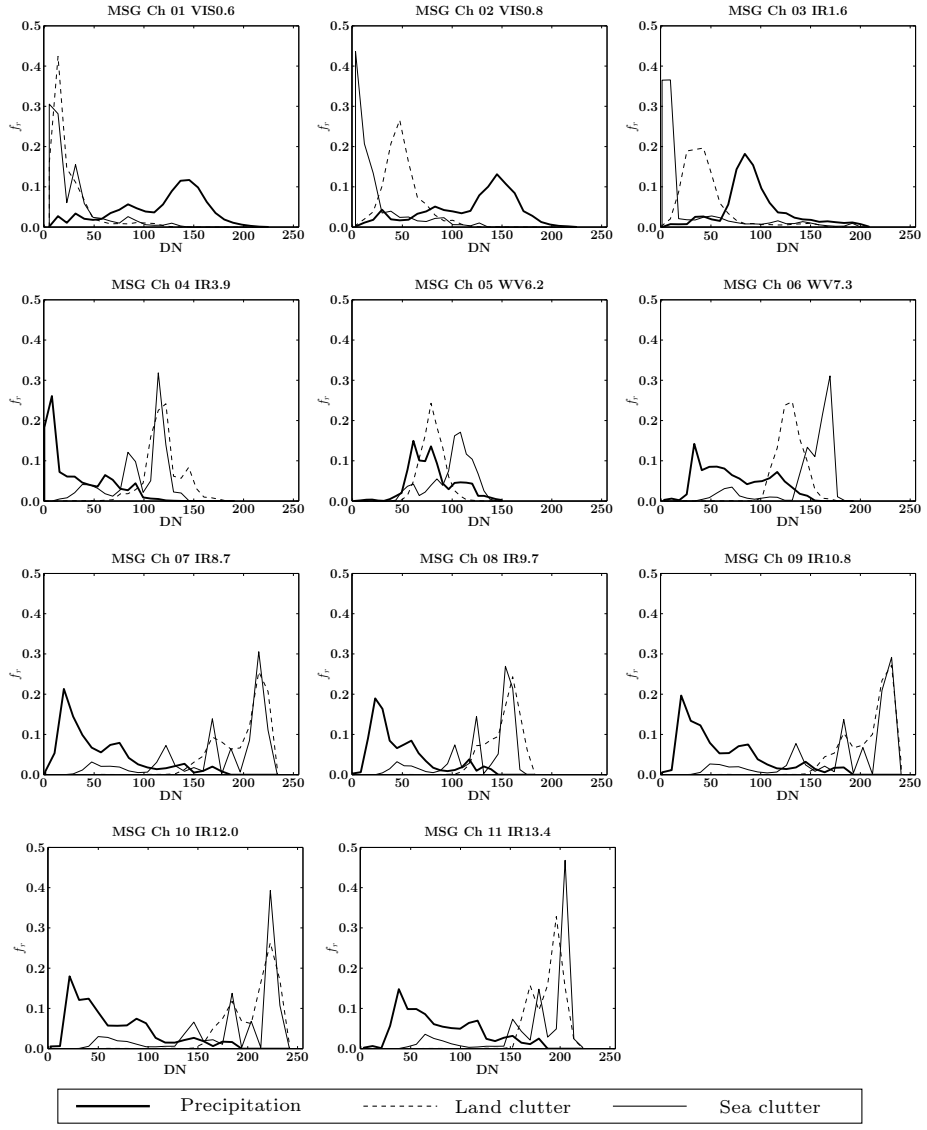

Figure B.5: Data Fusion features. Day features and weak echoes. Three classes.

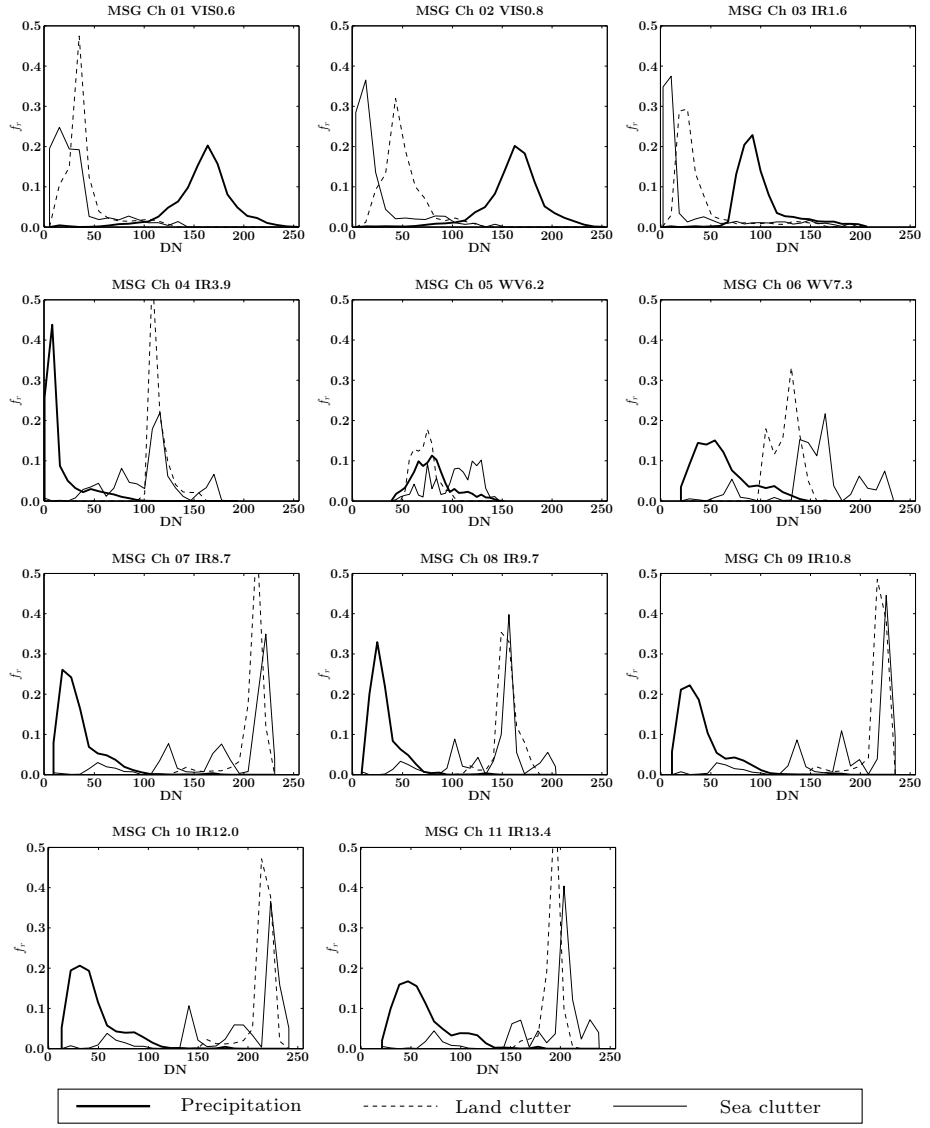

Figure B.6: Data Fusion features. Day features and strong echoes. Three classes.

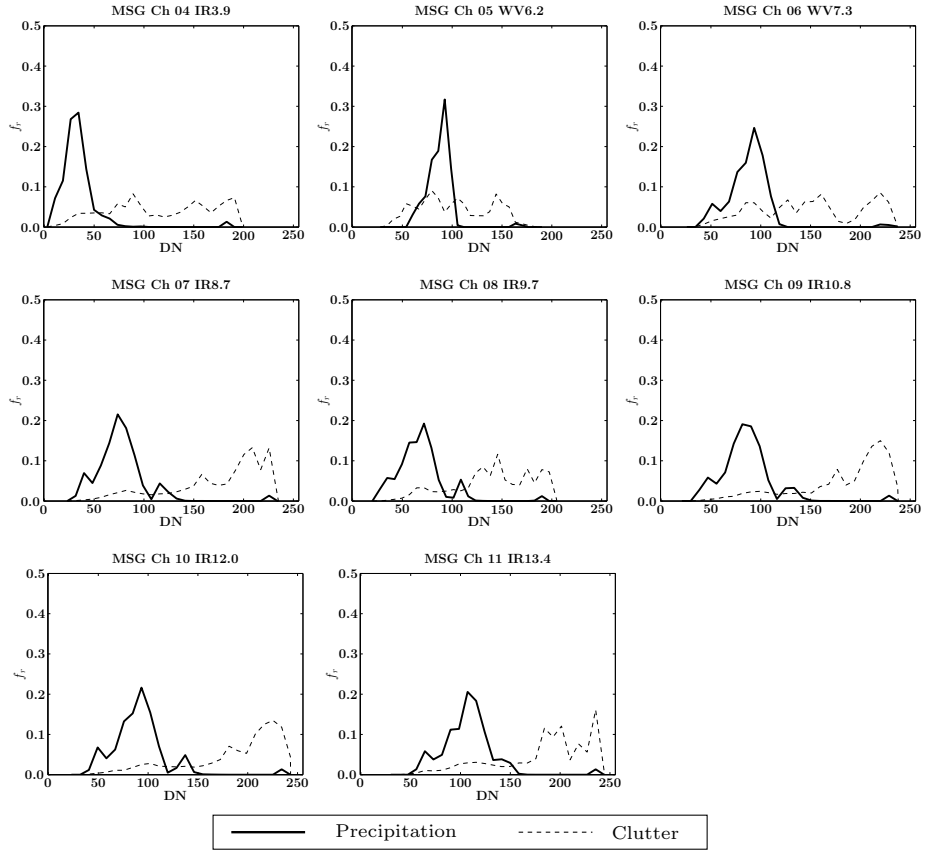

Figure B.7: Data Fusion features. Night features and all echoes. Two classes.

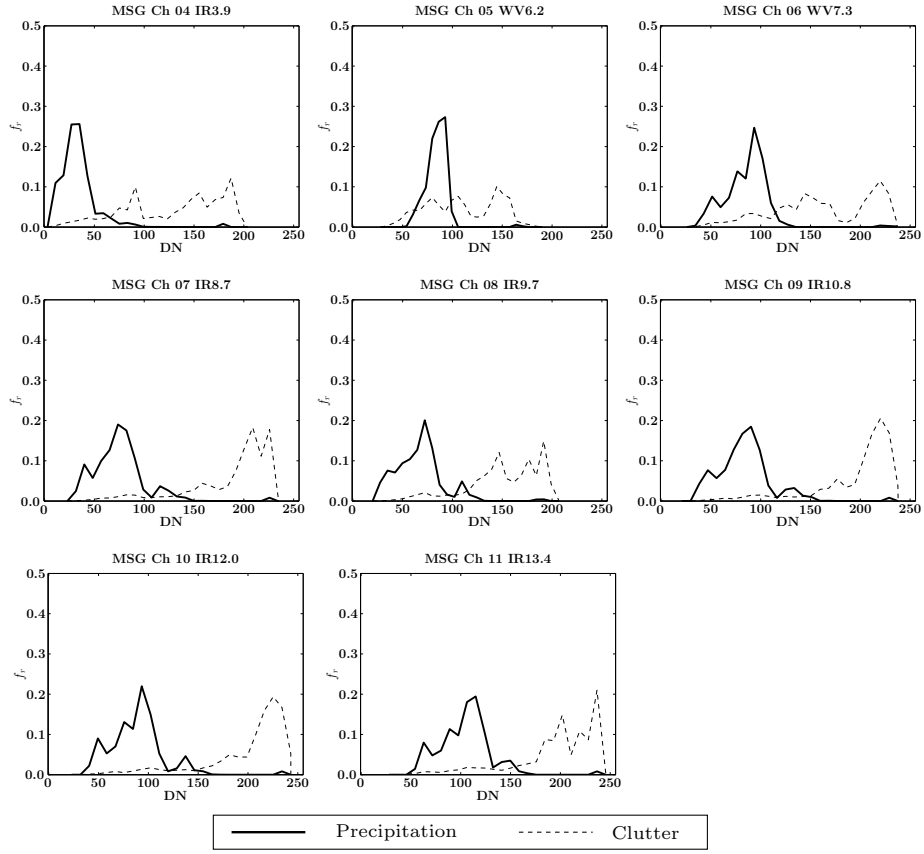

Figure B.8: Data Fusion features. Night features and weak echoes. Two classes.

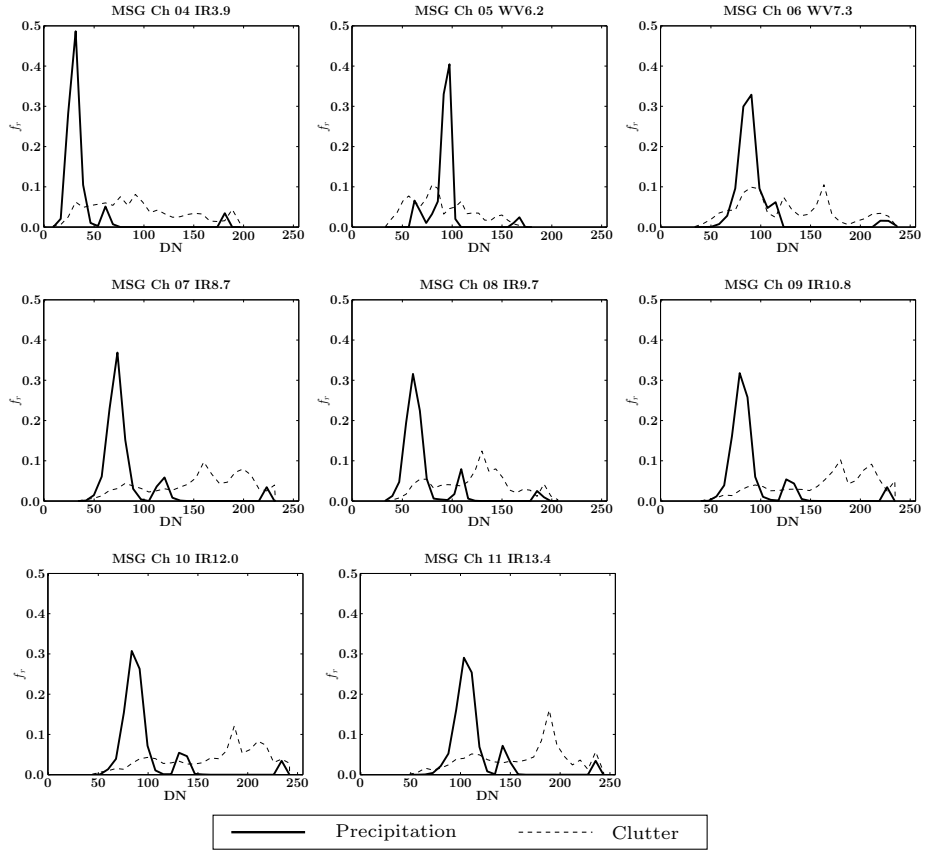

Figure B.9: Data Fusion features. Night features and strong echoes. Two classes.

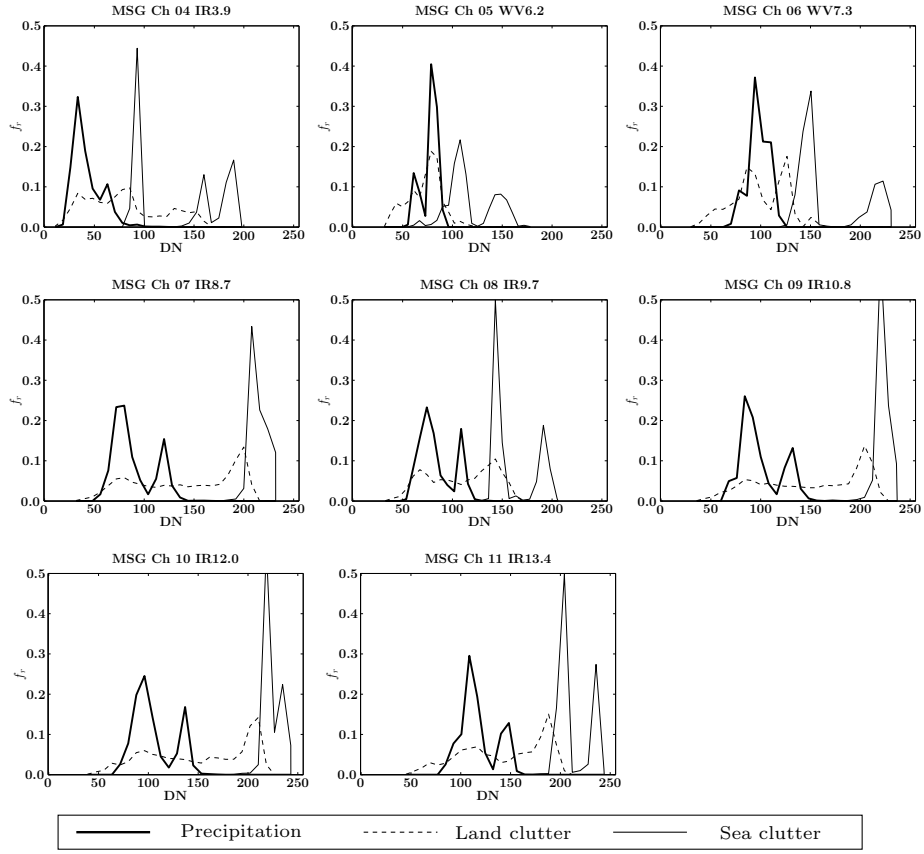

Figure B.10: Data Fusion features. Night features and all echoes. Three classes.

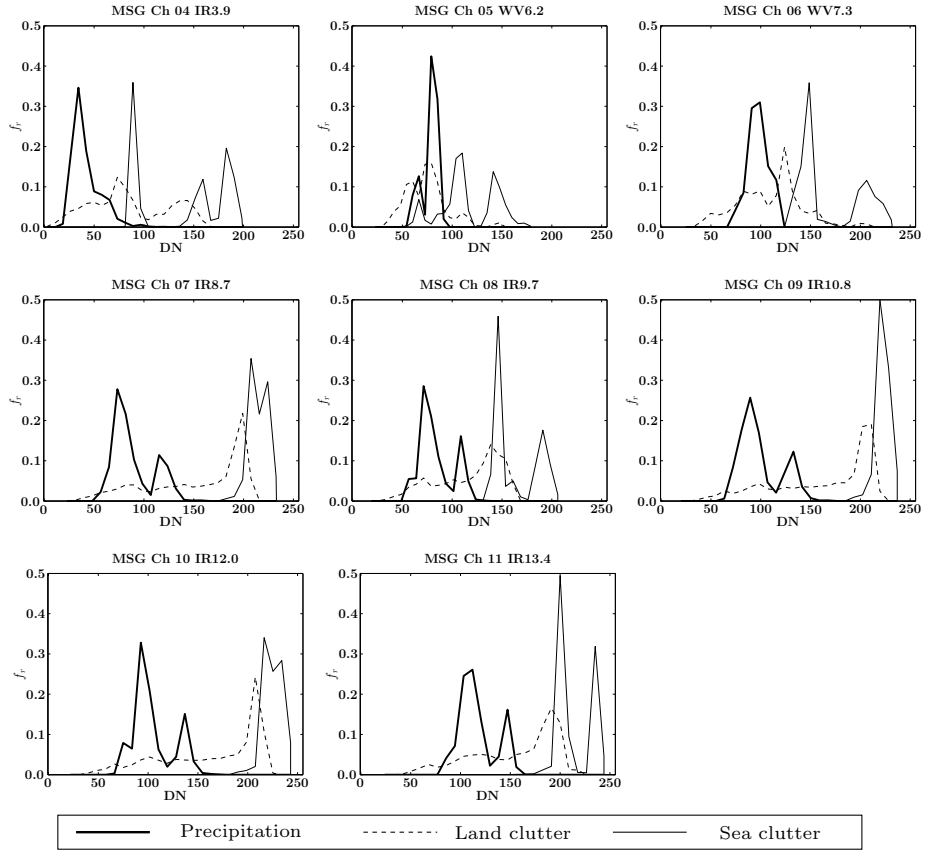

Figure B.11: Data Fusion features. Night features and weak echoes. Three classes.

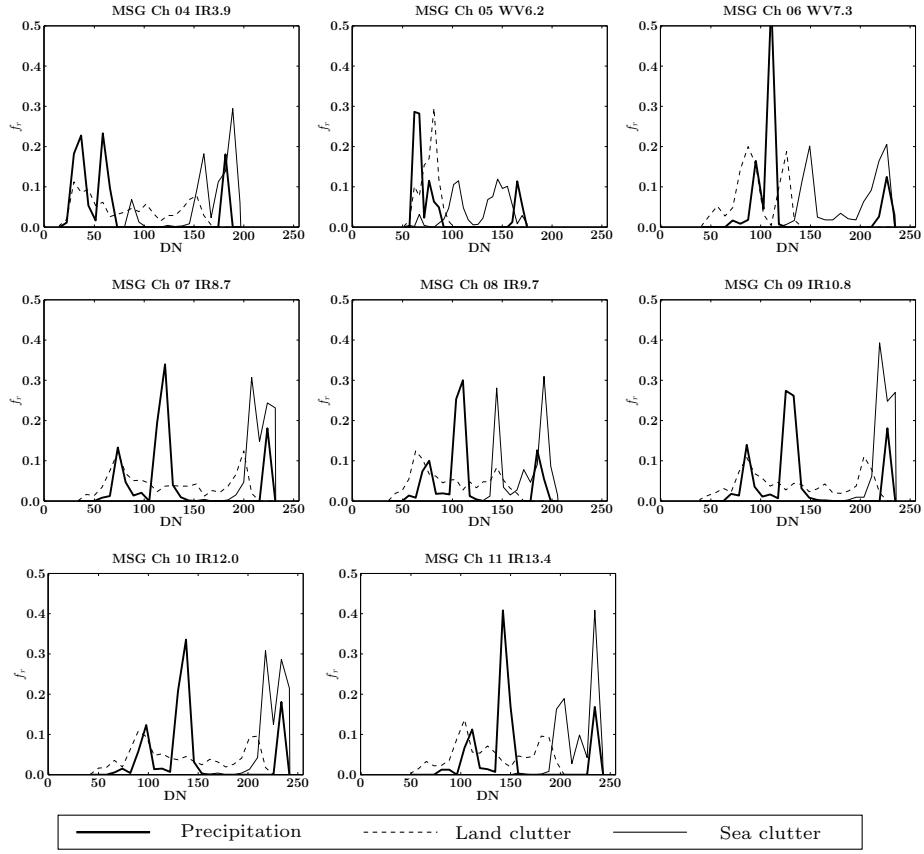

Figure B.12: Data Fusion features. Night features and strong echoes. Three classes.

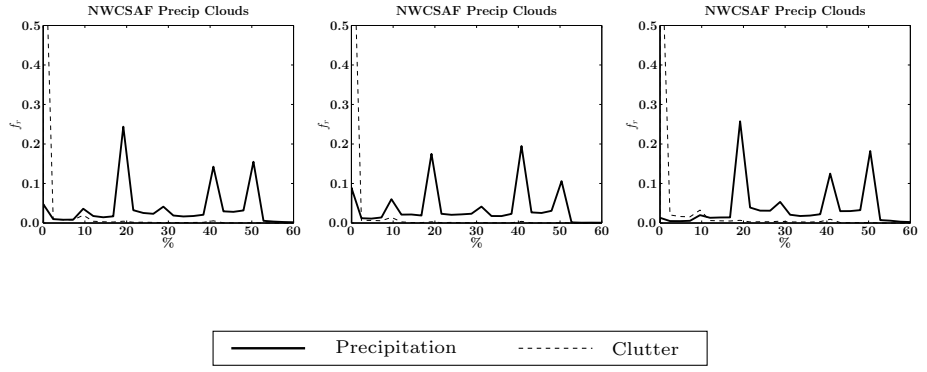

Figure B.13: Nowcasting SAF method. All echoes (left), weak echoes (middle) and strong echoes (right). Two classes.

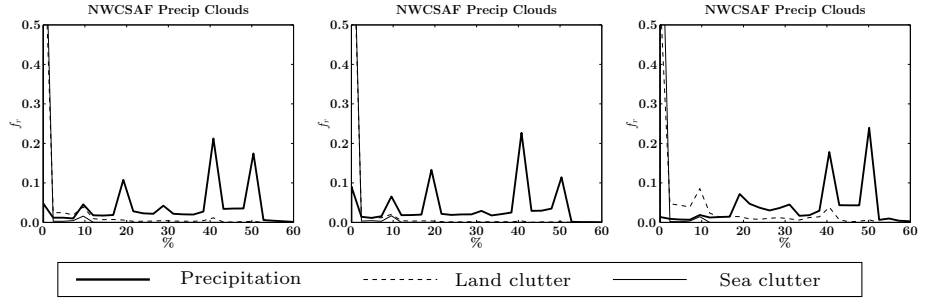

Figure B.14: Nowcasting SAF method. All echoes (left), weak echoes (middle) and strong echoes (right). Three classes.

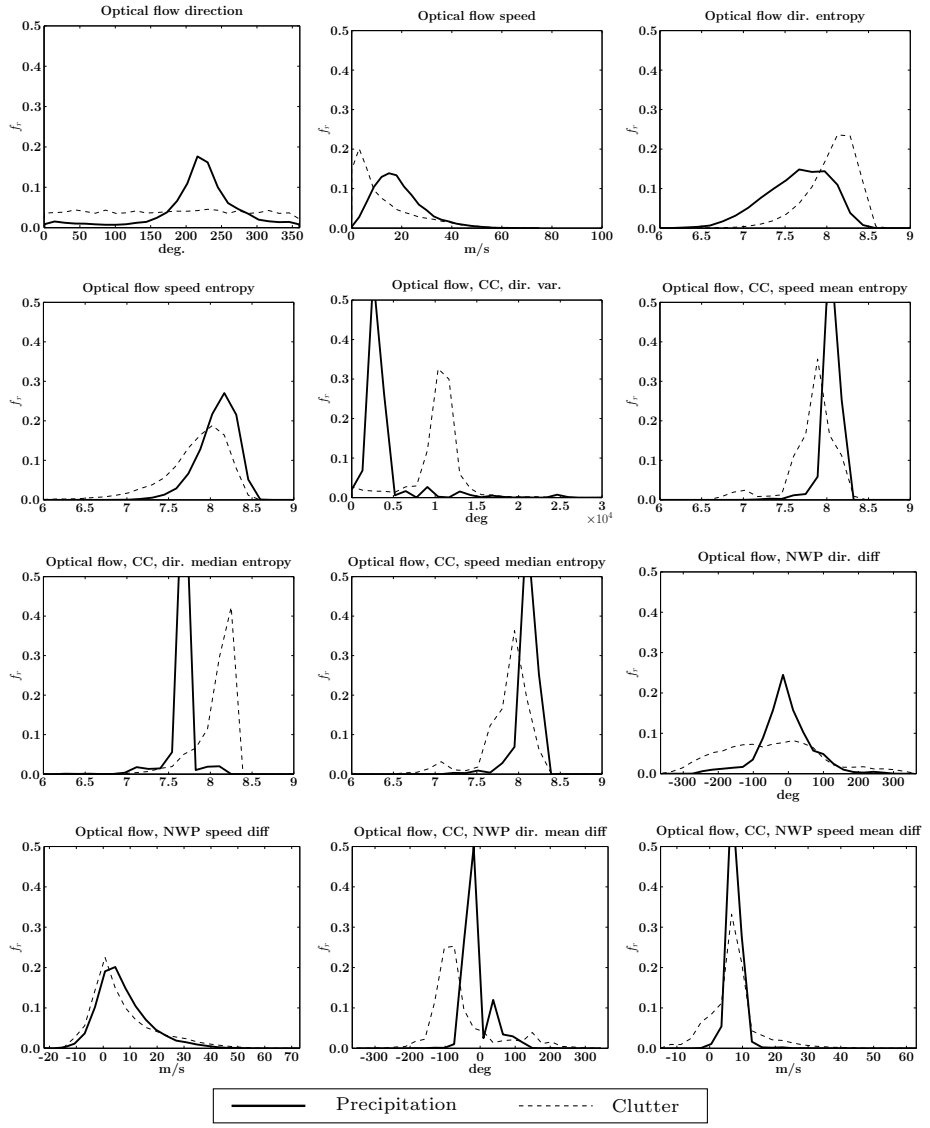

Figure B.15: Spatio-temporal features. All echoes. Two classes.

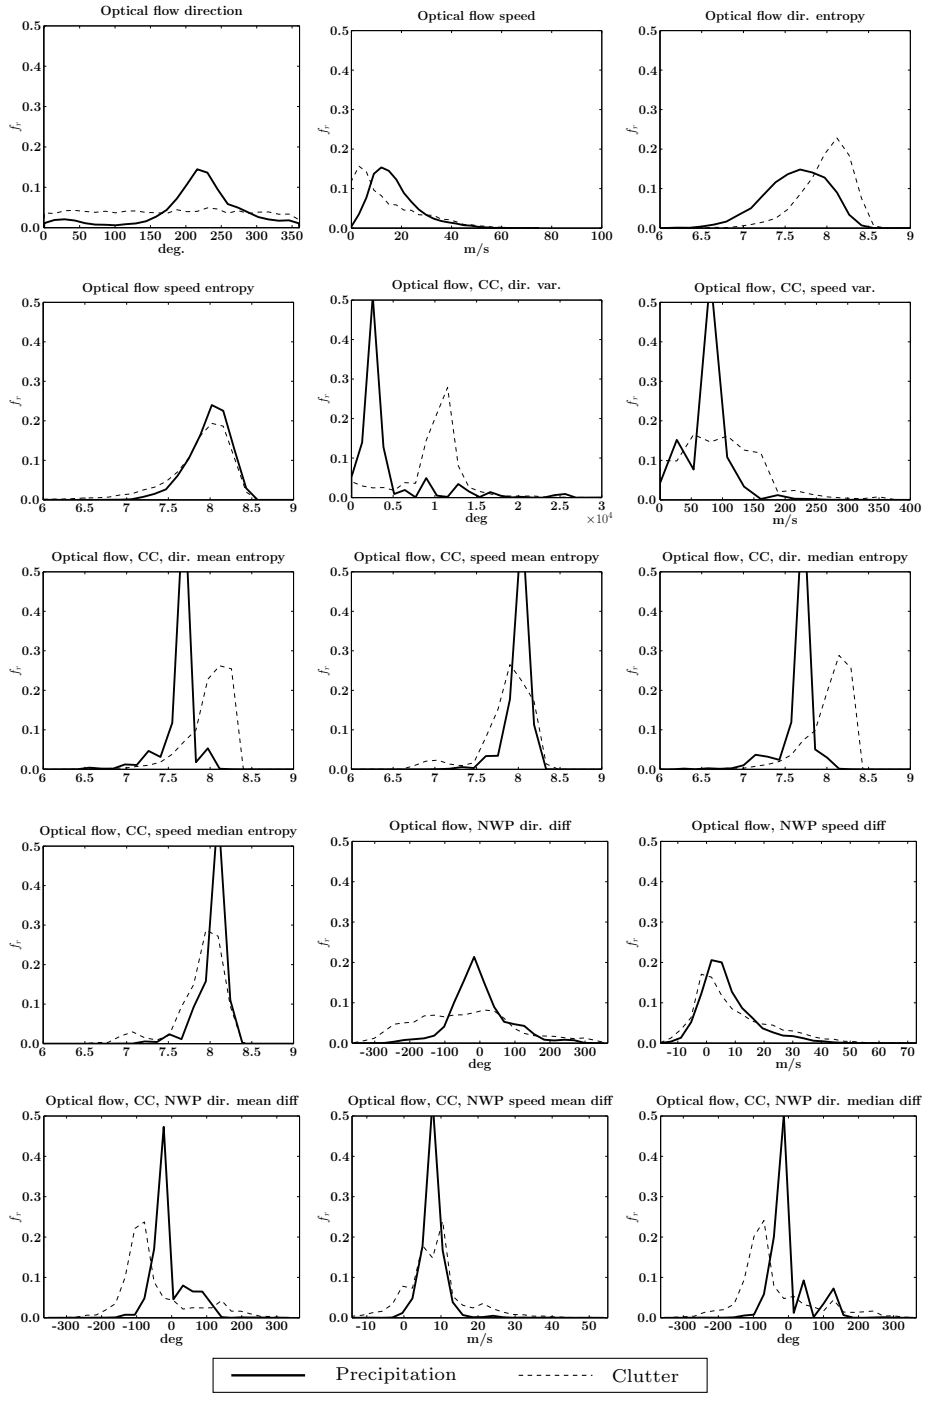

Figure B.16: Spatio-temporal features. Weak echoes. Two classes.

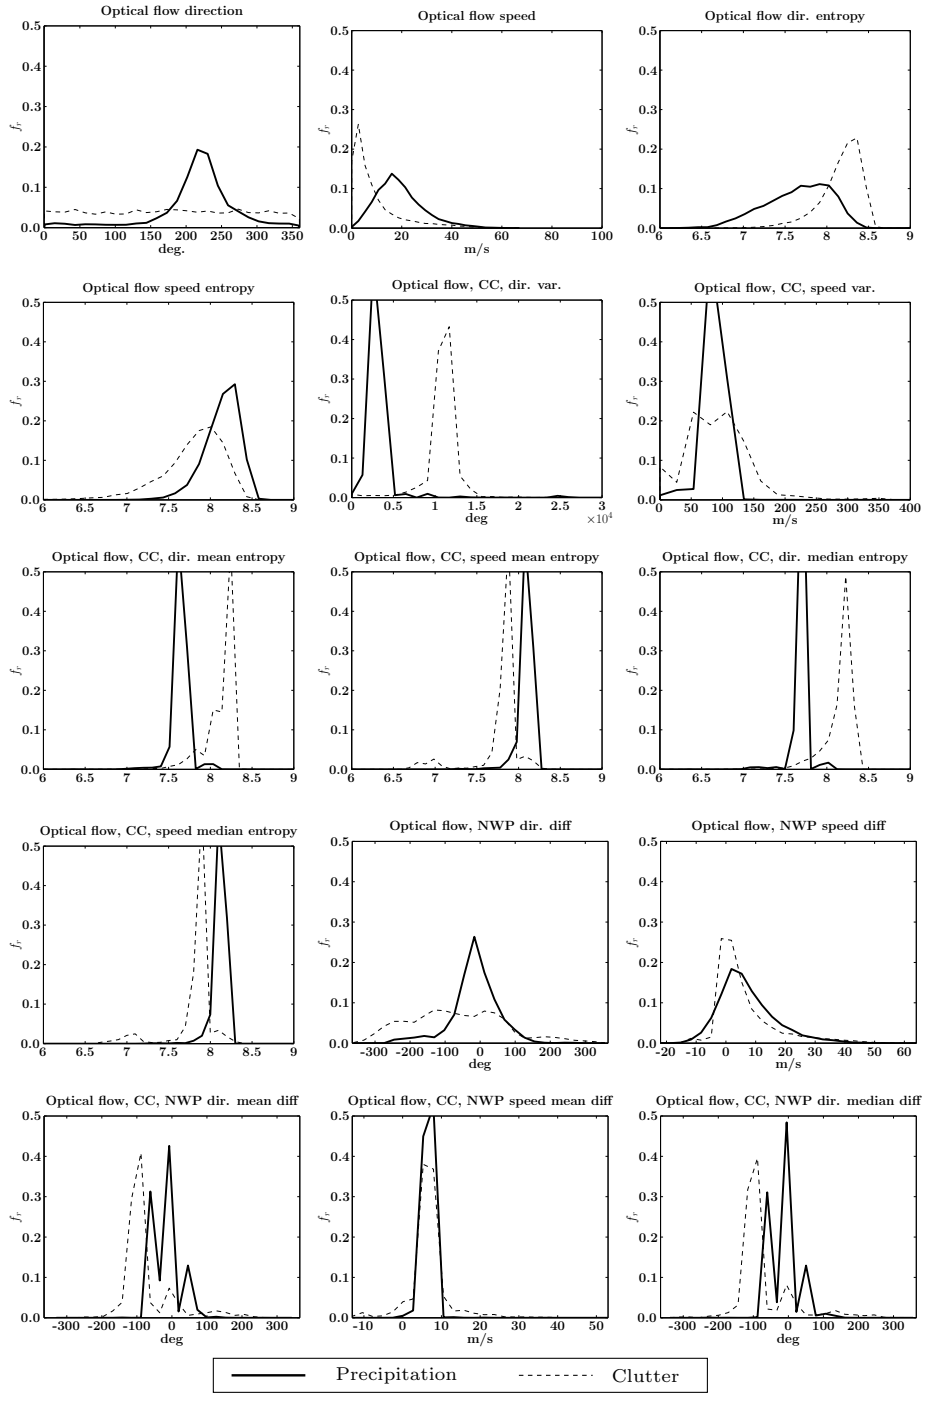

Figure B.17: Spatio-temporal features. Strong echoes. Two classes.

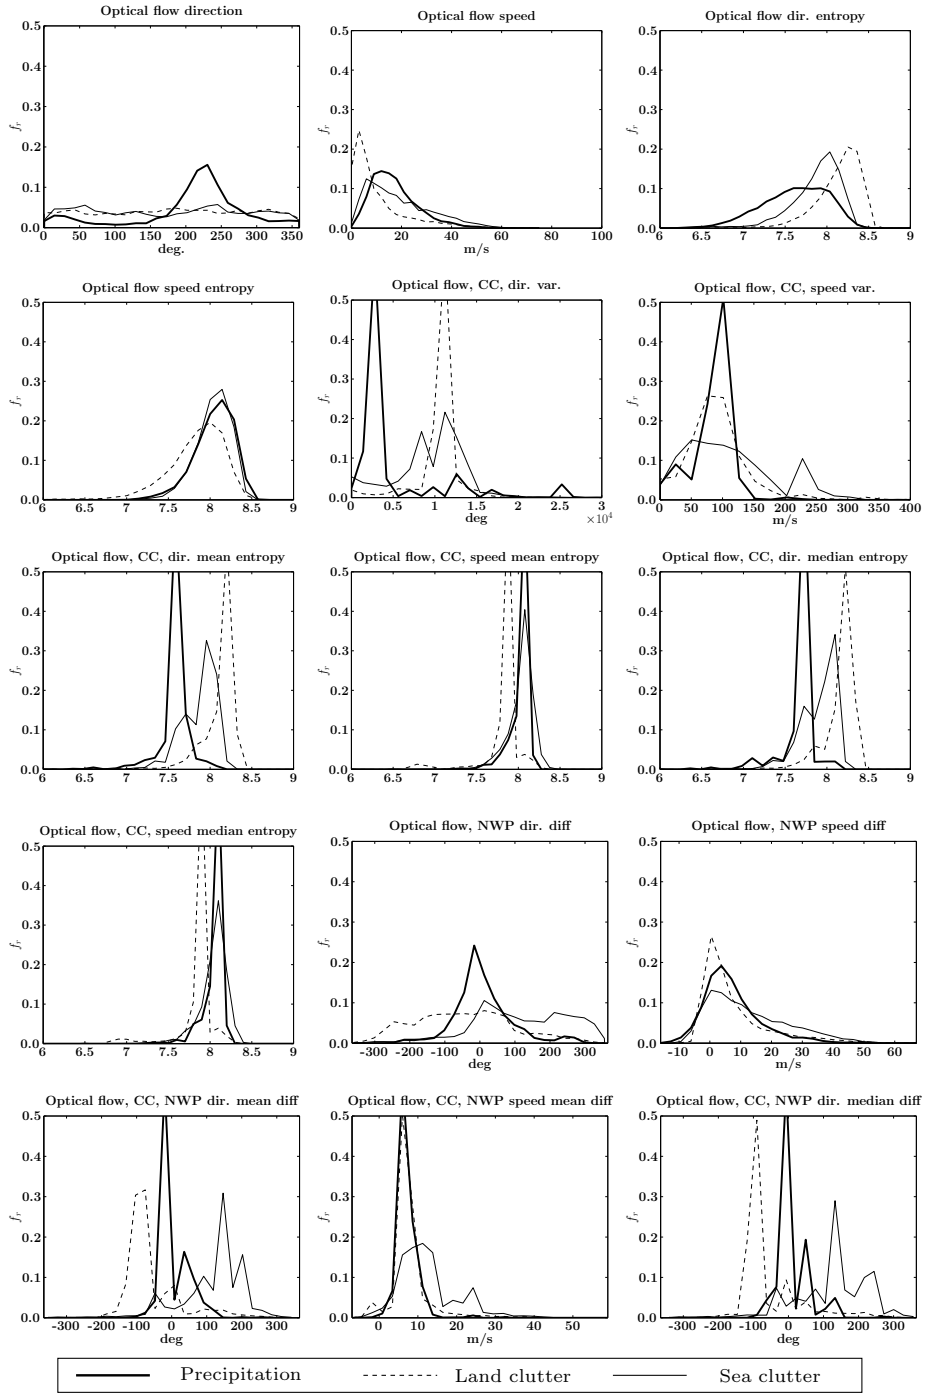

Figure B.18: Spatio-temporal features. All echoes. Three classes.

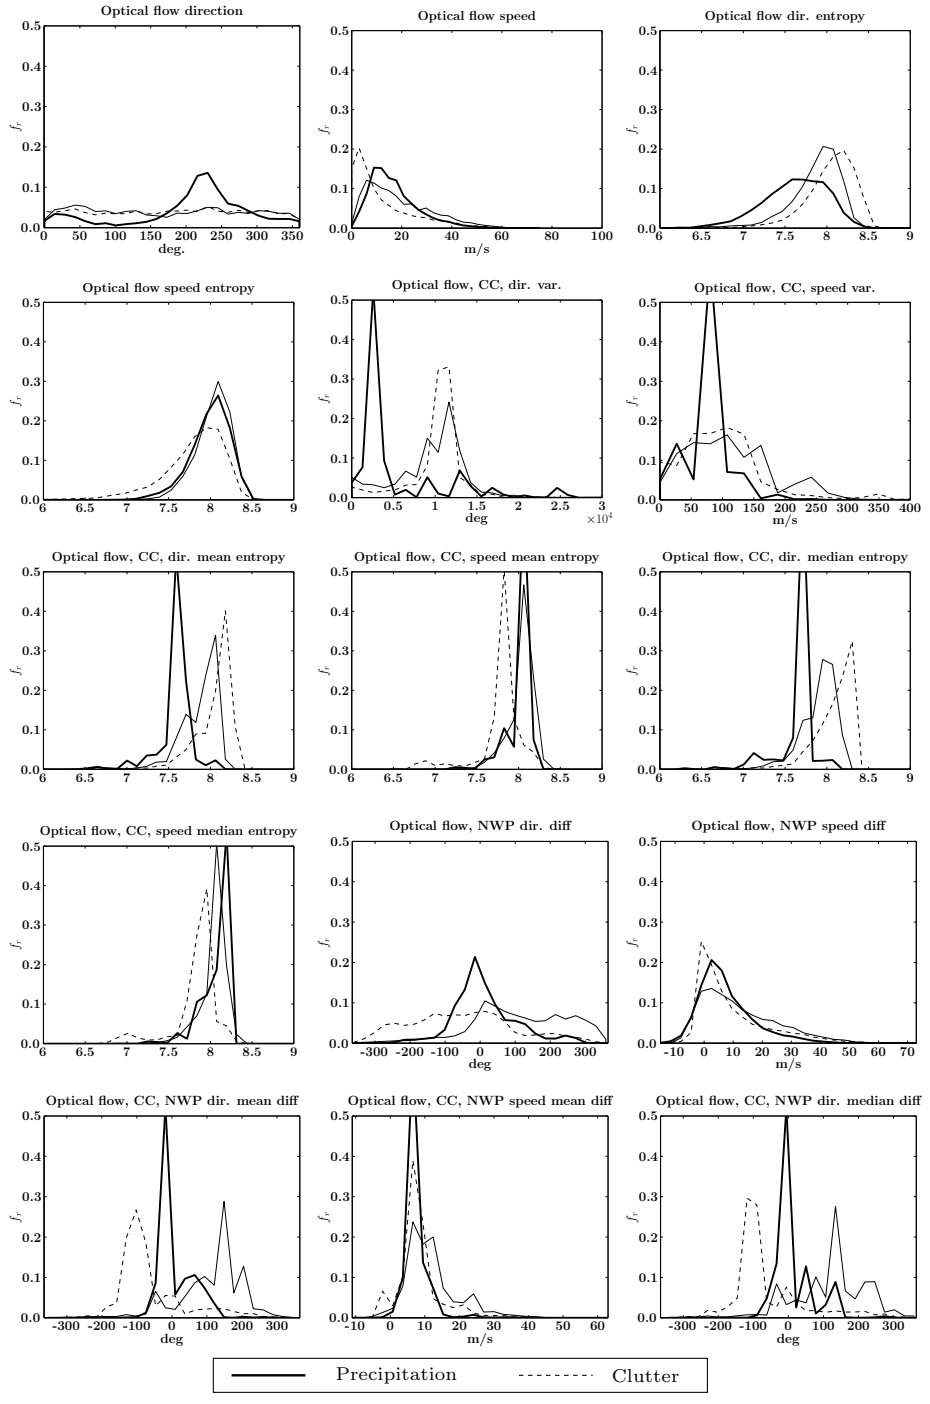

Figure B.19: Spatio-temporal features. Weak echoes. Three classes.

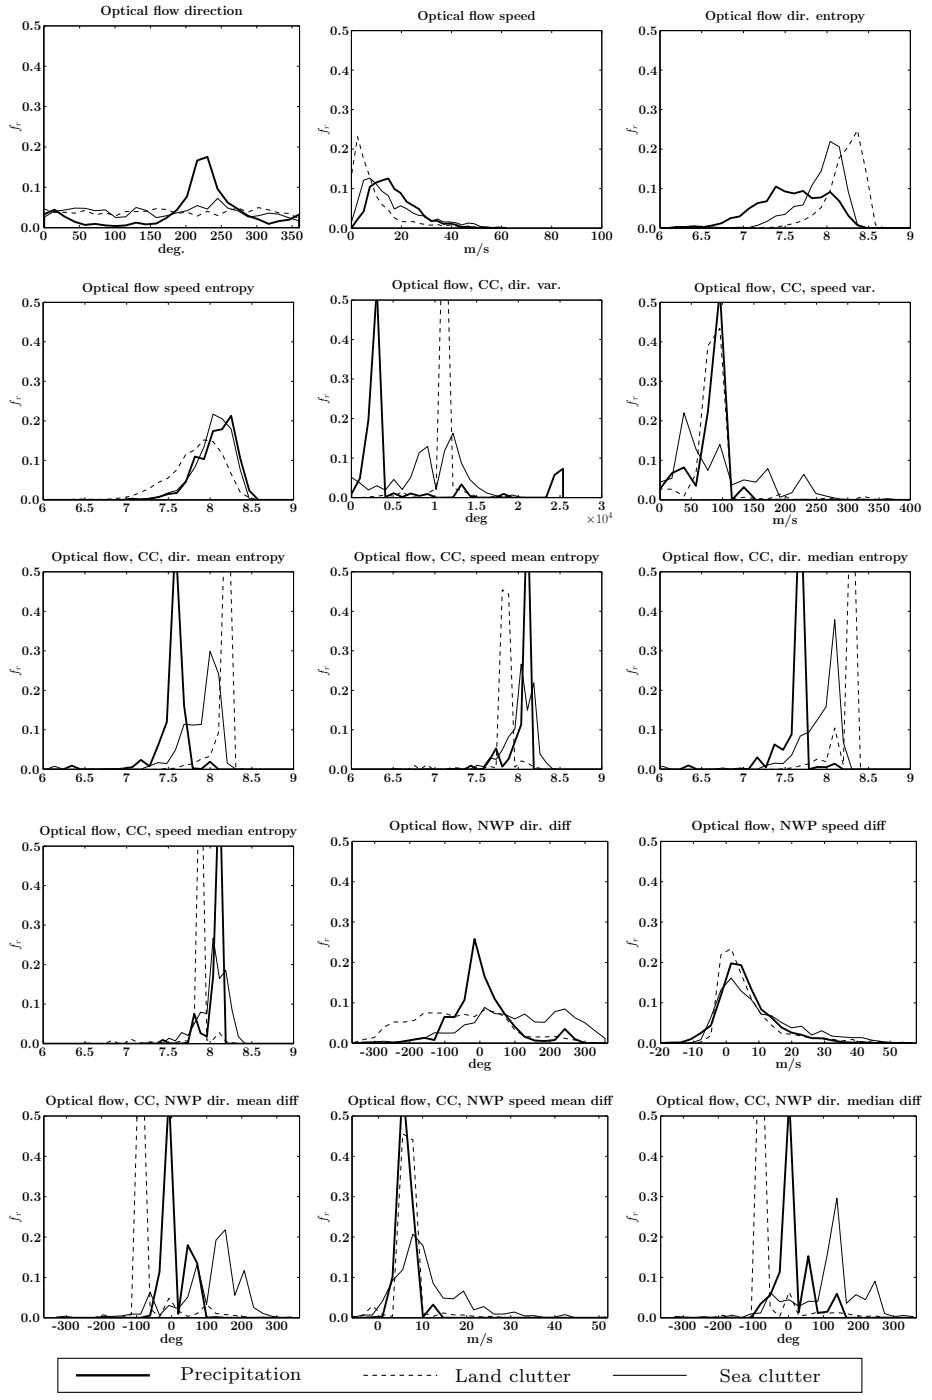

Figure B.20: Spatio-temporal features. Strong echoes. Three classes.

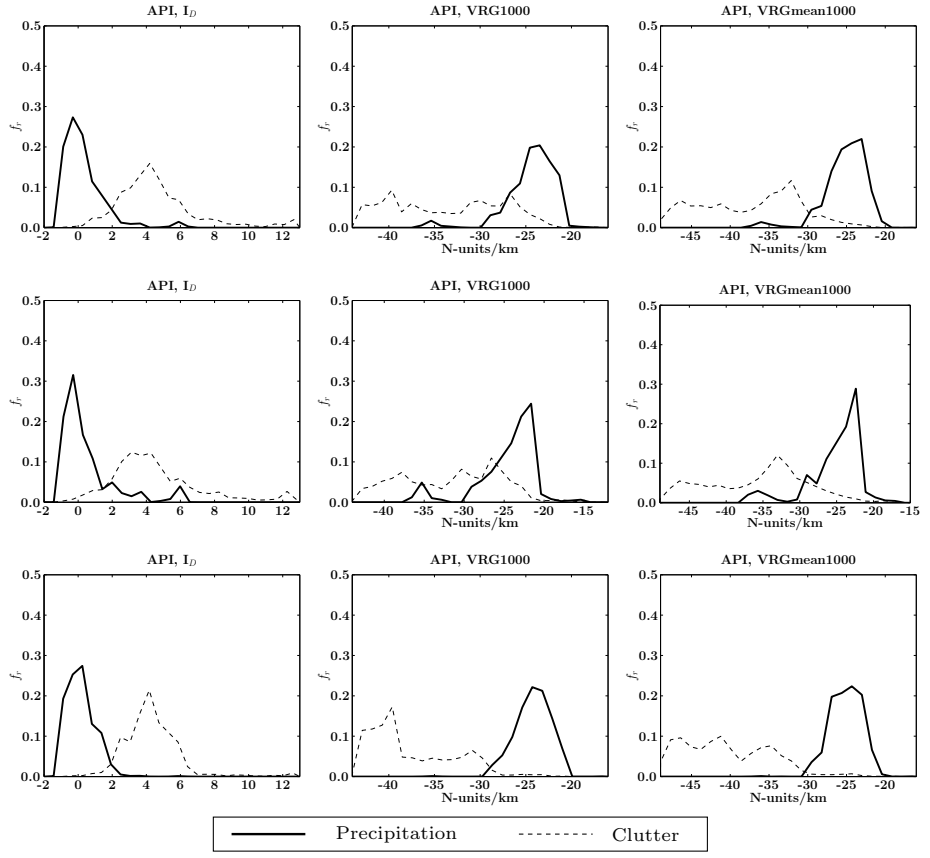

Figure B.21: Propagation method. All echoes (top row), weak echoes (middle row) and strong echoes (bottom row). Two classes.

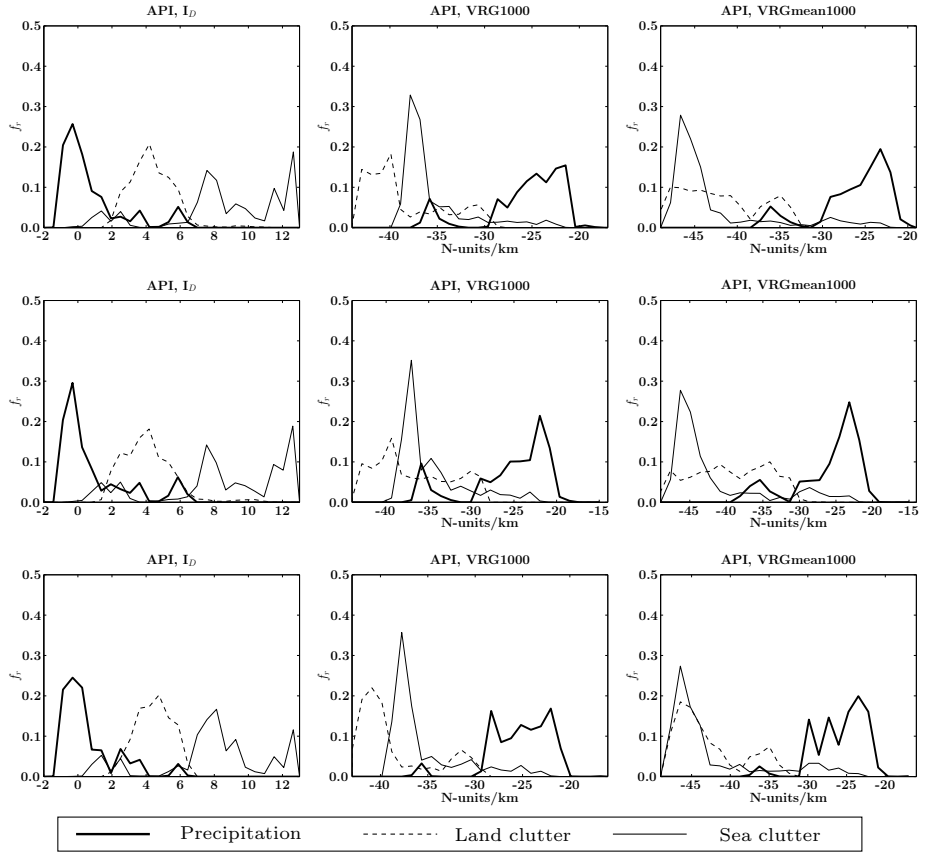

Figure B.22: Propagation method. All echoes (top row), weak echoes (middle row) and strong echoes (bottom row). Three classes.

Appendix C

Feature conditional probabilities

Feature conditional probabilities

This appendix contains conditional probability histograms for the features which are used in the methods.

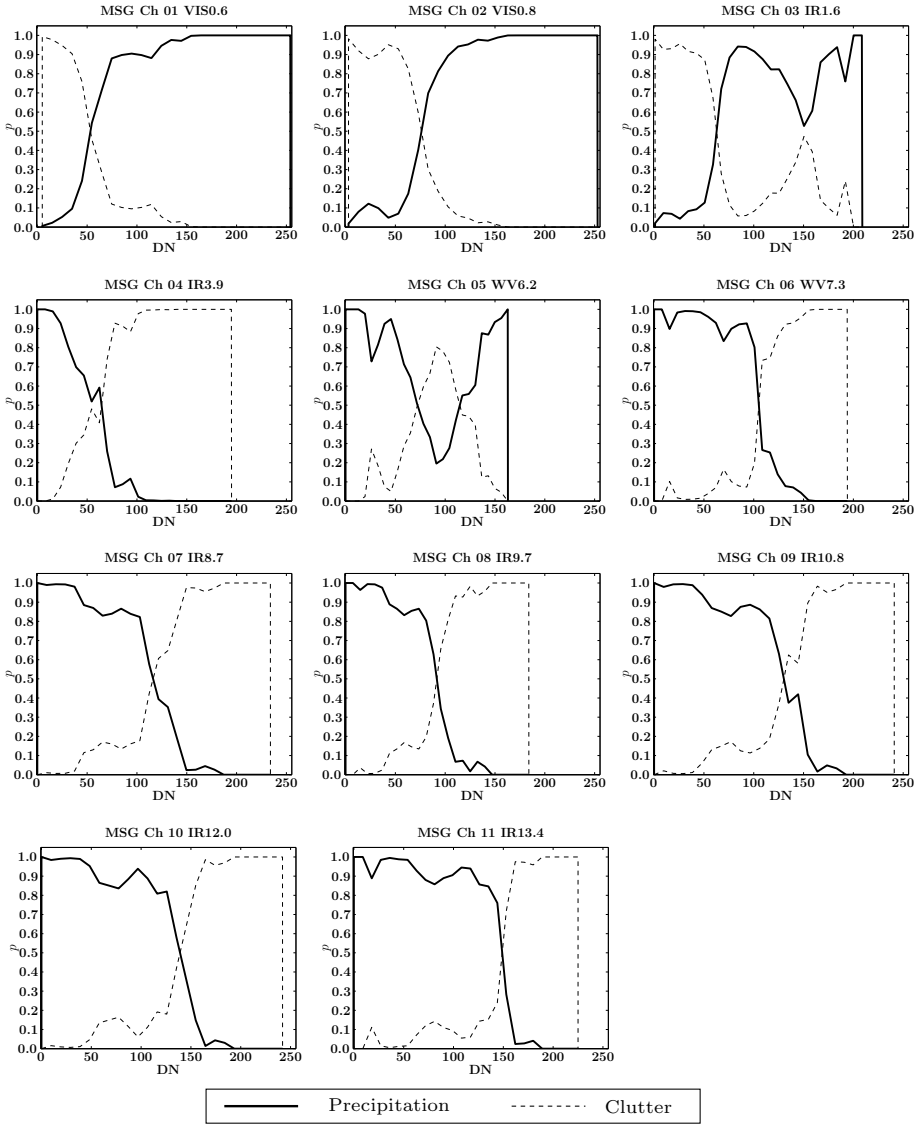

Figure C.1: Data Fusion features. Day features and all echoes. Two classes.

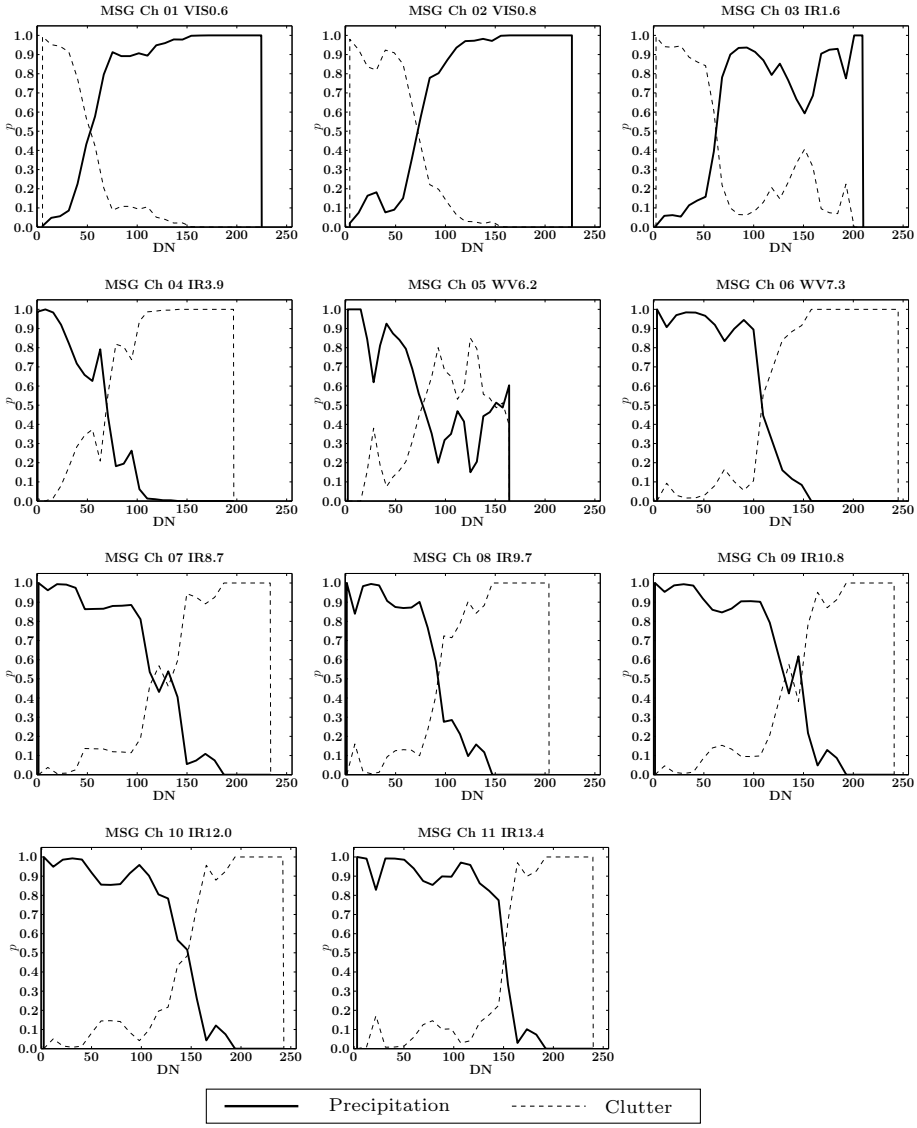

Figure C.2: Data Fusion features. Day features and weak echoes. Two classes.

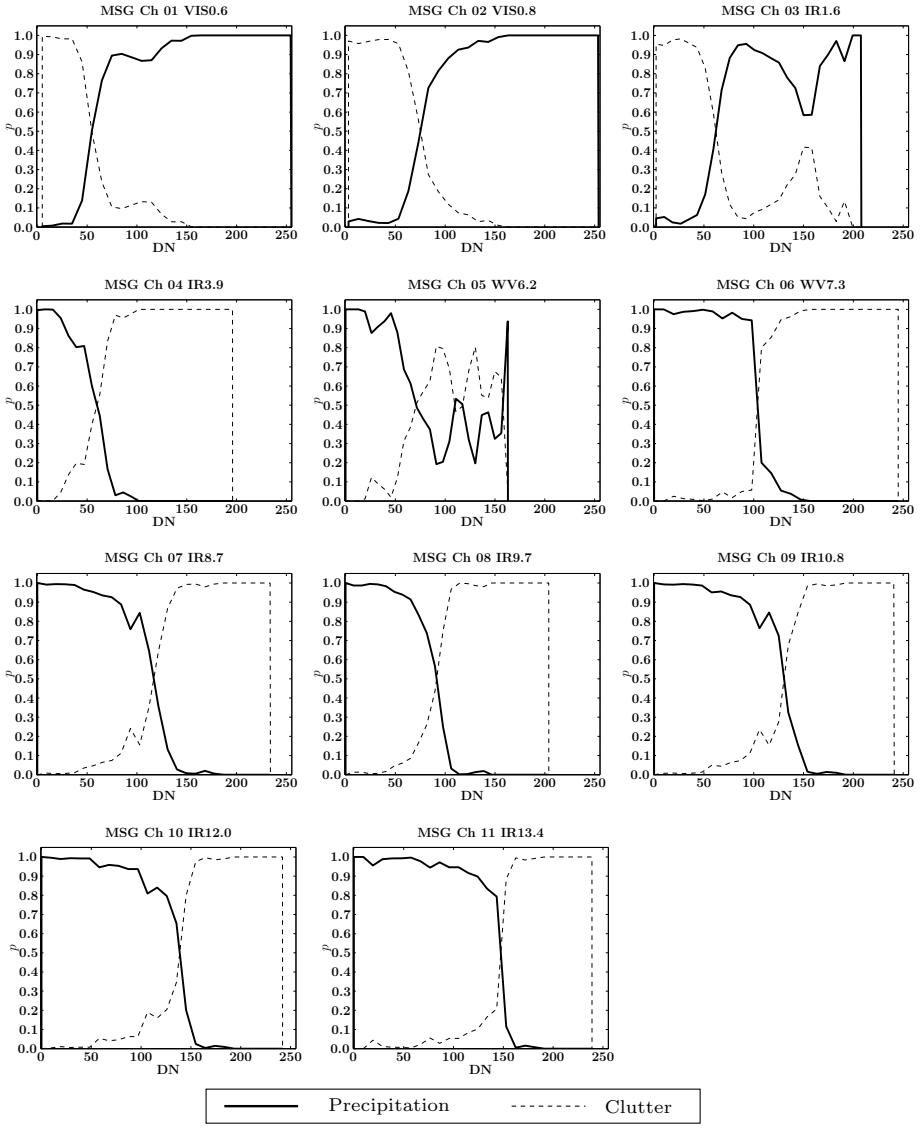

Figure C.3: Data Fusion features. Day features and strong echoes. Two classes.

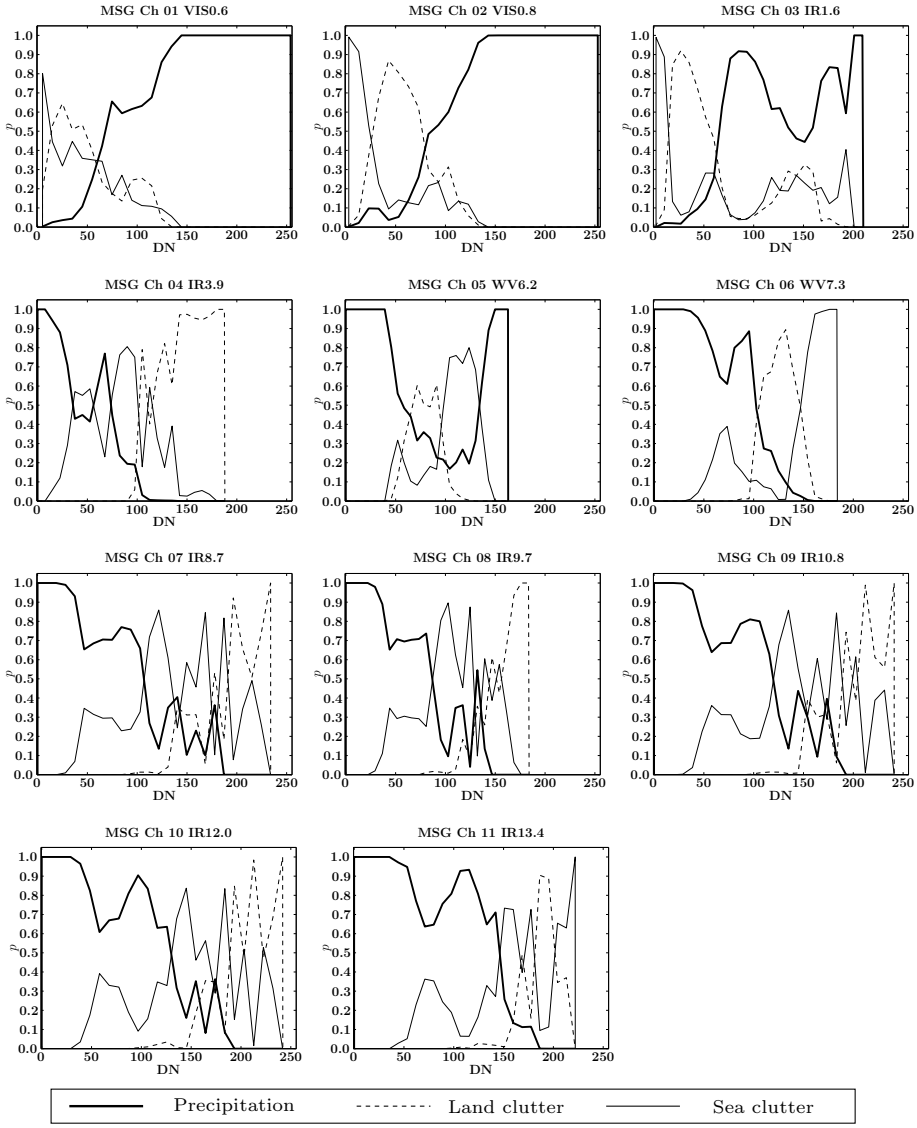

Figure C.4: Data Fusion features. Day features and all echoes. Three classes.

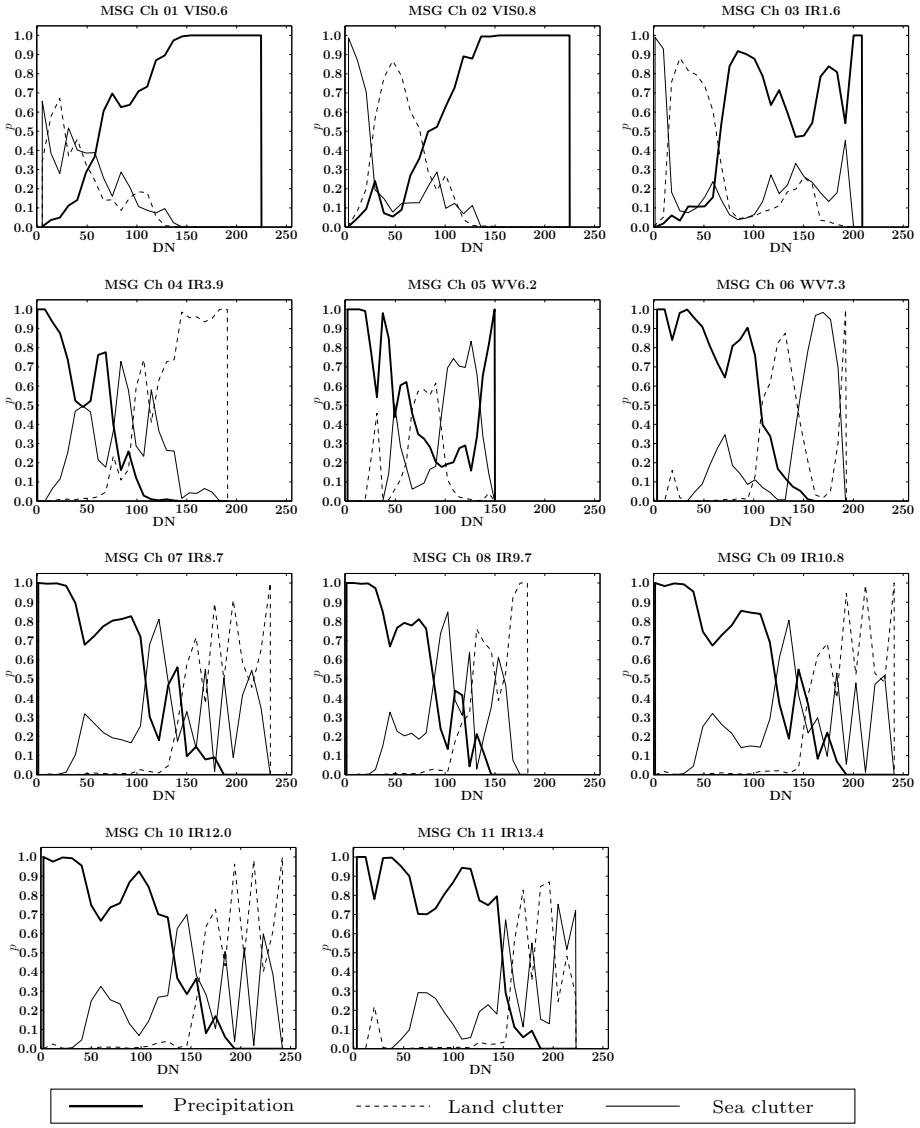

Figure C.5: Data Fusion features. Day features and weak echoes. Three classes.

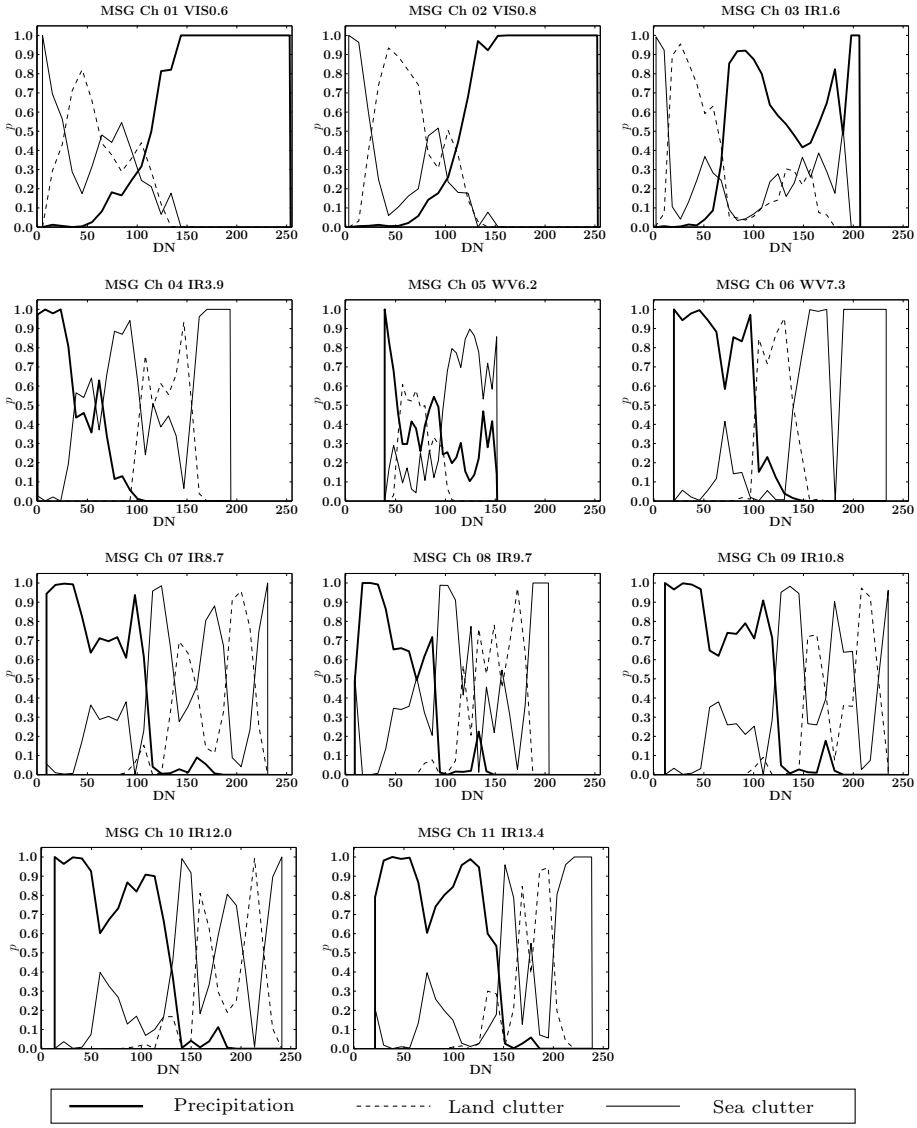

Figure C.6: Data Fusion features. Day features and strong echoes. Three classes.

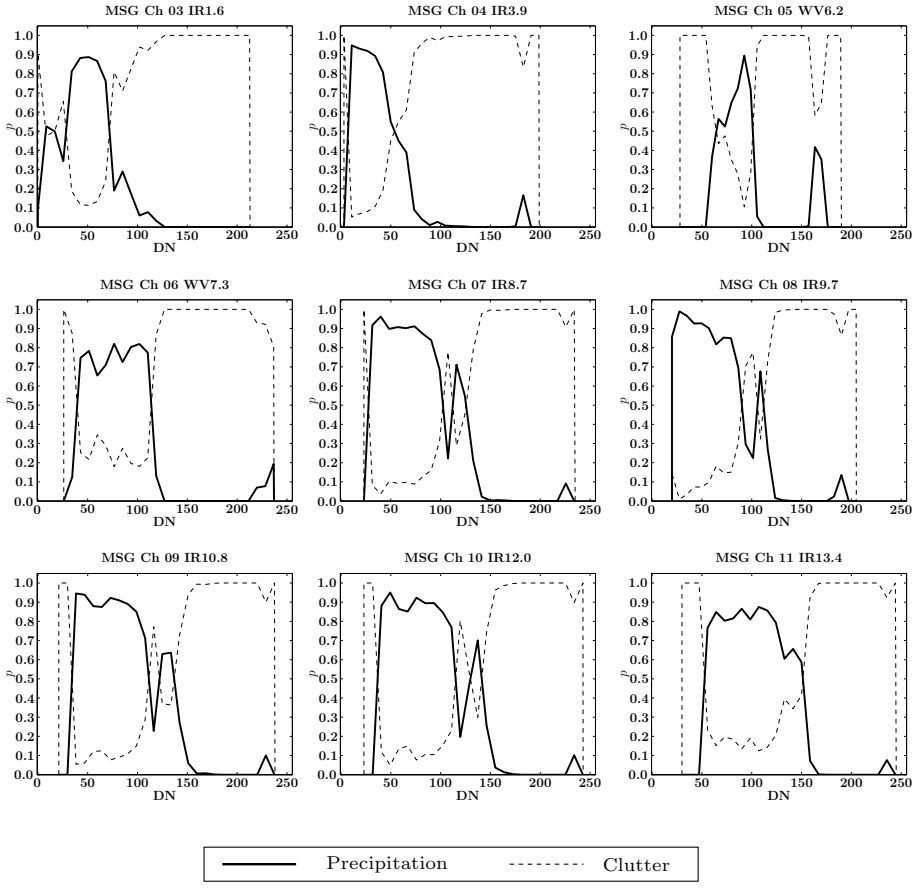

Figure C.7: Data Fusion features. Night features and all echoes. Two classes.

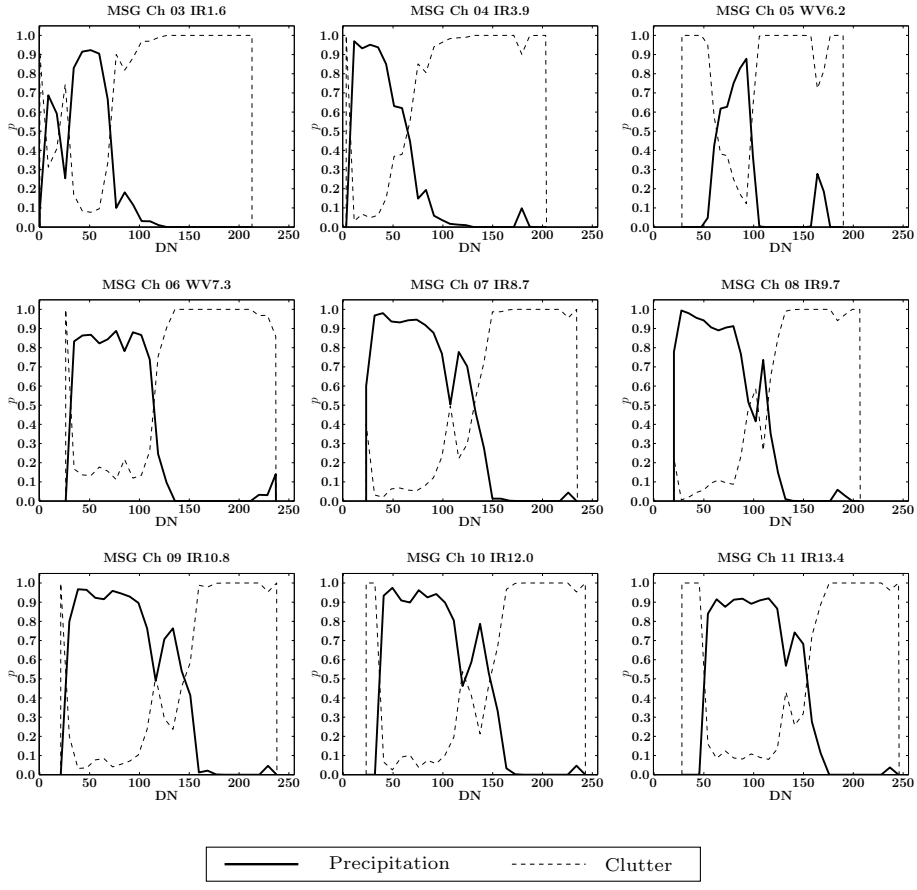

Figure C.8: Data Fusion features. Night features and weak echoes. Two classes.

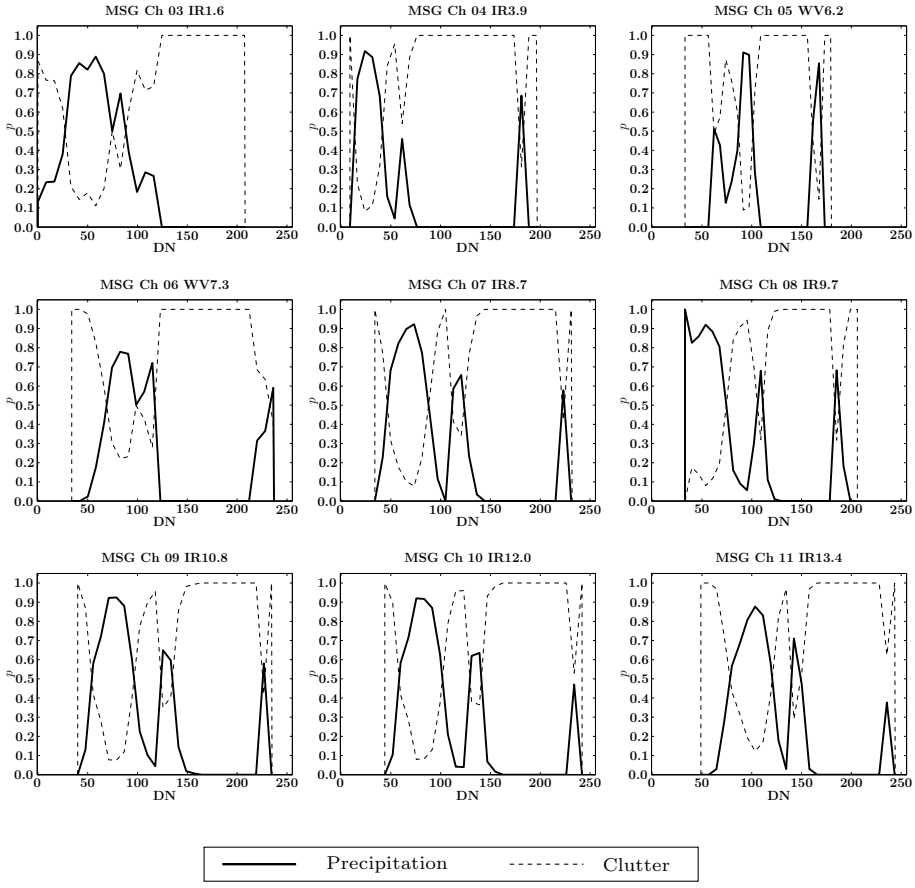

Figure C.9: Data Fusion features. Night features and strong echoes. Two classes.

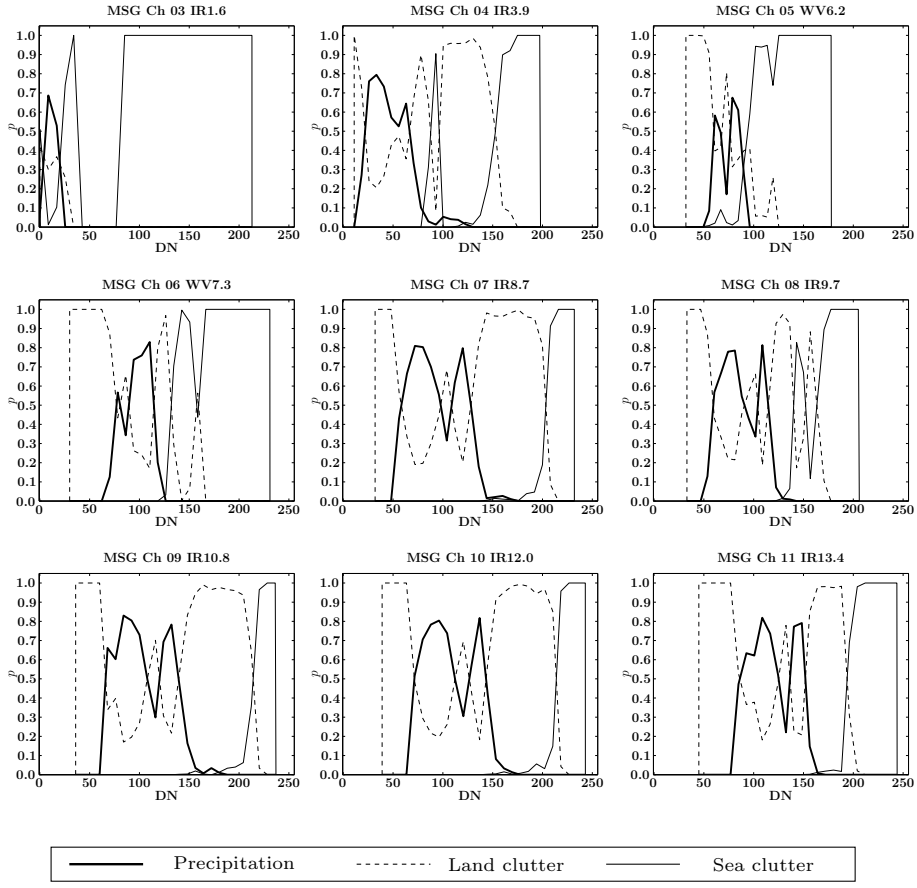

Figure C.10: Data Fusion features. Night features and all echoes. Three classes.

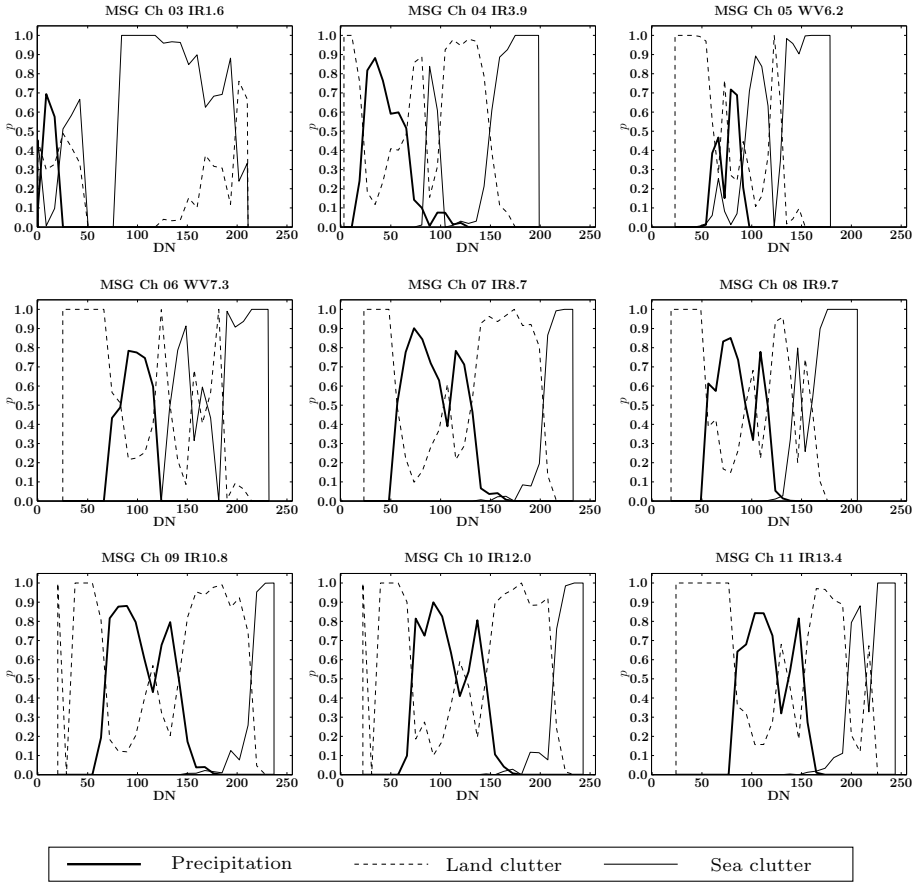

Figure C.11: Data Fusion features. Night features and weak echoes. Three classes.

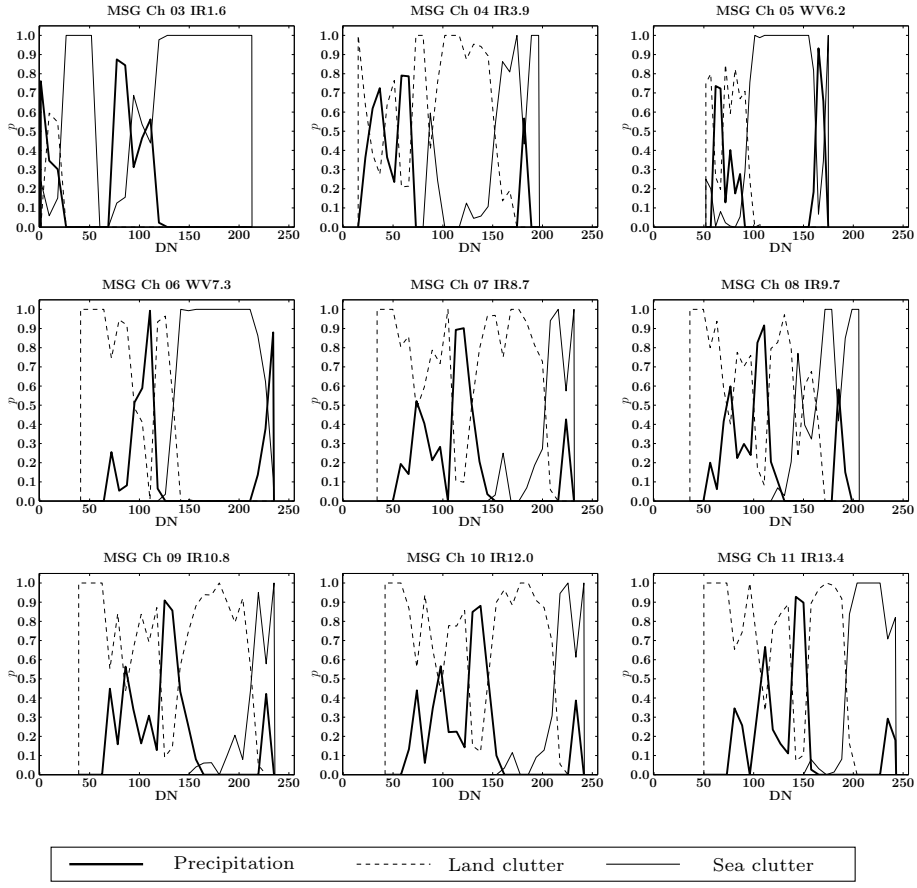

Figure C.12: Data Fusion features. Night features and strong echoes. Three classes.

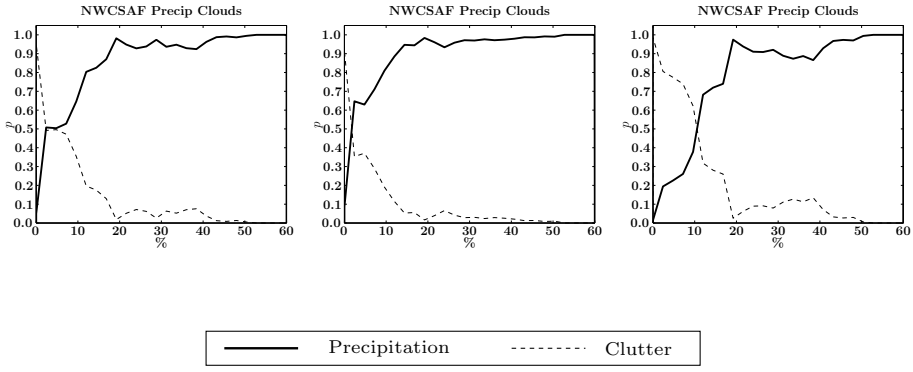

Figure C.13: Nowcasting SAF method. All echoes (left), weak echoes (middle) and strong echoes (right). Two classes.

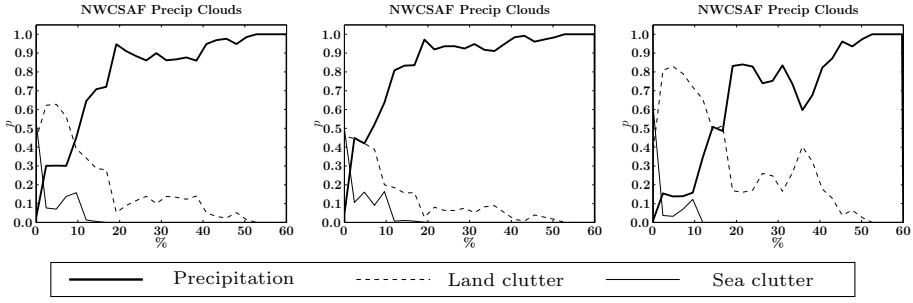

Figure C.14: Nowcasting SAF method. All echoes (left), weak echoes (middle) and strong echoes (right). Three classes.

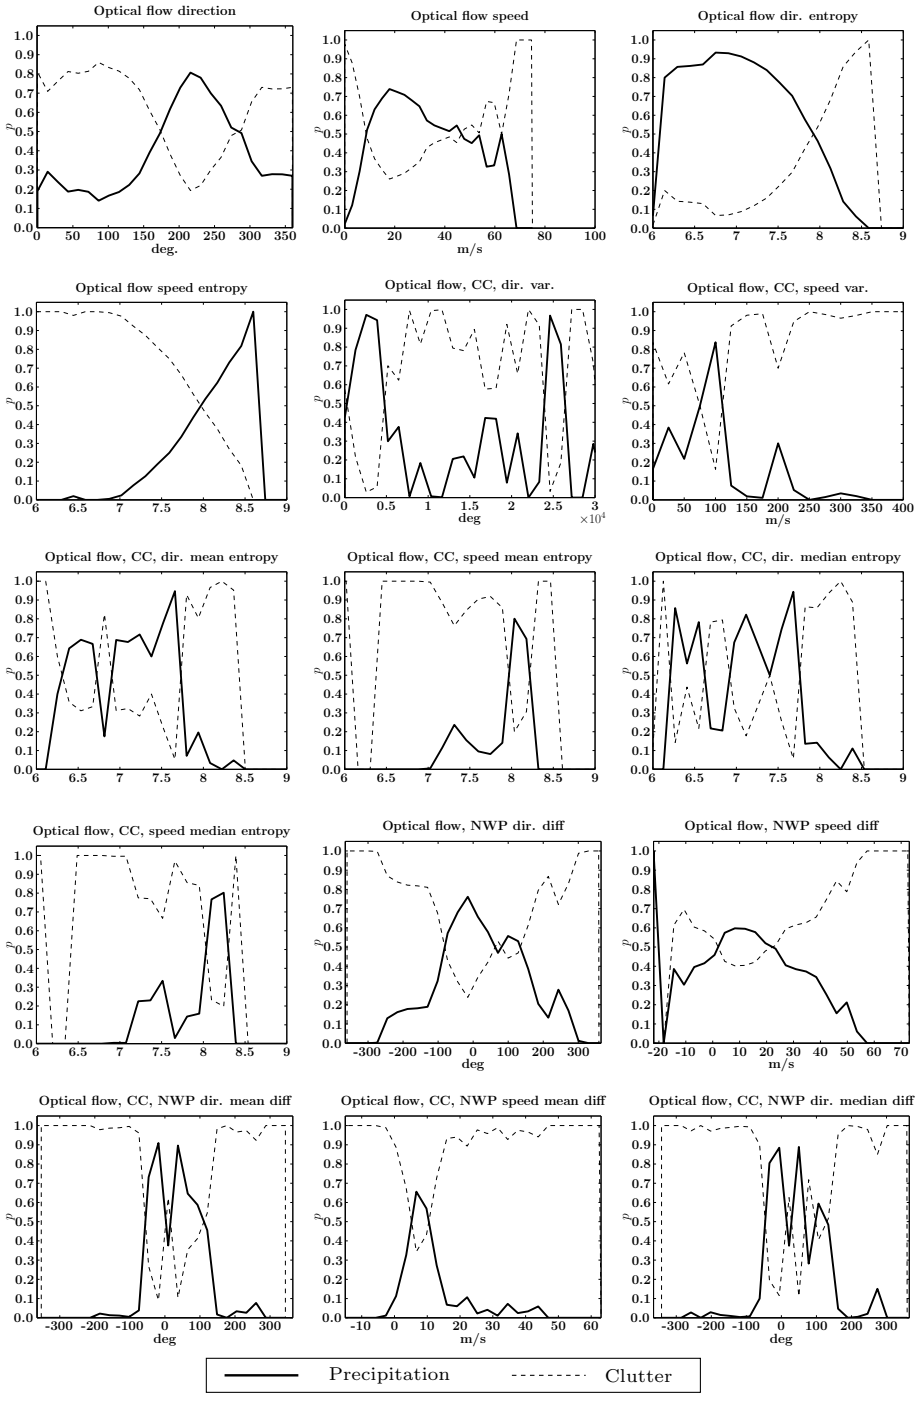

Figure C.15: Spatio-temporal features. All echoes. Two classes.

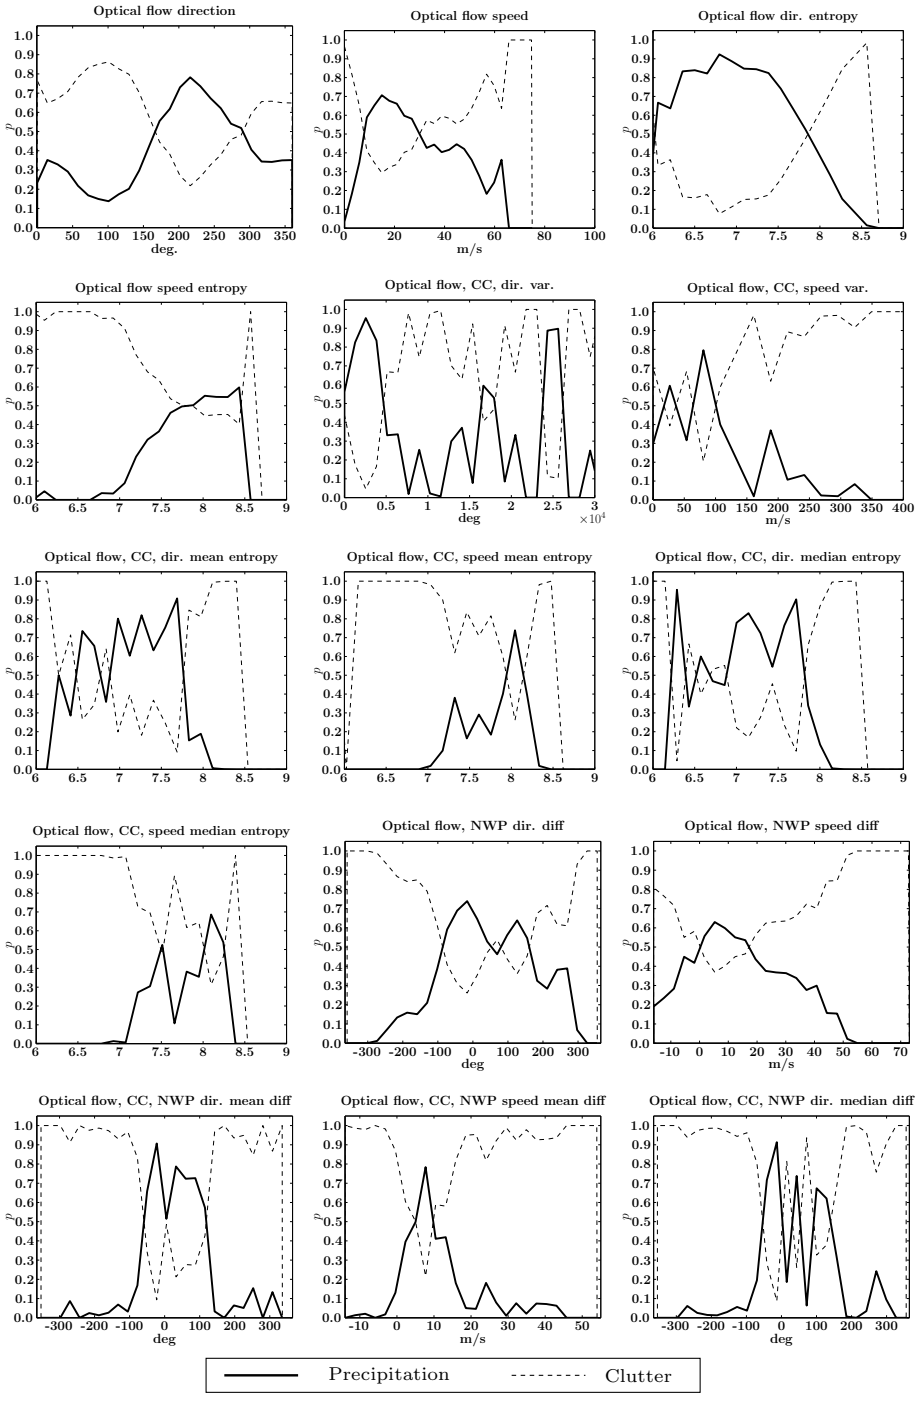

Figure C.16: Spatio-temporal features. Weak echoes. Two classes.

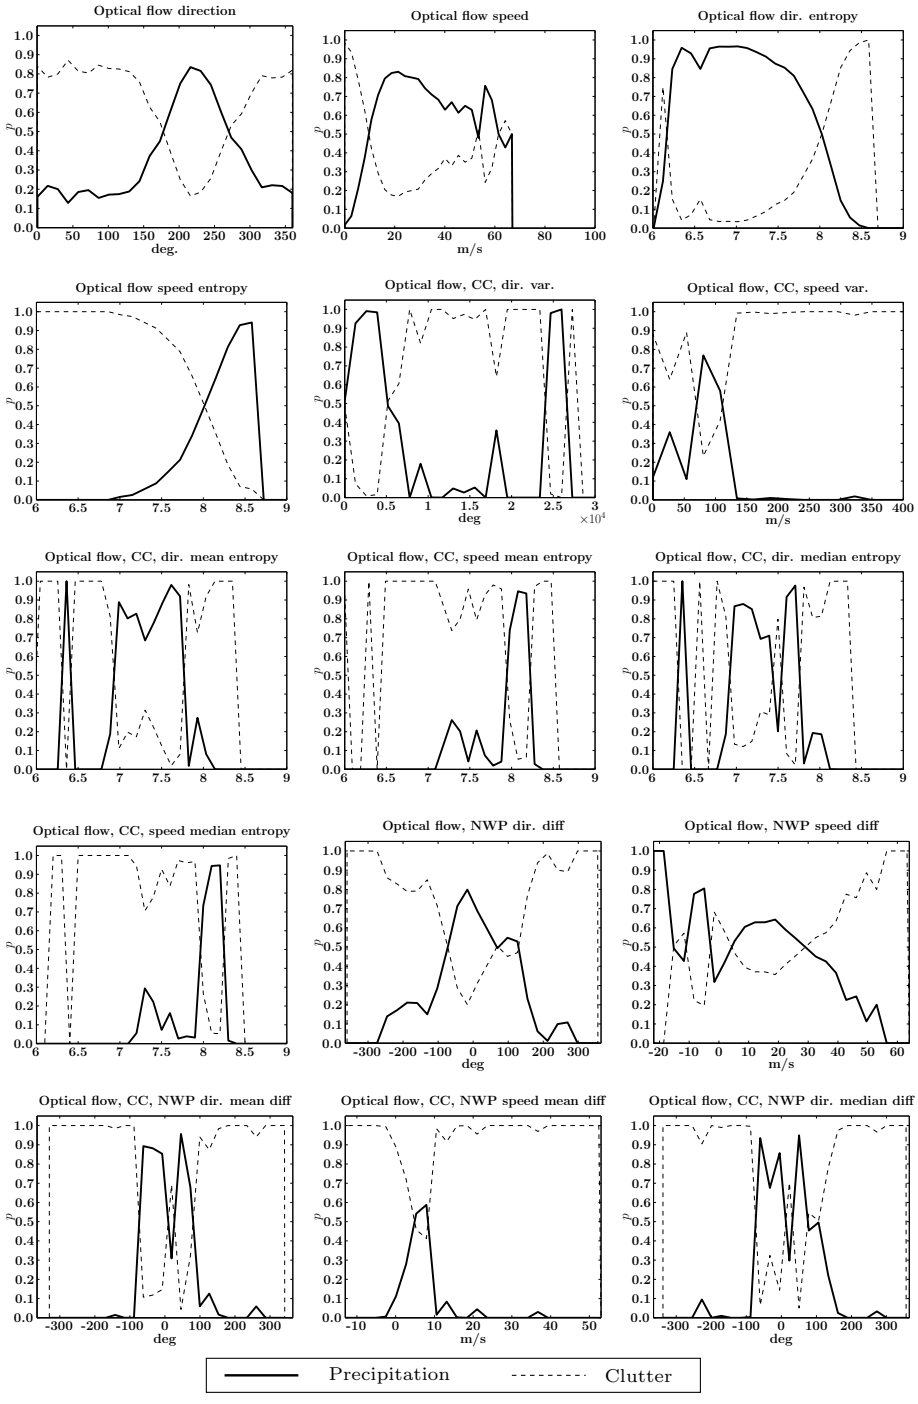

Figure C.17: Spatio-temporal features. Strong echoes. Two classes.

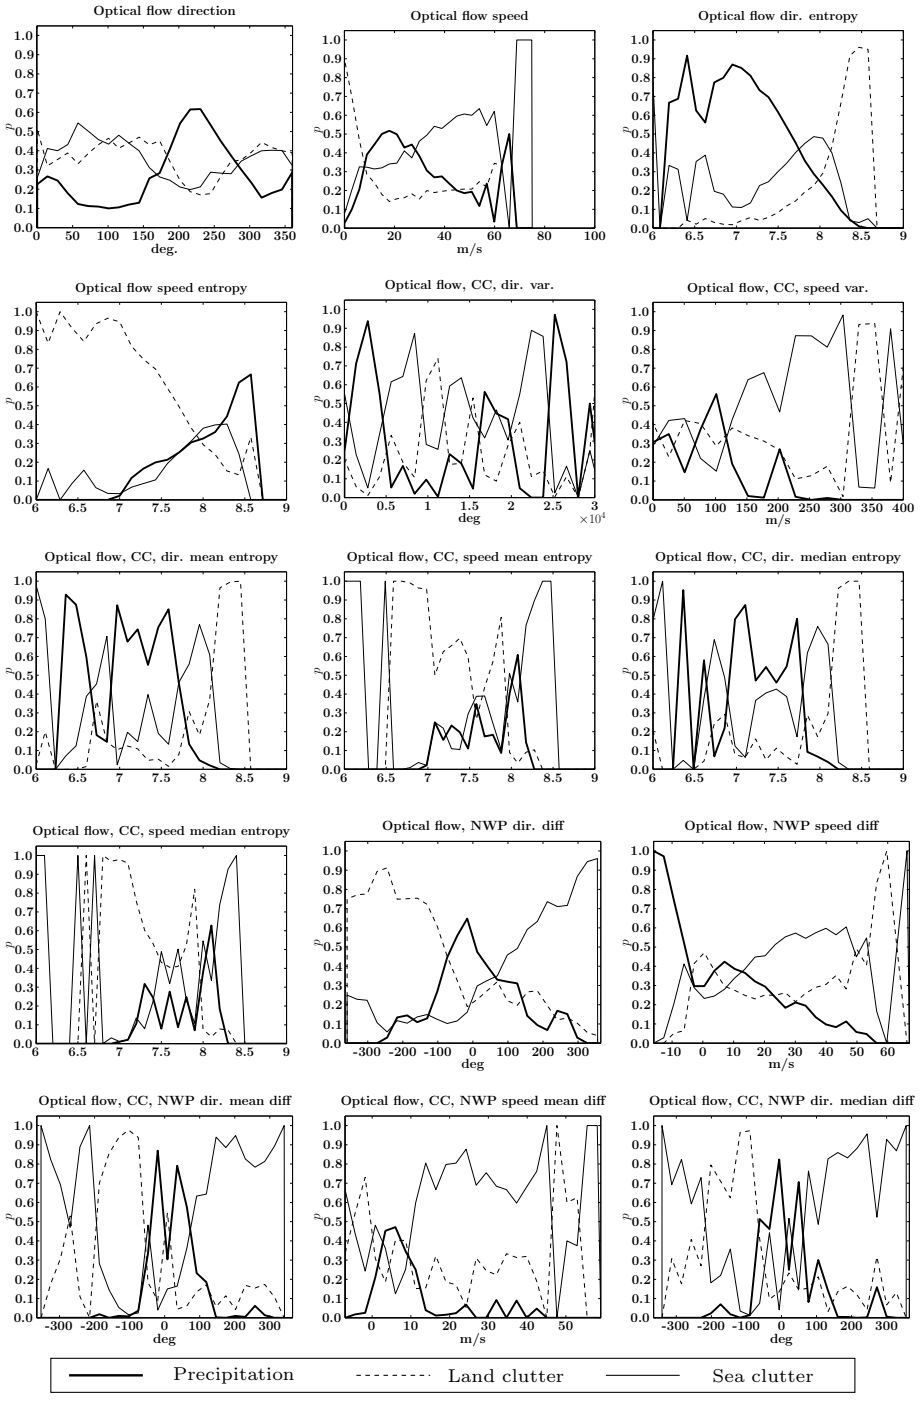

Figure C.18: Spatio-temporal features. All echoes. Three classes.

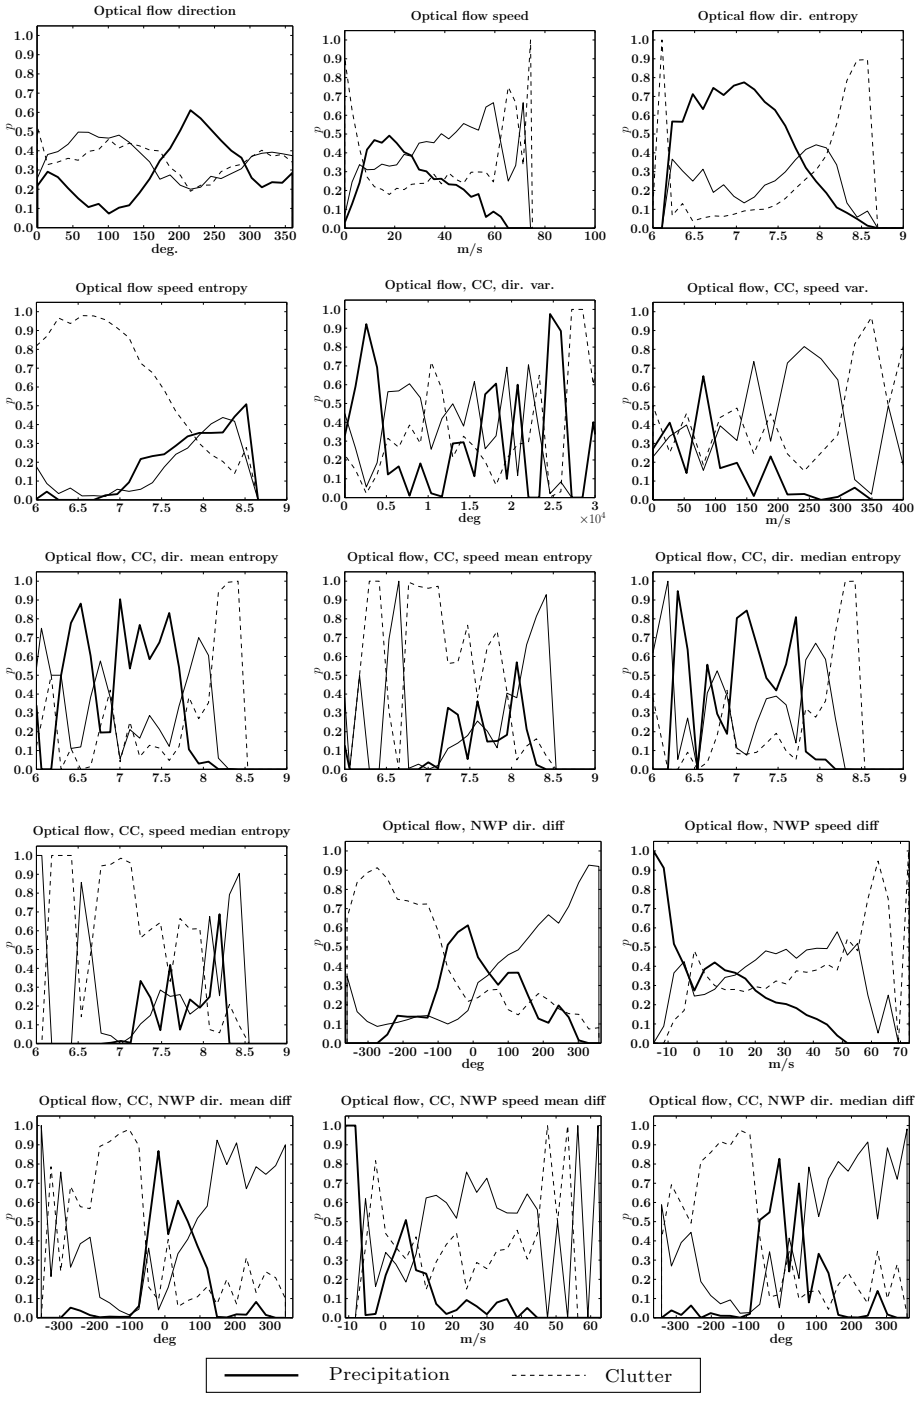

Figure C.19: Spatio-temporal features. Weak echoes. Three classes.

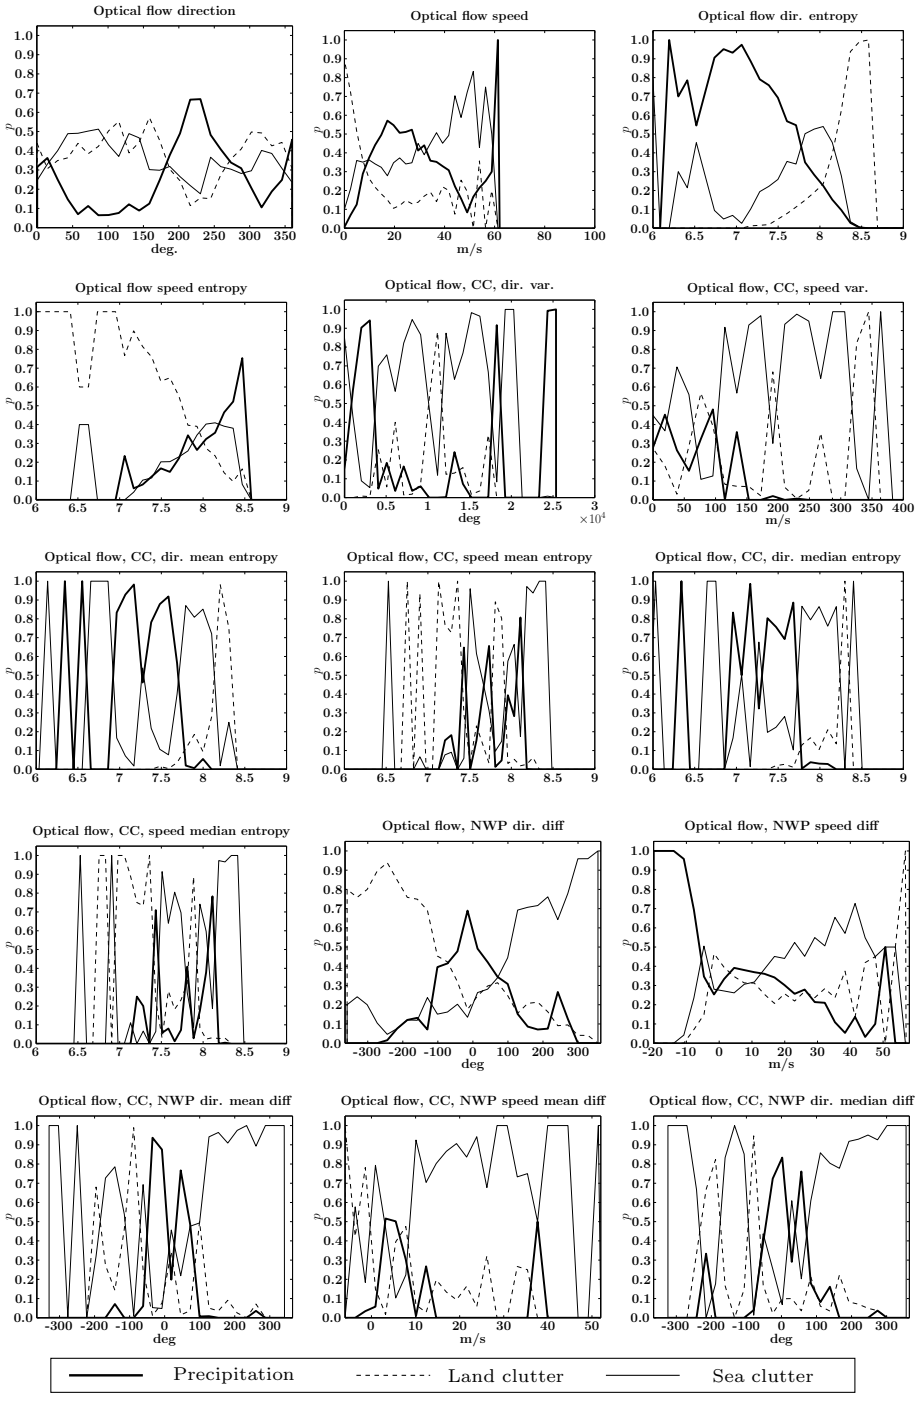

Figure C.20: Spatio-temporal features. Strong echoes. Three classes.

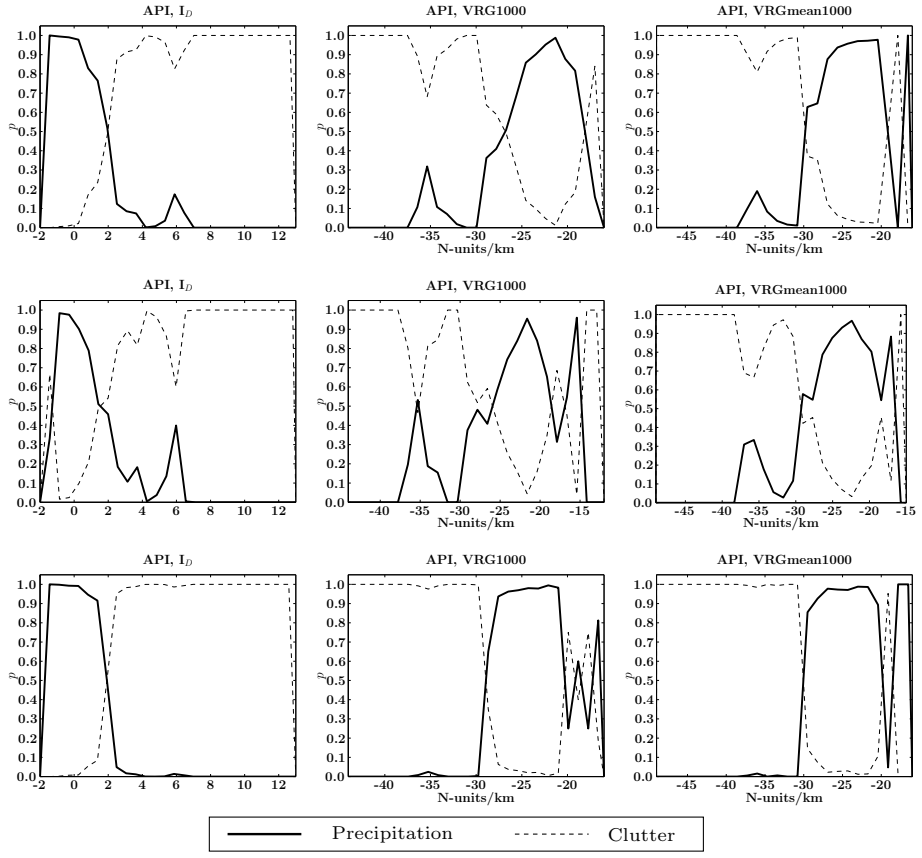

Figure C.21: Propagation method. All echoes (top row), weak echoes (middle row) and strong echoes (bottom row). Two classes.

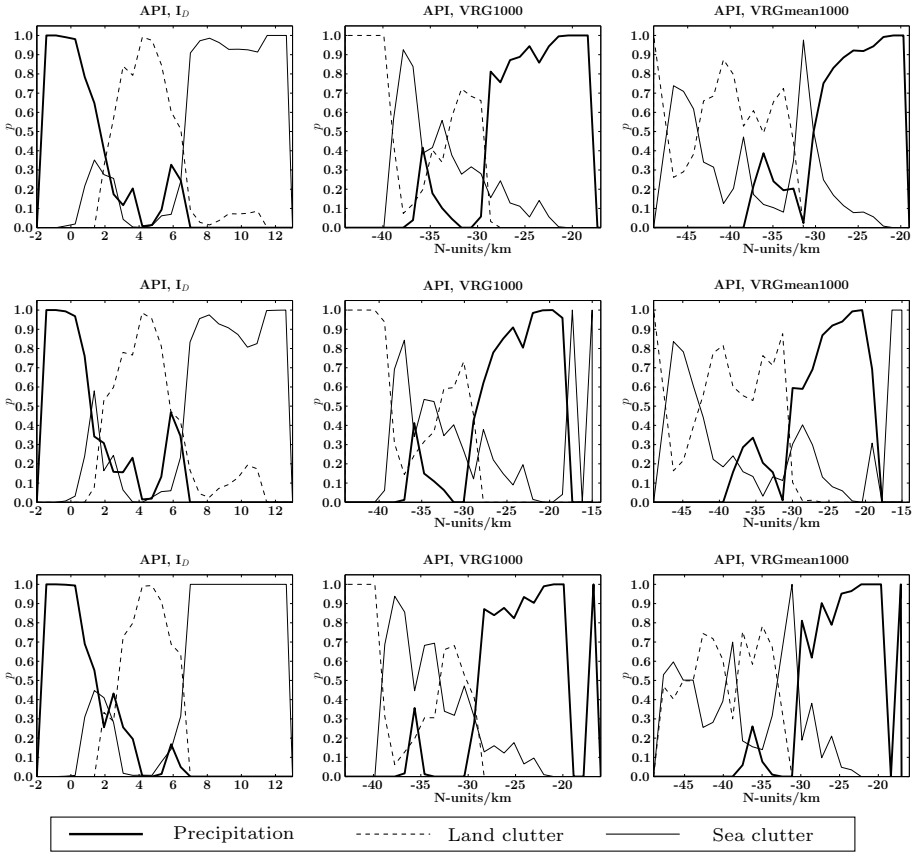

Figure C.22: Propagation method. All echoes (top row), weak echoes (middle row) and strong echoes (bottom row). Three classes.

Appendix D

Feature selection results

Feature selection

This appendix contains the results of feature selection.

On each page, the top table lists the highest classification accuracy in separation of clutter and precipitation for increasing number of features. In italics are the best classification accuracy obtained and in bold is the chosen accuracy which is either the best accuracy or the one with the fewest features closer then 0.2 % to the highest accuracy. The numbers in parenthesis are the features which are listed in the lower left table (in bold are the feature numbers which were chosen). The finally chosen feature subset combination from the exhaustive search is shown in bold in the right hand column.

The graph on the lower right shows the feature selection curves of each feature selection method.

Table D.1: Data fusion method, day features, all echoes, three classes. Feature selection results. See page 187 for full description.

# of feats.	Classification accuracy (Feature combination)			
	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	96.6 % (8)	96.5 % (9)	(8)	96.6 % (8)
2	96.9 % (8 10)	97.0 % (9 10)	(2 8)	97.1 % (1 4)
3	97.2 % (8 9 10)	97.3 % (5 9 10)	(1 2 8)	97.3 % (5 9 10)
4	97.5 % (5 8 9 10)	97.5 % (5 6 9 10)	(1 2 6 8)	97.5 % (5 6 9 10)
5	97.8 % (5 6 8 9 10)	97.8 % (5 6 8 9 10)	(1 2 3 6 8)	97.8 % (5 6 8 9 10)
6	98.0 % (1 5 6 8 9 10)	98.0 % (1 5 6 8 9 10)	(1 2 3 4 6 8)	98.0 % (1 5 6 8 9 10)
7	**98.1** % (1 4 5 6 8 9 10)	**98.1** % (1 3 5 6 8 9 10)	(1 2 3 4 6 7 8)	**98.1 % (1 4 5 6 8 9 10)**
8	98.1 % (1 2 4 5 6 8 9 10)	98.2 % (1 2 3 5 6 8 9 10)	(1 2 3 4 6 7 8 9)	98.2 % (1 2 3 5 6 8 9 10)
9	98.1 % (1 2 3 4 5 6 8 9 10)	98.2 % (1 2 3 5 6 8 9 10 11)	(1 2 3 4 6 7 8 9 10)	98.2 % (1 2 3 5 6 8 9 10 11)
10	*98.2* % (1 2 3 4 5 6 8 9 10 11)	*98.2* % (1 2 3 4 5 6 8 9 10 11)	(1 2 3 4 6 7 8 9 10 11)	*98.2 % (1 2 3 4 5 6 8 9 10 11)*
11	98.2 % (1 2 3 4 5 6 7 8 9 10 11)	98.2 % (1 2 3 4 5 6 7 8 9 10 11)	(1 2 3 4 5 6 7 8 9 10 11)	98.2 % (1 2 3 4 5 6 7 8 9 10 11)

Feat. #	Feature name
1	MSG Ch 01 VIS0.6
2	MSG Ch 02 VIS0.8
3	MSG Ch 03 IR1.6
4	MSG Ch 04 IR3.9
5	MSG Ch 05 WV6.2
6	MSG Ch 06 WV7.3
7	MSG Ch 07 IR8.7
8	MSG Ch 08 IR9.7
9	MSG Ch 09 IR10.8
10	MSG Ch 10 IR12.0
11	MSG Ch 11 IR13.4

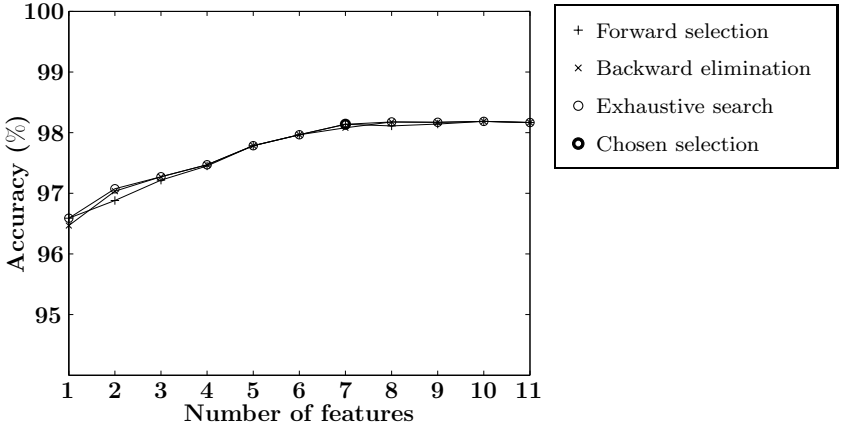

Table D.2: Data fusion method, day features, all echoes, two classes. Feature selection results. See page 187 for full description.

# of	Classification accuracy (Feature combination)			
feats.	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	95.6 % (8)	95.0 % (10)	(8)	95.6 % (8)
2	96.2 % (5 8)	96.1 % (5 10)	(1 10)	96.2 % (5 8)
3	96.2 % (5 8 9)	96.5 % (5 9 10)	(1 2 10)	96.5 % (5 9 10)
4	96.6 % (5 8 9 10)	96.8 % (5 6 9 10)	(1 2 10 11)	96.8 % (5 6 9 10)
5	**96.8** % (5 6 8 9 10)	**96.9** % (1 5 6 9 10)	(1 2 5 6 10)	**96.9 % (1 5 6 9 10)**
6	96.8 % (1 5 6 8 9 10)	*97.0* % (1 3 5 6 9 10)	(1 2 3 5 6 10)	*97.0 % (1 3 5 6 9 10)*
7	*96.9* % (1 4 5 6 8 9 10)	97.0 % (1 2 3 5 6 9 10)	(1 2 3 5 6 7 9)	97.0 % (1 2 3 5 6 9 10)
8	96.8 % (1 4 5 6 7 8 9 10)	97.0 % (1 2 3 5 6 8 9 10)	(1 2 3 4 5 6 7 9)	97.0 % (1 2 3 5 6 7 9 10)
9	96.7 % (1 4 5 6 7 8 9 10 11)	97.0 % (1 2 3 5 6 8 9 10 11)	(1 2 3 5 6 7 9 10 11)	97.0 % (1 2 3 5 6 8 9 10 11)
10	96.7 % (1 3 4 5 6 7 8 9 10 11)	96.9 % (1 2 3 5 6 7 8 9 10 11)	(1 2 3 5 6 7 8 9 10 11)	96.9 % (1 2 3 5 6 7 8 9 10 11)
11	96.8 % (1 2 3 4 5 6 7 8 9 10 11)	96.8 % (1 2 3 4 5 6 7 8 9 10 11)	(1 2 3 4 5 6 7 8 9 10 11)	96.8 % (1 2 3 4 5 6 7 8 9 10 11)

Feat. #	**Feature name**
1	MSG Ch 01 VIS0.6
2	MSG Ch 02 VIS0.8
3	MSG Ch 03 IR1.6
4	MSG Ch 04 IR3.9
5	MSG Ch 05 WV6.2
6	MSG Ch 06 WV7.3
7	MSG Ch 07 IR8.7
8	MSG Ch 08 IR9.7
9	MSG Ch 09 IR10.8
10	MSG Ch 10 IR12.0
11	MSG Ch 11 IR13.4

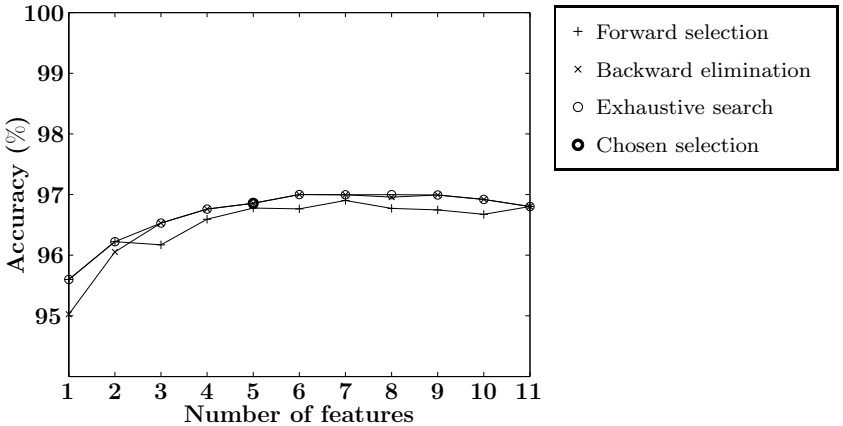

Table D.3: Data fusion method, day features, weak echoes, three classes. Feature selection results. See page 187 for full description.

# of feats.	Classification accuracy (Feature combination)			
	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	94.9 % (9)	92.9 % (4)	(9)	94.9 % (9)
2	95.3 % (9 10)	95.3 % (1 4)	(1 2)	95.4 % (1 8)
3	96.2 % (8 9 10)	96.3 % (1 4 10)	(1 2 7)	96.3 % (1 4 9)
4	96.6 % (8 9 10 11)	97.0 % (1 4 8 10)	(1 2 3 7)	97.0 % (1 4 8 10)
5	97.0 % (5 8 9 10 11)	97.5 % (1 4 8 9 10)	(1 2 3 7 8)	97.5 % (1 4 8 9 10)
6	97.3 % (1 5 8 9 10 11)	97.7 % (1 4 5 8 9 10)	(1 2 3 6 7 8)	97.7 % (1 3 8 9 10 11)
7	97.9 % (1 4 5 8 9 10 11)	97.9 % (1 4 5 8 9 10 11)	(1 2 3 6 8 10 11)	97.9 % (1 4 5 8 9 10 11)
8	97.9 % (1 3 4 5 8 9 10 11)	97.9 % (1 3 4 5 8 9 10 11)	(1 2 3 6 7 8 10 11)	97.9 % (1 3 4 5 8 9 10 11)
9	**98.0** % (1 2 3 4 5 8 9 10 11)	**98.0** % (1 2 3 4 5 8 9 10 11)	(1 2 3 6 7 8 9 10 11)	**98.0 % (1 2 3 4 5 8 9 10 11)**
10	98.1 % (1 2 3 4 5 6 8 9 10 11)	98.1 % (1 2 3 4 5 6 8 9 10 11)	(1 2 3 4 6 7 8 9 10 11)	98.1 % (1 2 3 4 5 6 8 9 10 11)
11	*98.2* % (1 2 3 4 5 6 7 8 9 10 11)	*98.2* % (1 2 3 4 5 6 7 8 9 10 11)	(1 2 3 4 5 6 7 8 9 10 11)	*98.2 % (1 2 3 4 5 6 7 8 9 10 11)*

Feat. #	Feature name
1	MSG Ch 01 VIS0.6
2	MSG Ch 02 VIS0.8
3	MSG Ch 03 IR1.6
4	MSG Ch 04 IR3.9
5	MSG Ch 05 WV6.2
6	MSG Ch 06 WV7.3
7	MSG Ch 07 IR8.7
8	MSG Ch 08 IR9.7
9	MSG Ch 09 IR10.8
10	MSG Ch 10 IR12.0
11	MSG Ch 11 IR13.4

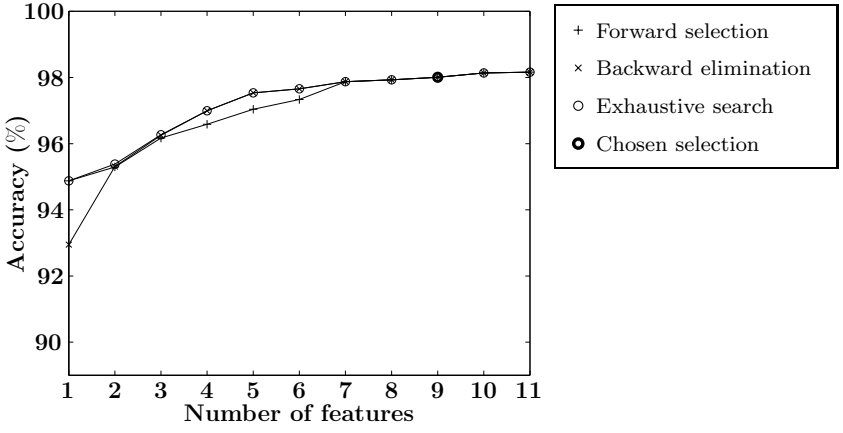

Table D.4: Data fusion method, day features, weak echoes, two classes. Feature selection results. See page 187 for full description.

# of	Classification accuracy (Feature combination)			
feats.	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	93.8 % (8)	93.7 % (9)	(9)	93.8 % (8)
2	93.7 % (6 8)	94.6 % (9 10)	(1 9)	94.6 % (9 10)
3	93.5 % (6 8 9)	94.8 % (5 9 10)	(1 2 9)	94.8 % (5 9 10)
4	94.4 % (6 8 9 10)	95.1 % (1 5 9 10)	(1 2 3 9)	95.1 % (1 5 9 10)
5	94.8 % (6 8 9 10 11)	95.1 % (1 5 9 10 11)	(1 2 6 10 11)	95.3 % (5 6 9 10 11)
6	95.3 % (5 6 8 9 10 11)	95.3 % (1 4 5 9 10 11)	(1 2 6 8 10 11)	95.3 % (1 4 5 9 10 11)
7	95.4 % (1 5 6 8 9 10 11)	***96.0*** % (1 4 5 8 9 10 11)	(1 2 3 6 8 10 11)	***96.0 % (1 4 5 8 9 10 11)***
8	***95.8*** % (1 4 5 6 8 9 10 11)	95.8 % (1 2 4 5 8 9 10 11)	(1 2 3 6 7 8 10 11)	95.8 % (1 4 5 6 8 9 10 11)
9	95.7 % (1 4 5 6 7 8 9 10 11)	95.7 % (1 2 4 5 7 8 9 10 11)	(1 2 3 6 7 8 9 10 11)	95.7 % (1 4 5 6 7 8 9 10 11)
10	95.7 % (1 2 4 5 6 7 8 9 10 11)	95.8 % (1 2 3 4 5 7 8 9 10 11)	(1 2 3 5 6 7 8 9 10 11)	95.8 % (1 2 3 4 5 7 8 9 10 11)
11	95.8 % (1 2 3 4 5 6 7 8 9 10 11)	95.8 % (1 2 3 4 5 6 7 8 9 10 11)	(1 2 3 4 5 6 7 8 9 10 11)	95.8 % (1 2 3 4 5 6 7 8 9 10 11)

Feat. #	Feature name
1	MSG Ch 01 VIS0.6
2	MSG Ch 02 VIS0.8
3	MSG Ch 03 IR1.6
4	MSG Ch 04 IR3.9
5	MSG Ch 05 WV6.2
6	MSG Ch 06 WV7.3
7	MSG Ch 07 IR8.7
8	MSG Ch 08 IR9.7
9	MSG Ch 09 IR10.8
10	MSG Ch 10 IR12.0
11	MSG Ch 11 IR13.4

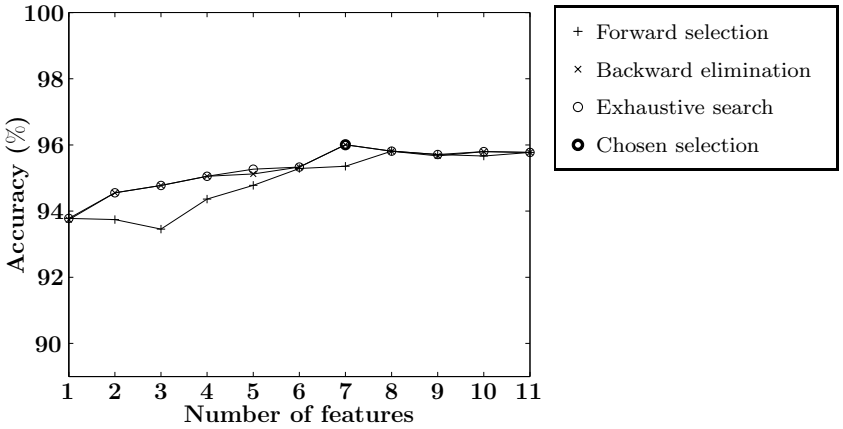

Table D.5: Data fusion method, day features, strong echoes, three classes. Feature selection results. See page 187 for full description.

# of feats.	Classification accuracy (Feature combination)			
	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	98.1 % (8)	95.3 % (6)	(11)	98.1 % (8)
2	98.7 % (5 8)	98.5 % (5 6)	(6 8)	98.7 % (1 4)
3	98.8 % (5 7 8)	98.7 % (2 5 6)	(2 6 8)	98.8 % (5 7 8)
4	98.7 % (2 5 7 8)	**98.9** % (2 3 5 6)	(1 2 6 8)	**98.9 % (2 3 5 8)**
5	**98.9** % (2 3 5 7 8)	98.9 % (2 3 5 6 10)	(1 2 4 6 8)	98.9 % (2 3 5 6 8)
6	98.9 % (2 3 5 6 7 8)	98.9 % (2 3 5 6 10 11)	(1 2 4 6 7 8)	*99.1 % (2 3 5 6 8 11)*
7	98.9 % (2 3 5 6 7 8 11)	99.0 % (2 3 5 6 9 10 11)	(1 2 4 6 7 8 9)	99.0 % (2 3 5 6 9 10 11)
8	99.0 % (2 3 5 6 7 8 10 11)	*99.0* % (2 3 5 6 8 9 10 11)	(1 2 3 4 6 8 9 11)	99.0 % (2 3 5 6 8 9 10 11)
9	*99.0* % (2 3 5 6 7 8 9 10 11)	99.0 % (1 2 3 5 6 8 9 10 11)	(1 2 3 4 6 7 8 9 11)	99.0 % (1 2 3 5 6 8 9 10 11)
10	98.9 % (1 2 3 5 6 7 8 9 10 11)	98.9 % (1 2 3 5 6 7 8 9 10 11)	(1 2 3 4 5 6 7 8 9 11)	98.9 % (1 2 3 5 6 7 8 9 10 11)
11	98.9 % (1 2 3 4 5 6 7 8 9 10 11)	98.9 % (1 2 3 4 5 6 7 8 9 10 11)	(1 2 3 4 5 6 7 8 9 10 11)	98.9 % (1 2 3 4 5 6 7 8 9 10 11)

Feat. #	Feature name
1	MSG Ch 01 VIS0.6
2	MSG Ch 02 VIS0.8
3	MSG Ch 03 IR1.6
4	MSG Ch 04 IR3.9
5	MSG Ch 05 WV6.2
6	MSG Ch 06 WV7.3
7	MSG Ch 07 IR8.7
8	MSG Ch 08 IR9.7
9	MSG Ch 09 IR10.8
10	MSG Ch 10 IR12.0
11	MSG Ch 11 IR13.4

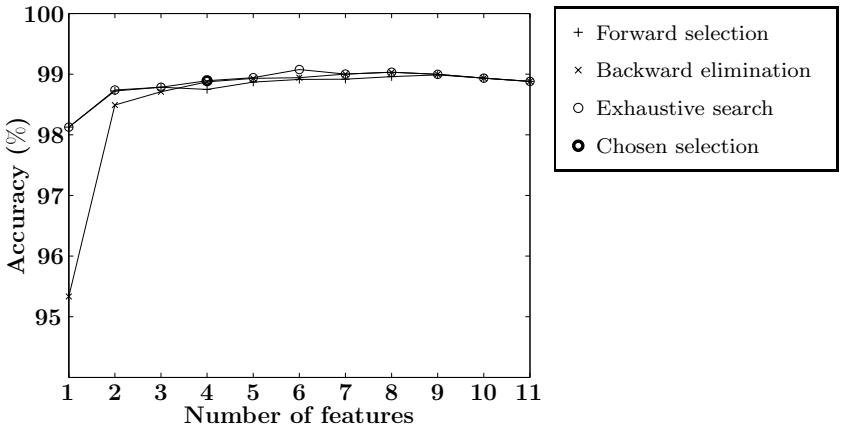

Table D.6: Data fusion method, day features, strong echoes, two classes. Feature selection results. See page 187 for full description.

# of	Classification accuracy (Feature combination)			
feats.	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	98.3 % (8)	98.1 % (10)	(9)	98.3 % (8)
2	98.8 % (5 8)	98.5 % (5 10)	(1 9)	98.8 % (5 8)
3	**98.8** % (1 5 8)	98.7 % (1 5 10)	(1 2 9)	**98.9 % (1 4 5)**
4	98.9 % (1 5 6 8)	98.8 % (1 5 9 10)	(1 5 6 9)	99.0 % (1 4 5 6)
5	*99.0* % (1 3 5 6 8)	**98.9** % (1 5 8 9 10)	(1 5 6 7 9)	99.0 % (1 3 5 6 8)
6	99.0 % (1 3 5 6 8 11)	99.0 % (1 3 5 8 9 10)	(1 2 5 6 7 9)	99.0 % (1 3 5 8 9 10)
7	98.9 % (1 3 5 6 7 8 11)	*99.0* % (1 3 5 6 8 9 10)	(1 2 5 6 7 8 9)	*99.0 % (1 3 5 6 8 9 10)*
8	98.9 % (1 3 5 6 7 8 10 11)	99.0 % (1 3 5 6 8 9 10 11)	(1 2 3 5 6 7 8 9)	99.0 % (1 3 5 6 8 9 10 11)
9	99.0 % (1 3 5 6 7 8 9 10 11)	99.0 % (1 3 5 6 7 8 9 10 11)	(1 2 3 4 5 6 7 8 9)	99.0 % (1 3 5 6 7 8 9 10 11)
10	98.9 % (1 2 3 5 6 7 8 9 10 11)	98.9 % (1 2 3 5 6 7 8 9 10 11)	(1 2 3 5 6 7 8 9 10 11)	98.9 % (1 2 3 5 6 7 8 9 10 11)
11	98.8 % (1 2 3 4 5 6 7 8 9 10 11)	98.8 % (1 2 3 4 5 6 7 8 9 10 11)	(1 2 3 4 5 6 7 8 9 10 11)	98.8 % (1 2 3 4 5 6 7 8 9 10 11)

Feat. #	Feature name
1	MSG Ch 01 VIS0.6
2	MSG Ch 02 VIS0.8
3	MSG Ch 03 IR1.6
4	MSG Ch 04 IR3.9
5	MSG Ch 05 WV6.2
6	MSG Ch 06 WV7.3
7	MSG Ch 07 IR8.7
8	MSG Ch 08 IR9.7
9	MSG Ch 09 IR10.8
10	MSG Ch 10 IR12.0
11	MSG Ch 11 IR13.4

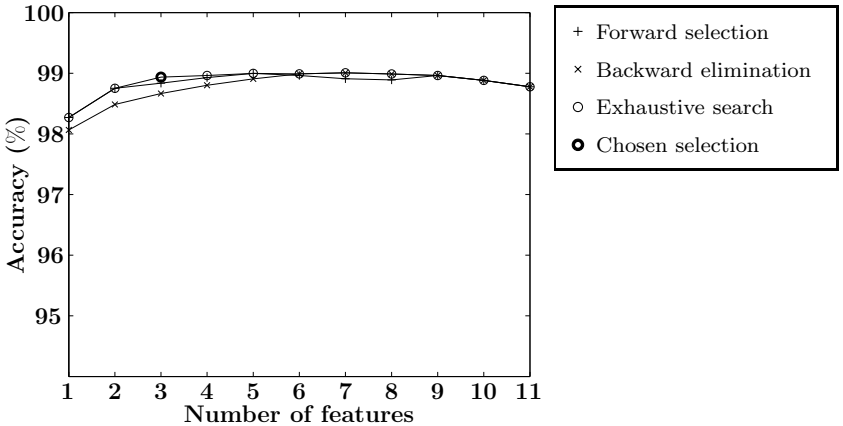

Table D.7: Data fusion method, night features, all echoes, three classes. Feature selection results. See page 187 for full description.

# of feats.	Classification accuracy (Feature combination)			
	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	77.0 % (4)	75.6 % (7)	(6)	77.0 % (4)
2	94.1 % (4 5)	94.0 % (3 7)	(4 5)	94.1 % (4 5)
3	***98.0*** % (4 5 6)	***99.0*** % (3 5 7)	(3 4 5)	***99.0 % (3 5 7)***
4	97.7 % (2 4 5 6)	98.9 % (3 5 6 7)	(3 4 5 6)	98.9 % (3 5 6 7)
5	97.4 % (2 3 4 5 6)	98.6 % (2 3 5 6 7)	(2 3 4 5 6)	98.6 % (2 3 5 6 7)
6	97.6 % (2 3 4 5 6 7)	98.3 % (2 3 5 6 7 8)	(2 3 4 5 6 7)	98.3 % (2 3 5 6 7 8)
7	97.2 % (2 3 4 5 6 7 8)	97.7 % (1 2 3 5 6 7 8)	(2 3 4 5 6 7 8)	97.7 % (1 2 3 5 6 7 8)
8	97.0 % (1 2 3 4 5 6 7 8)	97.0 % (1 2 3 4 5 6 7 8)	(1 2 3 4 5 6 7 8)	97.0 % (1 2 3 4 5 6 7 8)

Feat. #	Feature name
1	MSG Ch 04 IR3.9
2	MSG Ch 05 WV6.2
3	MSG Ch 06 WV7.3
4	MSG Ch 07 IR8.7
5	MSG Ch 08 IR9.7
6	MSG Ch 09 IR10.8
7	MSG Ch 10 IR12.0
8	MSG Ch 11 IR13.4

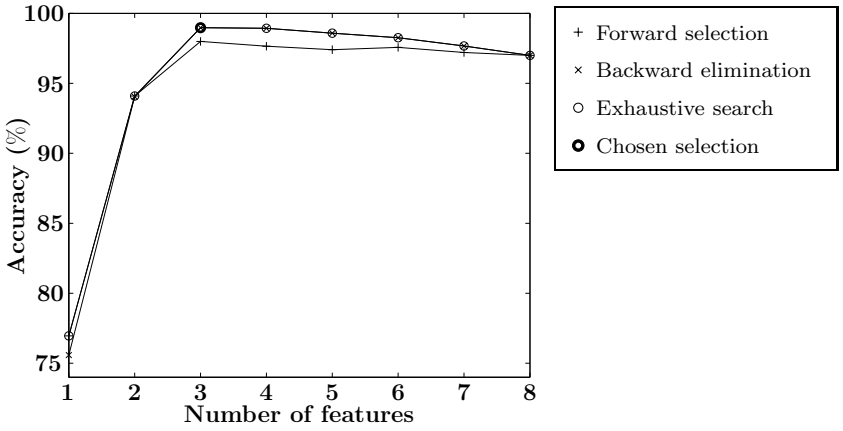

Table D.8: Data fusion method, night features, all echoes, two classes. Feature selection results. See page 187 for full description.

# of	Classification accuracy (Feature combination)			
feats.	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	81.6 % (4)	80.1 % (7)	(4)	81.6 % (4)
2	93.2 % (4 8)	93.6 % (7 8)	(7 8)	93.6 % (7 8)
3	94.1 % (4 7 8)	95.8 % (5 7 8)	(3 7 8)	95.8 % (5 7 8)
4	**97.4** % (4 5 7 8)	**97.4** % (4 5 7 8)	(1 3 7 8)	**97.4 % (4 5 7 8)**
5	*97.6* % (4 5 6 7 8)	*97.6* % (4 5 6 7 8)	(1 3 5 7 8)	*97.6 % (4 5 6 7 8)*
6	97.5 % (2 4 5 6 7 8)	97.5 % (2 4 5 6 7 8)	(1 3 4 5 7 8)	97.5 % (2 4 5 6 7 8)
7	97.1 % (2 3 4 5 6 7 8)	97.1 % (2 3 4 5 6 7 8)	(1 2 3 4 5 7 8)	97.1 % (2 3 4 5 6 7 8)
8	96.4 % (1 2 3 4 5 6 7 8)	96.4 % (1 2 3 4 5 6 7 8)	(1 2 3 4 5 6 7 8)	96.4 % (1 2 3 4 5 6 7 8)

Feat. #	**Feature name**
1	MSG Ch 04 IR3.9
2	MSG Ch 05 WV6.2
3	MSG Ch 06 WV7.3
4	MSG Ch 07 IR8.7
5	MSG Ch 08 IR9.7
6	MSG Ch 09 IR10.8
7	MSG Ch 10 IR12.0
8	MSG Ch 11 IR13.4

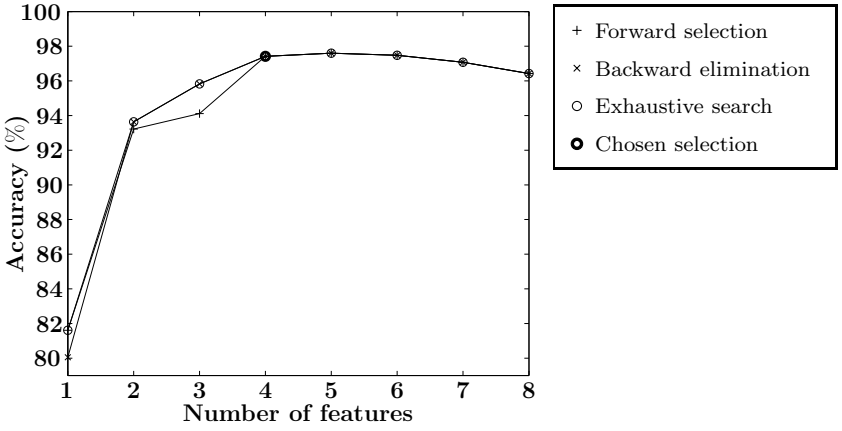

Table D.9: Data fusion method, night features, weak echoes, three classes. Feature selection results. See page 187 for full description.

# of feats.	Classification accuracy (Feature combination)			
	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	85.7 % (4)	84.7 % (7)	(6)	85.7 % (4)
2	96.5 % (4 5)	96.6 % (3 7)	(4 5)	96.6 % (3 6)
3	***98.9*** % (4 5 6)	**99.2** % (3 5 7)	(2 4 5)	**99.2 % (3 5 7)**
4	98.6 % (2 4 5 6)	*99.4* % (1 3 5 7)	(2 4 5 7)	*99.4 % (1 3 5 7)*
5	98.7 % (1 2 4 5 6)	99.4 % (1 3 5 6 7)	(2 4 5 6 7)	99.4 % (1 3 5 6 7)
6	98.7 % (1 2 4 5 6 7)	99.3 % (1 2 3 5 6 7)	(2 3 4 5 6 7)	99.3 % (1 2 3 5 6 7)
7	98.8 % (1 2 3 4 5 6 7)	99.2 % (1 2 3 5 6 7 8)	(2 3 4 5 6 7 8)	99.2 % (1 2 3 5 6 7 8)
8	98.8 % (1 2 3 4 5 6 7 8)	98.8 % (1 2 3 4 5 6 7 8)	(1 2 3 4 5 6 7 8)	98.8 % (1 2 3 4 5 6 7 8)

Feat. #	Feature name
1	MSG Ch 04 IR3.9
2	MSG Ch 05 WV6.2
3	MSG Ch 06 WV7.3
4	MSG Ch 07 IR8.7
5	MSG Ch 08 IR9.7
6	MSG Ch 09 IR10.8
7	MSG Ch 10 IR12.0
8	MSG Ch 11 IR13.4

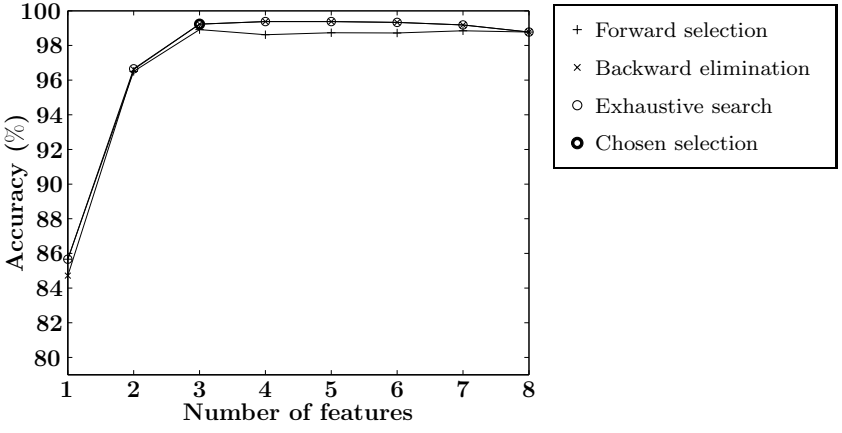

Table D.10: Data fusion method, night features, weak echoes, two classes. Feature selection results. See page 187 for full description.

# of	Classification accuracy (Feature combination)			
feats.	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	89.2 % (4)	88.3 % (7)	(4)	89.2 % (4)
2	95.2 % (4 8)	95.3 % (7 8)	(7 8)	95.3 % (7 8)
3	95.6 % (4 5 8)	96.7 % (5 7 8)	(3 7 8)	96.7 % (5 7 8)
4	97.2 % (4 5 7 8)	97.2 % (4 5 7 8)	(1 3 7 8)	97.2 % (4 5 7 8)
5	***97.4*** % (4 5 6 7 8)	***97.4*** % (4 5 6 7 8)	(1 3 5 7 8)	***97.4 % (4 5 6 7 8)***
6	97.3 % (2 4 5 6 7 8)	97.3 % (3 4 5 6 7 8)	(1 2 3 5 7 8)	97.3 % (2 4 5 6 7 8)
7	97.1 % (1 2 4 5 6 7 8)	97.3 % (1 3 4 5 6 7 8)	(1 2 3 5 6 7 8)	97.3 % (1 3 4 5 6 7 8)
8	97.1 % (1 2 3 4 5 6 7 8)	97.1 % (1 2 3 4 5 6 7 8)	(1 2 3 4 5 6 7 8)	97.1 % (1 2 3 4 5 6 7 8)

Feat. #	**Feature name**
1	MSG Ch 04 IR3.9
2	MSG Ch 05 WV6.2
3	MSG Ch 06 WV7.3
4	MSG Ch 07 IR8.7
5	MSG Ch 08 IR9.7
6	MSG Ch 09 IR10.8
7	MSG Ch 10 IR12.0
8	MSG Ch 11 IR13.4

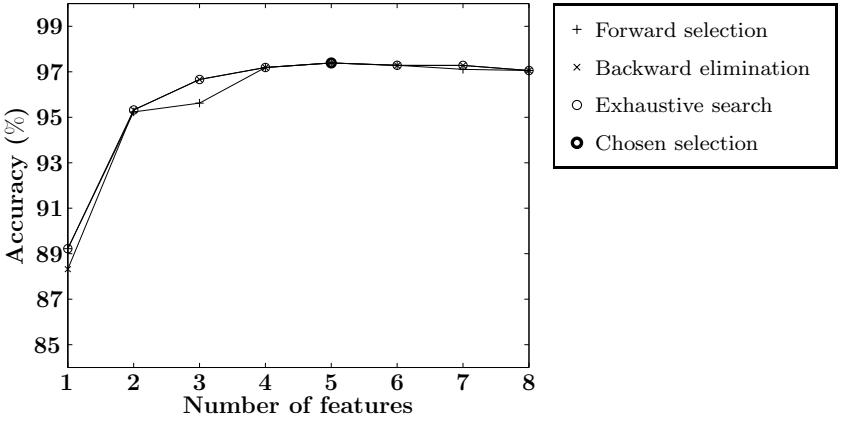

Table D.11: Data fusion method, night features, strong echoes, three classes. Feature selection results. See page 187 for full description.

# of feats.	Classification accuracy (Feature combination)			
	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	96.3 % (2)	94.8 % (3)	(2)	96.3 % (2)
2	84.1 % (2 8)	79.3 % (1 3)	(6 8)	93.8 % (4 5)
3	91.0 % (2 7 8)	83.7 % (1 3 5)	(1 6 8)	96.4 % (5 7 8)
4	96.6 % (2 5 7 8)	95.8 % (1 3 5 8)	(1 5 6 8)	98.5 % (2 3 5 7)
5	99.0 % (2 3 5 7 8)	98.9 % (1 2 3 5 8)	(1 3 4 5 6)	99.1 % (1 2 3 5 7)
6	99.3 % (1 2 3 5 7 8)	99.3 % (1 2 3 5 6 8)	(1 3 4 5 6 8)	99.3 % (1 2 3 5 6 8)
7	**99.6** % (1 2 3 4 5 7 8)	**99.6** % (1 2 3 4 5 6 8)	(1 2 3 4 5 6 8)	**99.6 % (1 2 3 4 5 6 8)**
8	*99.7* % (1 2 3 4 5 6 7 8)	*99.7* % (1 2 3 4 5 6 7 8)	(1 2 3 4 5 6 7 8)	*99.7 % (1 2 3 4 5 6 7 8)*

Feat. #	Feature name
1	MSG Ch 04 IR3.9
2	MSG Ch 05 WV6.2
3	MSG Ch 06 WV7.3
4	MSG Ch 07 IR8.7
5	MSG Ch 08 IR9.7
6	MSG Ch 09 IR10.8
7	MSG Ch 10 IR12.0
8	MSG Ch 11 IR13.4

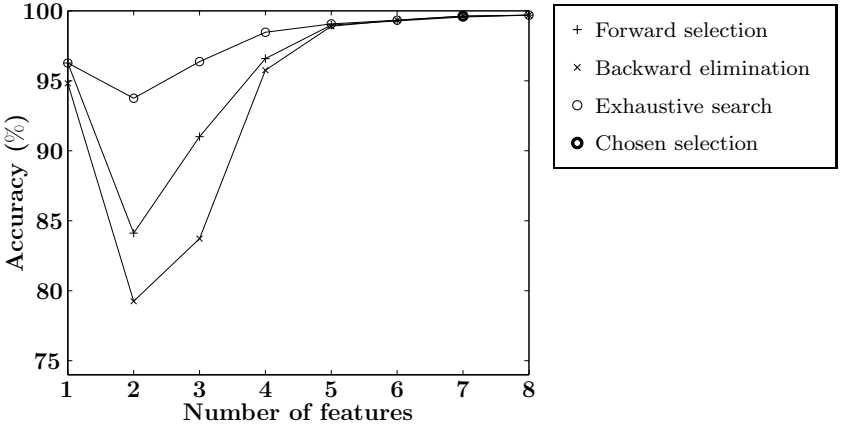

Table D.12: Data fusion method, night features, strong echoes, two classes. Feature selection results. See page 187 for full description.

# of	Classification accuracy (Feature combination)			
feats.	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	73.7 % (4)	71.5 % (7)	(6)	73.7 % (4)
2	93.7 % (4 8)	99.2 % (7 8)	(7 8)	99.2 % (7 8)
3	99.4 % (4 7 8)	99.4 % (4 7 8)	(5 7 8)	99.4 % (4 7 8)
4	99.5 % (4 5 7 8)	99.5 % (4 5 7 8)	(3 4 7 8)	99.5 % (4 5 7 8)
5	**99.8** % (2 4 5 7 8)	**99.8** % (2 4 5 7 8)	(3 4 5 7 8)	**99.8 % (2 4 5 7 8)**
6	99.9 % (2 3 4 5 7 8)	99.9 % (2 3 4 5 7 8)	(1 3 4 5 7 8)	99.9 % (2 3 4 5 7 8)
7	*99.9* % (2 3 4 5 6 7 8)	*99.9* % (2 3 4 5 6 7 8)	(1 2 3 4 5 7 8)	*99.9 % (2 3 4 5 6 7 8)*
8	99.8 % (1 2 3 4 5 6 7 8)	99.8 % (1 2 3 4 5 6 7 8)	(1 2 3 4 5 6 7 8)	99.8 % (1 2 3 4 5 6 7 8)

Feat. #	**Feature name**
1	MSG Ch 04 IR3.9
2	MSG Ch 05 WV6.2
3	MSG Ch 06 WV7.3
4	MSG Ch 07 IR8.7
5	MSG Ch 08 IR9.7
6	MSG Ch 09 IR10.8
7	MSG Ch 10 IR12.0
8	MSG Ch 11 IR13.4

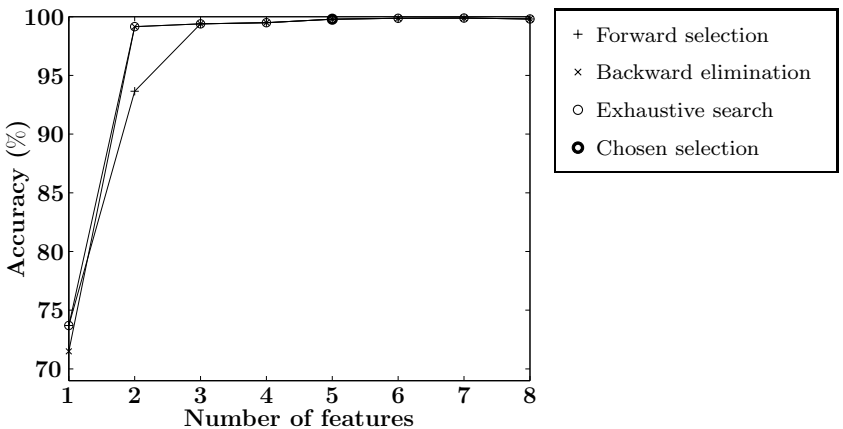

Table D.13: Spatio-temporal method, pixel-based, all echoes, three classes. Feature selection results. See page 187 for full description.

# of	Classification accuracy (Feature combination)			
feats.	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	81.3 % (1)	81.3 % (1)	(3)	81.3 % (1)
2	83.3 % (1 3)	82.5 % (1 4)	(1 3)	83.3 % (1 3)
3	84.8 % (1 3 4)	85.5 % (1 2 4)	(1 2 3)	85.5 % (1 2 4)
4	***86.2*** % (1 2 3 4)	***86.2*** % (1 2 3 4)	(1 2 3 4)	***86.2 % (1 2 3 4)***

Feat. #	**Feature name**
1	Optical flow dir. entropy
2	Optical flow vel. entropy
3	Optical flow, NWP dir. diff.
4	Optical flow, NWP vel. diff.

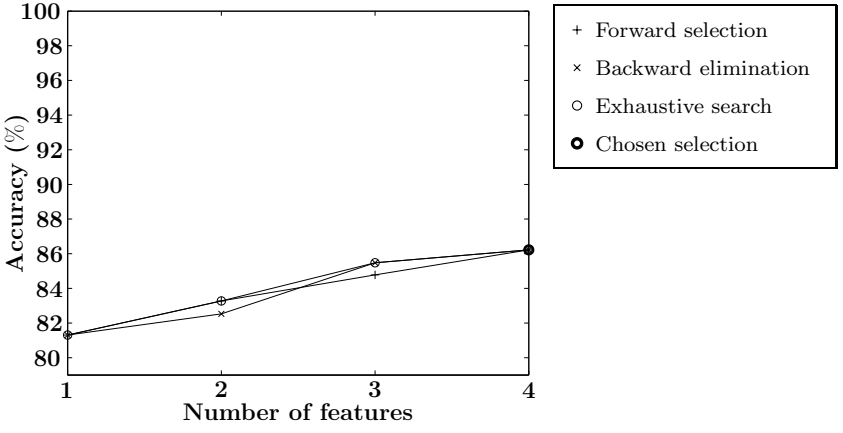

Table D.14: Spatio-temporal method, pixel-based, all echoes, two classes. Feature selection results. See page 187 for full description.

# of	Classification accuracy (Feature combination)			
feats.	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	75.2 % (1)	75.2 % (1)	(1)	75.2 % (1)
2	79.2 % (1 3)	79.2 % (1 3)	(1 2)	79.2 % (1 3)
3	81.0 % (1 3 4)	81.0 % (1 3 4)	(1 2 3)	81.0 % (1 3 4)
4	***81.7*** % (1 2 3 4)	***81.7*** % (1 2 3 4)	(1 2 3 4)	***81.7 % (1 2 3 4)***

Feat. #	**Feature name**
1	Optical flow dir. entropy
2	Optical flow vel. entropy
3	Optical flow, NWP dir. diff.
4	Optical flow, NWP vel. diff.

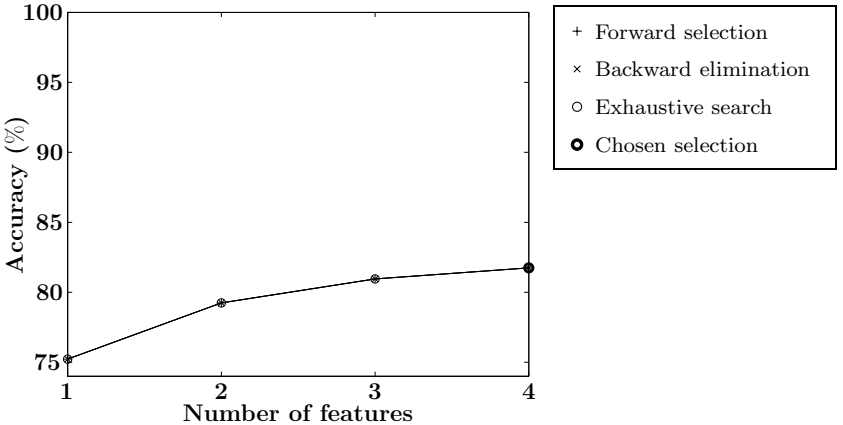

Table D.15: Spatio-temporal method, pixel-based, weak echoes, three classes. Feature selection results. See page 187 for full description.

# of	Classification accuracy (Feature combination)			
feats.	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	80.6 % (1)	80.6 % (1)	(3)	80.6 % (1)
2	81.9 % (1 4)	81.9 % (1 4)	(1 3)	81.9 % (1 4)
3	***84.5*** % (1 2 4)	***84.5*** % (1 2 4)	(1 2 3)	***84.5 % (1 2 4)***
4	83.3 % (1 2 3 4)	83.3 % (1 2 3 4)	(1 2 3 4)	83.3 % (1 2 3 4)

Feat. #	Feature name
1	Optical flow dir. entropy
2	Optical flow vel. entropy
3	Optical flow, NWP dir. diff.
4	Optical flow, NWP vel. diff.

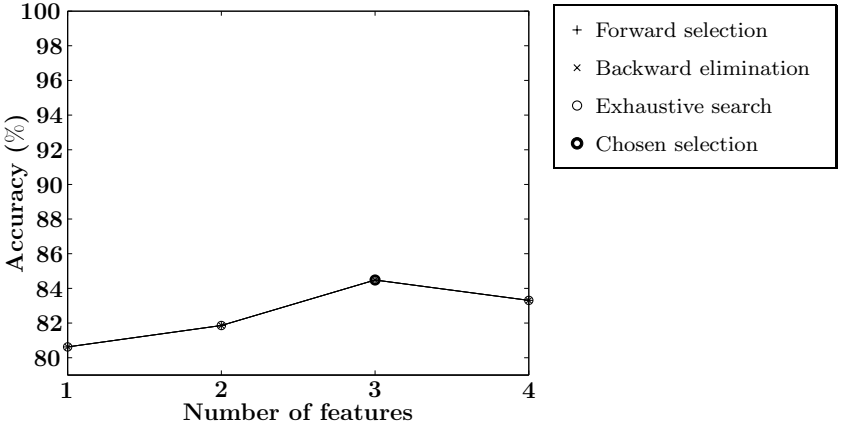

Table D.16: Spatio-temporal method, pixel-based, weak echoes, two classes. Feature selection results. See page 187 for full description.

# of	Classification accuracy (Feature combination)			
feats.	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	75.6 % (1)	75.6 % (1)	(1)	75.6 % (1)
2	78.9 % (1 4)	78.9 % (1 4)	(1 2)	78.9 % (1 4)
3	***80.5*** % (1 3 4)	***80.5*** % (1 3 4)	(1 2 3)	***80.5 % (1 3 4)***
4	77.4 % (1 2 3 4)	77.4 % (1 2 3 4)	(1 2 3 4)	77.4 % (1 2 3 4)

Feat. #	**Feature name**
1	Optical flow dir. entropy
2	Optical flow vel. entropy
3	Optical flow, NWP dir. diff.
4	Optical flow, NWP vel. diff.

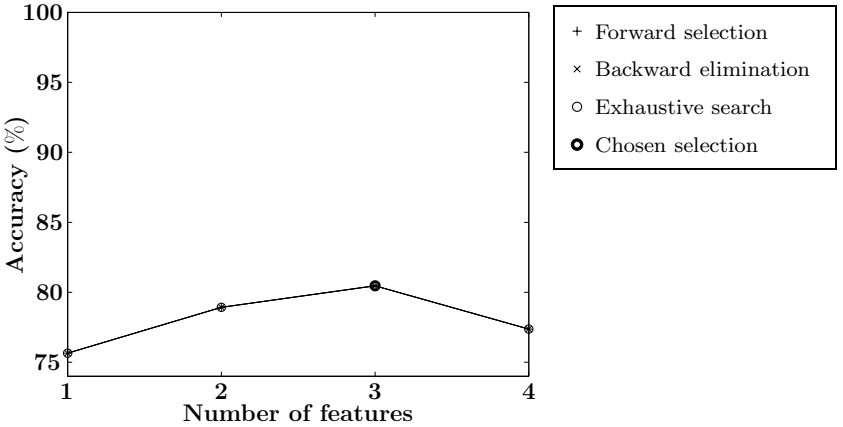

Table D.17: Spatio-temporal method, pixel-based, strong echoes, three classes. Feature selection results. See page 187 for full description.

# of	Classification accuracy (Feature combination)			
feats.	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	81.7 % (1)	81.7 % (1)	(1)	81.7 % (1)
2	85.7 % (1 3)	85.7 % (1 3)	(1 3)	85.7 % (1 3)
3	88.4 % (1 2 3)	88.4 % (1 2 3)	(1 2 3)	88.4 % (1 2 3)
4	***89.4*** % (1 2 3 4)	***89.4*** % (1 2 3 4)	(1 2 3 4)	***89.4 % (1 2 3 4)***

Feat. #	Feature name
1	Optical flow dir. entropy
2	Optical flow vel. entropy
3	Optical flow, NWP dir. diff.
4	Optical flow, NWP vel. diff.

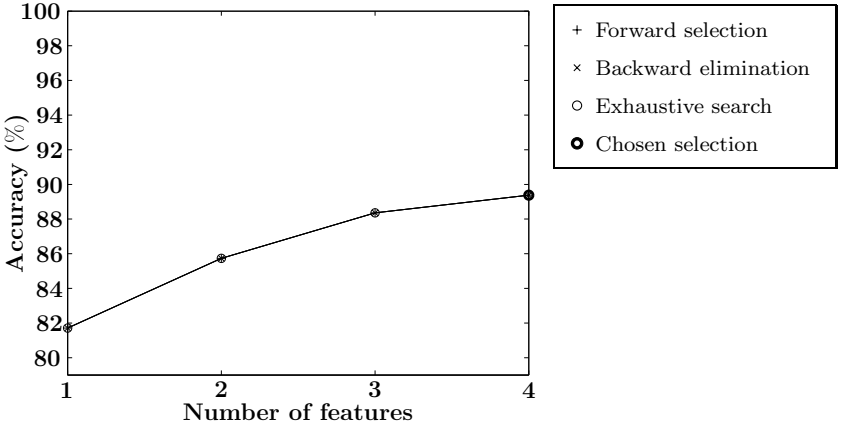

Table D.18: Spatio-temporal method, pixel-based, strong echoes, two classes. Feature selection results. See page 187 for full description.

# of	Classification accuracy (Feature combination)			
feats.	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	78.0 % (1)	78.0 % (1)	(1)	78.0 % (1)
2	85.5 % (1 2)	85.5 % (1 2)	(1 2)	85.5 % (1 2)
3	87.4 % (1 2 3)	87.4 % (1 2 3)	(1 2 3)	87.4 % (1 2 3)
4	***88.1*** % (1 2 3 4)	***88.1*** % (1 2 3 4)	(1 2 3 4)	***88.1 % (1 2 3 4)***

Feat. #	**Feature name**
1	Optical flow dir. entropy
2	Optical flow vel. entropy
3	Optical flow, NWP dir. diff.
4	Optical flow, NWP vel. diff.

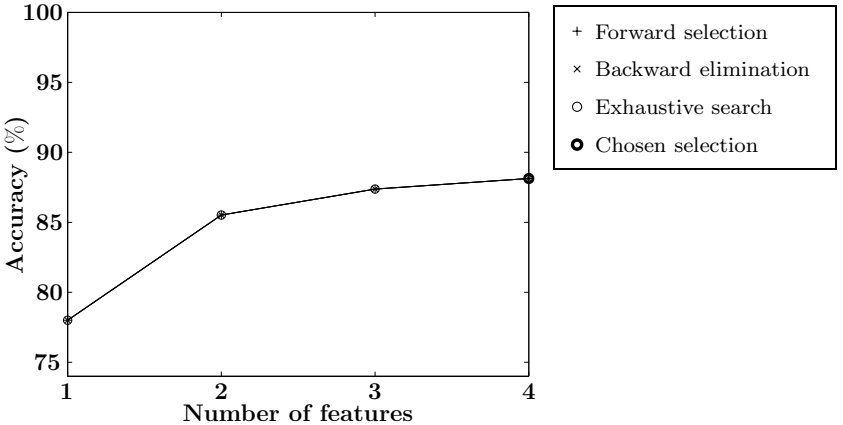

Table D.19: Spatio-temporal method, object-based, all echoes, three classes. Feature selection results. See page 187 for full description.

# of feats.	Classification accuracy (Feature combination)			
	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	90.8 % (3)	90.8 % (3)	(5)	90.8 % (3)
2	95.0 % (3 6)	95.0 % (3 6)	(3 5)	95.0 % (3 6)
3	96.5 % (3 4 6)	95.5 % (2 3 6)	(3 4 6)	96.5 % (3 4 6)
4	97.0 % (3 4 5 6)	97.4 % (2 3 5 6)	(3 4 5 6)	97.4 % (2 3 5 6)
5	***97.8*** % (2 3 4 5 6)	***97.8*** % (2 3 4 5 6)	(2 3 4 5 6)	***97.8 % (2 3 4 5 6)***
6	97.1 % (1 2 3 4 5 6)	97.1 % (1 2 3 4 5 6)	(1 2 3 4 5 6)	97.1 % (1 2 3 4 5 6)

Feat. #	Feature name
1	Optical flow, CC, dir. var.
2	Optical flow, CC, vel. var.
3	Optical flow, CC, dir. mean entropy
4	Optical flow, CC, vel. mean entropy
5	Optical flow, CC, NWP dir. mean diff.
6	Optical flow, CC, NWP vel. mean diff.

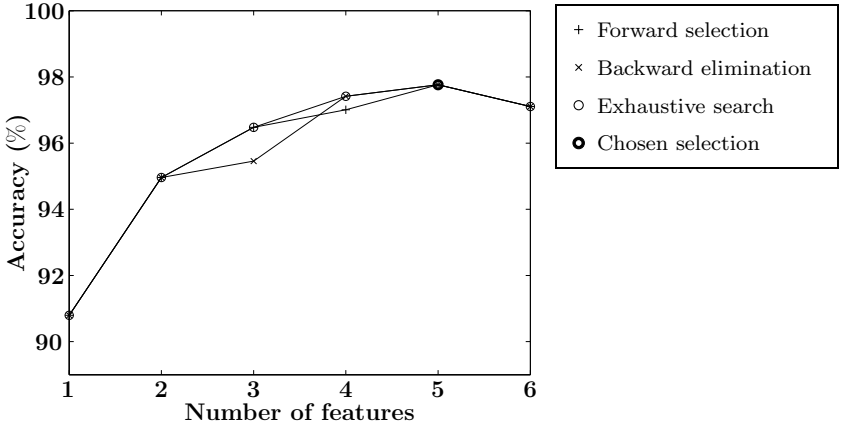

Table D.20: Spatio-temporal method, object-based, all echoes, two classes. Feature selection results. See page 187 for full description.

# of	Classification accuracy (Feature combination)			
feats.	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	90.3 % (1)	89.2 % (3)	(3)	90.3 % (1)
2	92.8 % (1 6)	93.7 % (3 6)	(3 4)	93.7 % (3 6)
3	95.8 % (1 5 6)	96.0 % (3 4 6)	(3 4 6)	96.0 % (3 4 6)
4	96.1 % (1 3 5 6)	96.6 % (3 4 5 6)	(2 3 4 6)	96.6 % (3 4 5 6)
5	96.7 % (1 3 4 5 6)	96.7 % (1 3 4 5 6)	(1 2 3 4 6)	96.7 % (1 3 4 5 6)
6	***97.0*** % (1 2 3 4 5 6)	***97.0*** % (1 2 3 4 5 6)	(1 2 3 4 5 6)	***97.0 % (1 2 3 4 5 6)***

Feat. #	**Feature name**
1	Optical flow, CC, dir. var.
2	Optical flow, CC, vel. var.
3	Optical flow, CC, dir. mean entropy
4	Optical flow, CC, vel. mean entropy
5	Optical flow, CC, NWP dir. mean diff.
6	Optical flow, CC, NWP vel. mean diff.

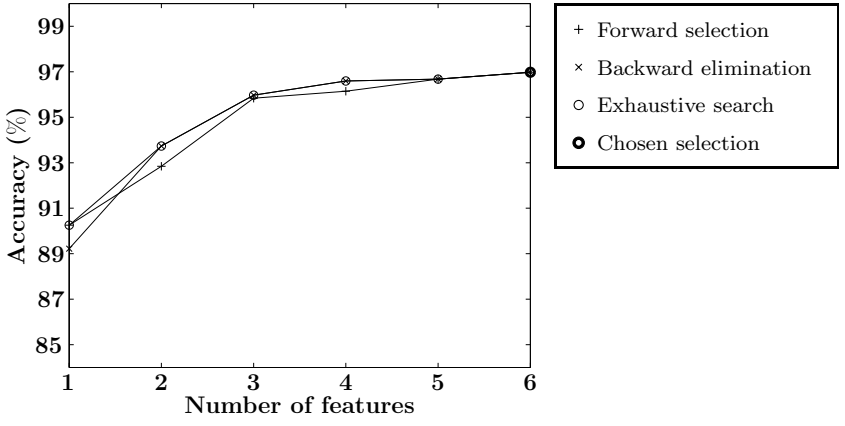

Table D.21: Spatio-temporal method, object-based, weak echoes, three classes. Feature selection results. See page 187 for full description.

# of feats.	Classification accuracy (Feature combination)			
	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	86.0 % (3)	86.0 % (3)	(5)	86.0 % (3)
2	91.9 % (3 6)	91.9 % (3 6)	(3 5)	91.9 % (3 6)
3	94.7 % (3 5 6)	94.7 % (3 5 6)	(3 4 5)	94.7 % (3 5 6)
4	***96.4*** % (3 4 5 6)	***96.4*** % (3 4 5 6)	(3 4 5 6)	***96.4 % (3 4 5 6)***
5	96.3 % (2 3 4 5 6)	96.3 % (2 3 4 5 6)	(2 3 4 5 6)	96.3 % (2 3 4 5 6)
6	96.2 % (1 2 3 4 5 6)	96.2 % (1 2 3 4 5 6)	(1 2 3 4 5 6)	96.2 % (1 2 3 4 5 6)

Feat. #	Feature name
1	Optical flow, CC, dir. var.
2	Optical flow, CC, vel. var.
3	Optical flow, CC, dir. mean entropy
4	Optical flow, CC, vel. mean entropy
5	Optical flow, CC, NWP dir. mean diff.
6	Optical flow, CC, NWP vel. mean diff.

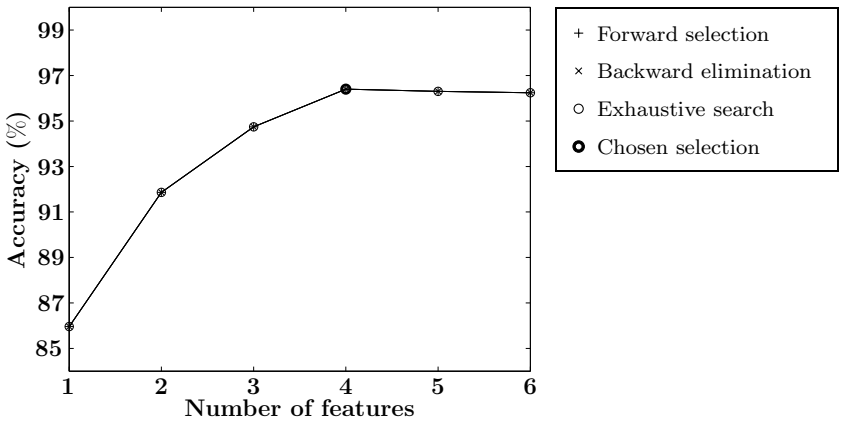

Table D.22: Spatio-temporal method, object-based, weak echoes, two classes. Feature selection results. See page 187 for full description.

# of	Classification accuracy (Feature combination)			
feats.	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	84.6 % (1)	84.1 % (3)	(3)	84.6 % (1)
2	88.7 % (1 5)	92.6 % (3 6)	(3 6)	92.6 % (3 6)
3	91.1 % (1 3 5)	94.4 % (3 5 6)	(3 4 6)	94.4 % (3 5 6)
4	95.0 % (1 3 5 6)	95.1 % (3 4 5 6)	(2 3 4 6)	95.1 % (3 4 5 6)
5	***95.7*** % (1 3 4 5 6)	***95.7*** % (1 3 4 5 6)	(1 2 3 4 6)	***95.7 % (1 3 4 5 6)***
6	95.7 % (1 2 3 4 5 6)	95.7 % (1 2 3 4 5 6)	(1 2 3 4 5 6)	95.7 % (1 2 3 4 5 6)

Feat. #	Feature name
1	Optical flow, CC, dir. var.
2	Optical flow, CC, vel. var.
3	Optical flow, CC, dir. mean entropy
4	Optical flow, CC, vel. mean entropy
5	Optical flow, CC, NWP dir. mean diff.
6	Optical flow, CC, NWP vel. mean diff.

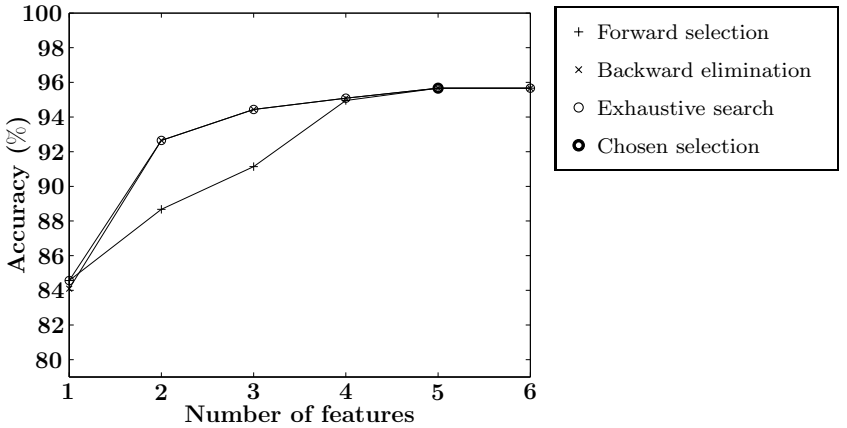

Table D.23: Spatio-temporal method, object-based, strong echoes, three classes. Feature selection results. See page 187 for full description.

# of	Classification accuracy (Feature combination)			
feats.	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	95.7 % (3)	95.7 % (3)	(3)	95.7 % (3)
2	97.7 % (3 6)	97.7 % (3 6)	(3 5)	97.7 % (3 6)
3	98.0 % (3 5 6)	98.0 % (3 5 6)	(3 5 6)	98.0 % (3 5 6)
4	98.2 % (3 4 5 6)	98.2 % (3 4 5 6)	(3 4 5 6)	98.2 % (3 4 5 6)
5	***98.5*** % (2 3 4 5 6)	***98.5*** % (2 3 4 5 6)	(2 3 4 5 6)	***98.5 % (2 3 4 5 6)***
6	97.9 % (1 2 3 4 5 6)	97.9 % (1 2 3 4 5 6)	(1 2 3 4 5 6)	97.9 % (1 2 3 4 5 6)

Feat. #	**Feature name**
1	Optical flow, CC, dir. var.
2	Optical flow, CC, vel. var.
3	Optical flow, CC, dir. mean entropy
4	Optical flow, CC, vel. mean entropy
5	Optical flow, CC, NWP dir. mean diff.
6	Optical flow, CC, NWP vel. mean diff.

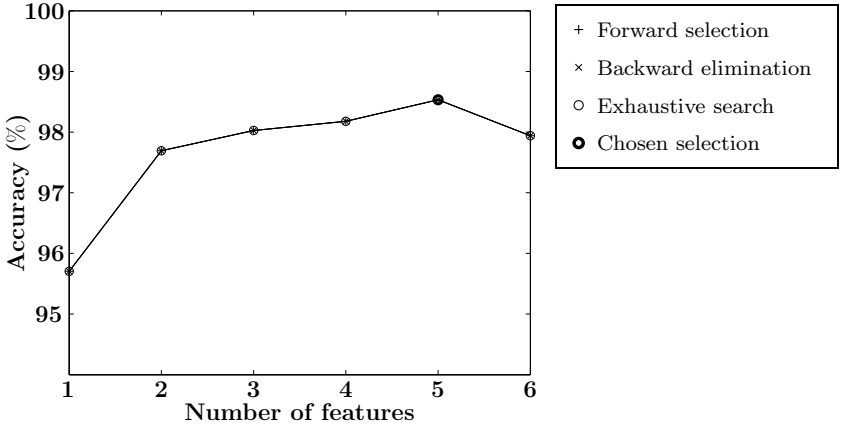

Table D.24: Spatio-temporal method, object-based, strong echoes, two classes. Feature selection results. See page 187 for full description.

# of feats.	Classification accuracy (Feature combination)			
	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	96.6 % (1)	95.2 % (3)	(1)	96.6 % (1)
2	97.9 % (1 6)	97.6 % (3 6)	(1 3)	97.9 % (1 6)
3	98.5 % (1 5 6)	**98.5** % (3 4 6)	(3 4 6)	**98.5 % (3 4 6)**
4	98.4 % (1 3 5 6)	98.7 % (3 4 5 6)	(1 3 4 6)	98.7 % (3 4 5 6)
5	***98.7*** % (1 3 4 5 6)	*98.7* % (1 3 4 5 6)	(1 2 3 4 6)	*98.7 % (1 3 4 5 6)*
6	98.6 % (1 2 3 4 5 6)	98.6 % (1 2 3 4 5 6)	(1 2 3 4 5 6)	98.6 % (1 2 3 4 5 6)

Feat. #	Feature name
1	Optical flow, CC, dir. var.
2	Optical flow, CC, vel. var.
3	Optical flow, CC, dir. mean entropy
4	Optical flow, CC, vel. mean entropy
5	Optical flow, CC, NWP dir. mean diff.
6	Optical flow, CC, NWP vel. mean diff.

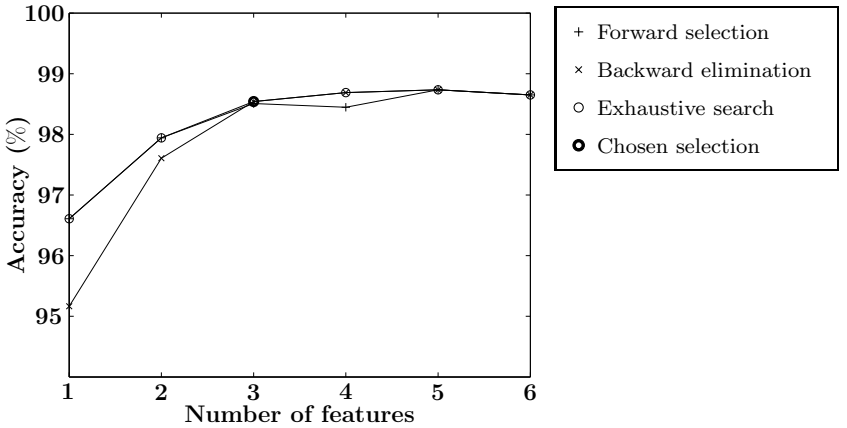

Table D.25: Propagation method, all echoes, three classes. Feature selection results. See page 187 for full description.

# of	Classification accuracy (Feature combination)			
feats.	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	93.5 % (3)	93.5 % (3)	(4)	93.5 % (3)
2	94.8 % (2 3)	94.8 % (2 3)	(2 4)	94.8 % (2 3)
3	95.4 % (1 2 3)	95.4 % (1 2 3)	(2 3 4)	95.4 % (1 2 3)
4	***96.4*** % (1 2 3 4)	***96.4*** % (1 2 3 4)	(1 2 3 4)	***96.4 % (1 2 3 4)***

Feat. #	**Feature name**
1	API, I_D
2	API, VRG1000
3	API, VRGmean1000
4	API, VRGmin10000

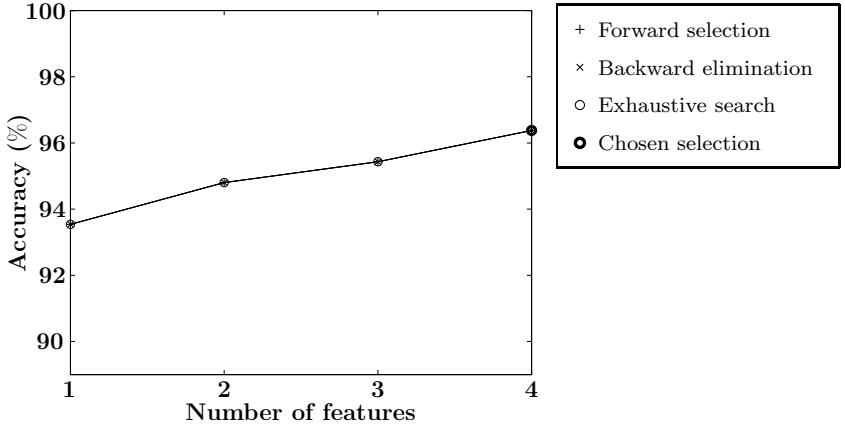

Table D.26: Propagation method, all echoes, two classes. Feature selection results. See page 187 for full description.

# of feats.	Classification accuracy (Feature combination)			
	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	93.5 % (3)	93.5 % (3)	(3)	93.5 % (3)
2	***94.4*** % (2 3)	***94.4*** % (2 3)	(3 4)	***94.4 % (2 3)***
3	93.9 % (1 2 3)	93.9 % (1 2 3)	(1 3 4)	93.9 % (1 2 3)
4	93.6 % (1 2 3 4)	93.6 % (1 2 3 4)	(1 2 3 4)	93.6 % (1 2 3 4)

Feat. #	Feature name
1	API, I_D
2	API, VRG1000
3	API, VRGmean1000
4	API, VRGmin10000

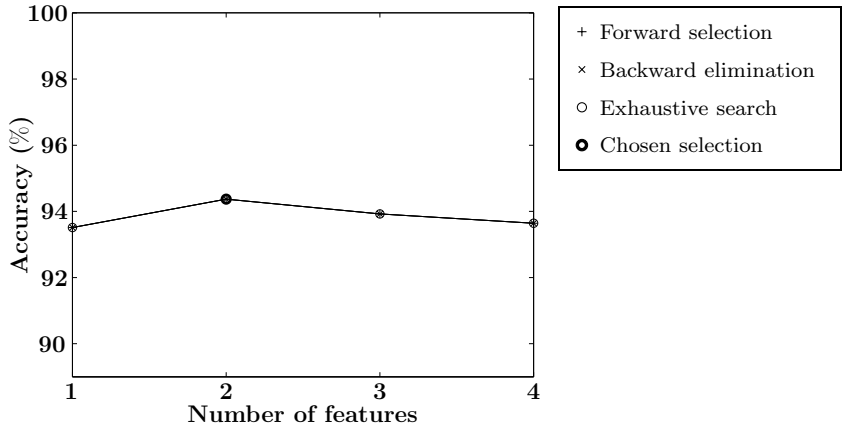

Table D.27: Propagation method, weak echoes, three classes. Feature selection results. See page 187 for full description.

# of	Classification accuracy (Feature combination)			
feats.	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	87.4 % (3)	87.4 % (3)	(4)	87.4 % (3)
2	91.3 % (2 3)	91.3 % (2 3)	(2 3)	91.3 % (2 3)
3	92.4 % (1 2 3)	92.4 % (1 2 3)	(2 3 4)	92.4 % (1 2 3)
4	***94.3*** % (1 2 3 4)	***94.3*** % (1 2 3 4)	(1 2 3 4)	***94.3 % (1 2 3 4)***

Feat. #	**Feature name**
1	API, I_D
2	API, VRG1000
3	API, VRGmean1000
4	API, VRGmin10000

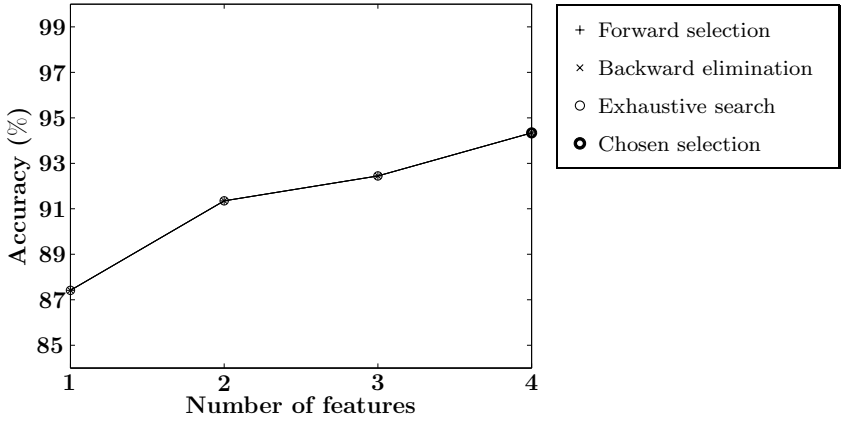

Table D.28: Propagation method, weak echoes, two classes. Feature selection results. See page 187 for full description.

# of	Classification accuracy (Feature combination)			
feats.	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	86.5 % (3)	86.5 % (3)	(3)	86.5 % (3)
2	90.3 % (2 3)	90.3 % (2 3)	(2 3)	90.3 % (2 3)
3	90.2 % (1 2 3)	90.2 % (1 2 3)	(1 3 4)	90.2 % (1 2 3)
4	***91.0*** % (1 2 3 4)	***91.0*** % (1 2 3 4)	(1 2 3 4)	***91.0 % (1 2 3 4)***

Feat. #	**Feature name**
1	API, I_D
2	API, VRG1000
3	API, VRGmean1000
4	API, VRGmin10000

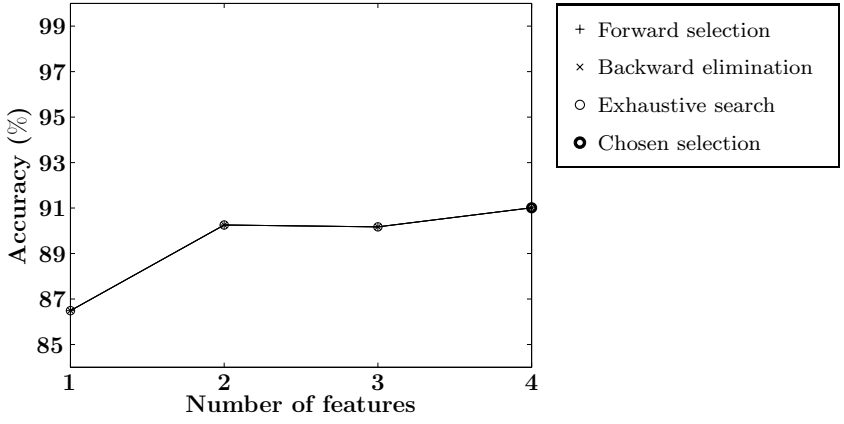

Table D.29: Propagation method, strong echoes, three classes. Feature selection results. See page 187 for full description.

# of	Classification accuracy (Feature combination)			
feats.	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	98.2 % (3)	**98.2** % (3)	(1)	98.2 % (3)
2	***98.6*** % (2 3)	98.1 % (1 3)	(1 3)	***98.6 % (2 3)***
3	97.8 % (2 3 4)	*98.2* % (1 3 4)	(1 2 4)	98.2 % (1 3 4)
4	97.9 % (1 2 3 4)	97.9 % (1 2 3 4)	(1 2 3 4)	97.9 % (1 2 3 4)

Feat. #	**Feature name**
1	API, I_D
2	API, VRG1000
3	API, VRGmean1000
4	API, VRGmin10000

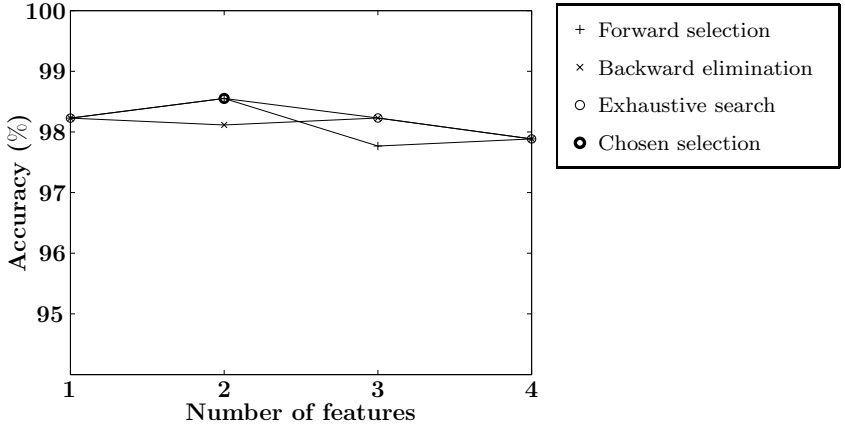

Table D.30: Propagation method, strong echoes, two classes. Feature selection results. See page 187 for full description.

# of	Classification accuracy (Feature combination)			
feats.	Forward selection	Backward elimination	Jeffries-Matusita	Exhaustive search
1	98.1 % (3)	97.7 % (2)	(2)	98.1 % (3)
2	***98.5*** % (2 3)	***98.1*** % (1 2)	(1 2)	***98.5 % (2 3)***
3	97.4 % (1 2 3)	97.9 % (1 2 4)	(1 3 4)	97.9 % (1 2 4)
4	97.4 % (1 2 3 4)	97.4 % (1 2 3 4)	(1 2 3 4)	97.4 % (1 2 3 4)

Feat. #	**Feature name**
1	API, I_D
2	API, VRG1000
3	API, VRGmean1000
4	API, VRGmin10000

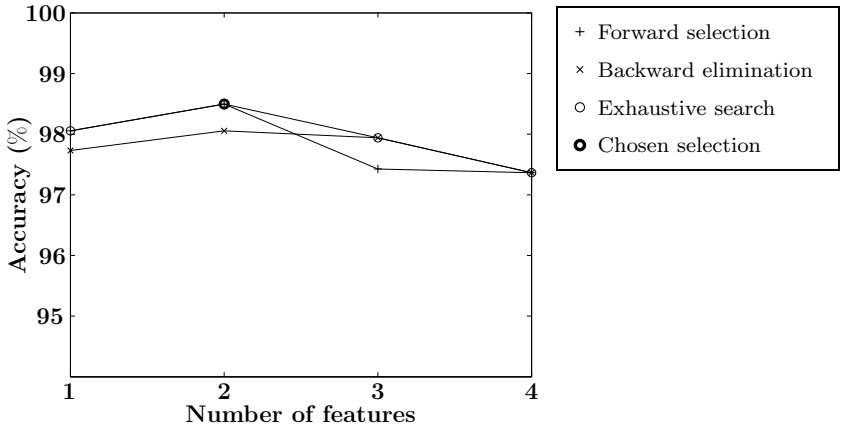

Bibliography

[1] H. Aanæs, J. R. Sveinsson, A. A. Nielsen, T. Bøvith, and J. A. Benediktsson. Model-based satellite image fusion. *IEEE Transactions on Geoscience and Remote Sensing*, 46(5):1336–1346, 2008.

[2] S. H. Abdul-Jauwad, P. Z. Khan, and T. O. Halawani. Prediction of radar coverage under anomalous propagation condition for a typical coastal site: A case study. *Radio Science*, 26(4):909–919, 1991.

[3] P. P. Alberoni, T. Andersson, P. Mezzasalma, D. B. Michelson, and S. Nanni. Use of the vertical reflectivity profile for identification of anomalous propagation. *Meteorological Applications*, 8(3):257–266, 2001.

[4] P. Anandan. A computational framework and an algorithm for the measurement of visual motion. *International Journal of Computer Vision*, 2(3):283–310, 1989.

[5] B. W. Atkinson and M. Zhu. Coastal effects on radar propagation in atmospheric ducting conditions. *Meteorological Applications*, 13(1):53–62, 2006.

[6] D. Atlas, editor. *Radar in Meteorology: Battan Memorial and 40th Anniversary Radar Meteorology Conference*. American Meteorological Society, Boston, MA, USA, 1990.

[7] A. E. Barrios. A terrain parabolic equation model for propagation in the troposphere. *IEEE Transactions on Antennas and Propagation*, 42(1):90–98, 1994.

[8] J. L. Barron, S. S. Beauchemin, and D. J. Fleet. On optical flow. In I. Plander, editor, *6th Int. Conf. on Artificial Intelligence and Information-Control Systems of Robots (AIICSR)*, pages 3–14, Bratislava, Slovakia, September 1994. World Scientific.

[9] D. Bebbington, S. Rae, J. Bech, B. Codina, and M. Picanyol. Modelling of weather radar echoes from anomalous propagation using a hybrid parabolic equation method and NWP model data. *Natural Hazards and Earth System Sciences*, 7:391–398, 2007.

[10] J. Bech, B. Codina, and J. Lorente. Forecasting weather radar propagation conditions. *Meteorology and Atmospheric Physics*, 2006.

[11] J. Bech, B. Codina, J. Lorente, and D. Bebbington. Monthly and daily variations of radar anomalous propagation conditions: How "normal" is normal propagation? In *Proceedings of the Second European Conference on Radar Meteorology (ERAD 2002)*, pages 35–39, Delft, Netherlands, November 2002.

[12] J. Bech, U. Gjertsen, and G. Haase. Modelling weather radar beam propagation and topographical blockage at northern high latitudes. *Quarterly Journal of the Royal Meteorological Society*, 133(626):1191–1204, 2007.

[13] U. C. Benz, P. Hofmann, G. Willhauck, I. Lingenfelder, and M. Heynen. Multi-resolution, object-oriented fuzzy analysis of remote sensing data for gis-ready information. *Journal of Photogrammetry and Remote Sensing*, 58(1-3):239–258, 2003.

[14] H.-H. Benzon and T. Bøvith. Simulation and prediction of weather radar clutter using a wave propagator on high-resolution NWP data. *IEEE Transactions on Antennas and Propagation*, Accepted for publication.

[15] H.-H. Benzon, A. S. Nielsen, and L. Olsen. An atmospheric wave optics propagator - theory and application. Danish Meteorological Institute Scientific Report 03–01, Danish Meteorological Institute, Copenhagen, Denmark, 2003.

[16] M. Berenguer, D. Sempere-Torres, C. Corral, and R. Sánchez-Diezma. A fuzzy logic technique for identifying nonprecipitating echoes in radar scans. *Journal of Atmospheric and Oceanic Technology*, 23(9):1157–1180, 2006.

[17] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer-Verlag, Berlin, Heidelberg, Germany, first edition, 2006.

[18] T. Bøvith, R. S. Gill, S. Overgaard, A. A. Nielsen, and L. K. Hansen. Detecting weather radar clutter using satellite-based nowcasting products. In *Proceedings of the Fourth European Conference on Radar Meteorology (ERAD 2006)*, pages 153–156, Barcelona, Spain, September 2006. Conference paper.

[19] T. Bøvith, A. A. Nielsen, L. K. Hansen, S. Overgaard, and R. S. Gill. Detecting weather radar clutter by information fusion with satellite images and numerical weather prediction model output. In *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium (IGARSS) 2006*, volume I, pages 511–514, Denver, Colorado, USA, August 2006. Conference paper.

[20] N. Bowler, C. Pierce, and A. Seed. Development of a precipitation nowcasting algorithm based upon optical flow techniques. *Journal of Hydrology*, 288(1-2):74–91, 2004.

[21] V. N. Bringi and V. Chandrasekar. *Polarimetric Doppler Weather Radar: Principles and Applications*. Cambridge University Press, New York, NY, USA, 2001.

[22] J. M. Carstensen, editor. *Digital Image Processing*. Technical University of Denmark (DTU), Lyngby, Denmark, fourth edition, 1999.

[23] M. Cheng and R. Brown. Delineation of precipitation areas by correlation of meteosat visible and infrared data with radar data. *Monthly Weather Review*, 123(9):2743–2757, 1995.

[24] Y.-H. Cho, G. Lee, K.-E. kim, and I. Zawadzki. Identification and removal of ground echoes and anomalous propagation using the characteristics of radar echoes. *Journal of Atmospheric and Oceanic Technology*, 23(9):1206–1222, 2006.

[25] C. G. Collier. *Applications of Weather Radar Systems: a guide to uses of radar data in meteorology and hydrology*. Praxis Publishing / John Wiley and Sons, New York, NY, USA, second edition, 1996.

[26] R. B. da Silveira and A. R. Holt. An automatic idedntification of clutter and anomalous propagation in polarization-diversity weather radar data using neural networks. *IEEE Transactions on Geoscience and Remote Sensing*, 39(8):1777–1788, 2001.

[27] G. D. Dockery. Modelling electromagnetic wave propagation in the troposphere using the parabolic equation. *IEEE Transactions on Antennas and Propagation*, 36(10):1464–1470, 1988.

[28] G. D. Dockery and J. R. Kuttler. An improved impedance boundary algorithm for fourier split-step solutions of the parabolic equation. *IEEE Transactions on Antennas and Propagation*, 44(12):1592–1599, 1996.

[29] R. J. Doviak and D. S. Zrnic. *Doppler Radar and Weather Observations*. Academic Press, San Diego, CA, USA, second edition, 1993.

[30] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification*. Wiley-Interscience, New York, NY, USA, second edition, 2000.

[31] E. E. Ebert and G. T. Weymouth. Incorporating satellite observations of "no rain" in an Australian daily rainfall analysis. *Journal of Applied Meteorology*, 38(1):44–56, 1999.

[32] R. S. Gill, T. Bøvith, and S. Overgaard. Sea clutter removal using radar elevation dependent second order texture parameters. In *International Symposium on Weather Radar and Hydrology (WRaH 2008)*, Grenoble, France, March 2008. Conference poster.

[33] R. S. Gill and S. Overgaard. Weather radar DMI pseudo CAPPI product versus the EEC BASE product. Danish Meteorological Institute Internal Report 05–10, Danish Meteorological Institute, Copenhagen, Denmark, 2005.

[34] R. S. Gill, S. Overgaard, and T. Bøvith. The Danish weather radar network. In *Fourth European Conference on Radar Meteorology (ERAD 2006)*, Barcelona, Spain, September 2006. Conference poster.

[35] M. Grecu and W. F. Krajewski. An efficient methodology for detection of anomalous propagation echoes in radar reflectivity data using neural networks. *Journal of Atmospheric and Oceanic Technology*, 17(2):121–129, 2000.

[36] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning; data mining, inference, and prediction*. Springer-Verlag, Berlin, Heidelberg, Germany, 2001.

[37] H. Hohti, J. Koistinen, P. Nurmi, E. Saltikoff, and K. Holmlund. Precipitation nowcasting using radar-derived atmospheric motion vectors. *Physics and Chemistry of the Earth, Part B: Hydrology Oceans and Atmosphere*, 25(10–12):1323–1327, 2000.

[38] K. Holmlund. The utilization of statistical properties of satellite-derived atmospheric motion vectors to derive quality indicators. *Weather and Forecasting*, 13(4):1093–1104, 1998.

[39] B. K. P. Horn and B. G. Schunck. Determining optical flow. *Artificial Intelligence*, 17(1–3):185–203, 1981.

[40] Intergovernmental Panel on Climate Change. *Fourth Assessment Report: Climate Change 2007: The AR4 Synthesis Report*. Geneva: IPCC, 2007.

[41] ITU. The radio refractive index: its formula and refractivity data. ITU-R Recommendations P.453-2, International Telecomminications Union, 1997.

[42] D. B. Johnson, P. Flament, and R. L. Bernstein. High-resolution satellite imagery for mesoscale meteorological studies. *Bulletin of the American Meteorological Society*, 75(1):5–33, 1994.

[43] D. E. Kerr, editor. *Propagation of Short Radio Waves*. Dover Publications, New York, NY, USA, 1951.

[44] C. Kessinger, S. Ellis, and J. V. Andel. The radar echo classifier: A fuzzy logic algorithm for the WSR-88D. In *3rd Conference on Artificial Intelligence Applications to the Environmental Science*, Long Beach, CA, USA, February 2003. AMS.

[45] C. Kessinger, S. Ellis, J. Van Andel, J. Yee, and J. Hubbert. Current and future plans for the ap clutter mitigation scheme. In *Preprints of the 20th International Conference on Interactive Information and Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology (20IIPS)*, Seattle, WA, USA, January 2004. American Meteorological Society.

[46] J. Kittler, M. Hatef, R. P. W. Duin, and J. Matas. On combining classifiers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(3):226–239, 1998.

[47] J. R. Kuttler and G. D. Dockery. Theoretical description of the parabolic approximation/fourier split-step method of representing electromagnetic propagation in the troposphere. *Radio Science*, 26(2):381–393, 1991.

[48] V. Lakshmanan, A. Fritz, T. Smith, K. Hondl, and G. Stumpf. An automated technique to quality control radar reflectivity data. *Journal of Applied Meteorology and Climatology*, 46(3):288–305, 2007.

[49] V. Lakshmanan and M. Valente. Real-time quality control of reflectivity data using satellite infrared channel and surface observations. In *Preprints of the 20th International Conference on Interactive Information and Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology (20IIPS)*, Seattle, WA, USA, January 2004. American Meteorological Society.

[50] R. Larsen. *Estimation of Visual Motion in Image Sequences*. Ph.D. Thesis, Department of Mathematical Modelling, Technical University of Denmark (DTU), Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby, Denmark, 1994.

[51] V. Levizzani, J. Schmetz, H. Lutz, J. Kerkmann, P. P. Alberoni, and M. Cervino. Precipitation estimations from geostationary orbit and prospects for Meteosat Second Generation. *Meteorological Applications*, 8(1):23–41, 2001.

[52] M. Levy. *Parabolic Equation Methods for Electromagnetic Wave Propagation*. Institution of Electrical Engineers, London, United Kingdom, 2000.

[53] T. Lindeberg. Scale-space theory: A basic tool for analysing structures at different scales. *Journal of Applied Statistics*, 21(2):224–270, 1994.

[54] S. Lovejoy and G. L. Austin. The delineation of rain areas from visible and IR satellite data from GATE and mid-latitudes. *Atmos.-Ocean*, 17:77–92, 1979.

[55] A. Magaldi, J. Bech, G. Delgado, and J. Lorente. Filtering weather radar AP echoes with MSG observations and NWP data. In *Proceedings of the Fourth European Conference on Radar Meteorology (ERAD 2006)*, pages 199–201, Barcelona, Spain, September 2006.

[56] P. Meischner, editor. *Weather Radar: Principles and Advanced Applications*. Springer-Verlag, Berlin, Heidelberg, Germany, 2004.

[57] D. B. Michelson and T. Andersson. Identification and suppression of anomalous propagation echoes in two-dimensional radar images. In *Proceedings of the 27th Conference on Radar Meteorology*, Vail, Colorado, USA, October 1995. American Meteorological Society.

[58] D. B. Michelson and D. Sunhede. Spurious weather radar echo identification and removal using multisource temperature information. *Meteorological Applications*, 11(1):1–14, 2004.

[59] NWCSAF/PGE04. Software user manual for PGE04 of the NWCSAF/MSG: Scientific part. Satellite Application Facility to NoWCasting and Very Short Range Forecasting 1.2, EUMETSAT, 2005.

[60] J. A. Pamment and B. J. Conway. Objective identification of echoes due to anomalous propagation in weather radar data. *Journal of Atmospheric and Oceanic Technology*, 15(1):98–113, 1998.

[61] G. S. Pankiewicz, C. J. Johnson, and D. L. Harrison. Improving radar observations of precipitation with a Meteosat neural network classifier. *Meteorology and Atmospheric Physics*, 76(1-2):9–22, 2001.

[62] M. Peura. Computer vision methods for anomaly removal. In *Proceedings of the Second European Conference on Radar Meteorology (ERAD 2002)*, pages 312–317, Delft, Netherlands, November 2002.

[63] C. Pohl and J. L. van Genderen. Multisensor image fusion in remote sensing: concepts, methods and applications. *International Journal of Remote Sensing*, 19(5):823–854, 1998.

[64] J. A. Richards and X. Jia. *Remote Sensing Digital Image Analysis: An Introduction*. Springer-Verlag, Berlin, Heidelberg, Germany, third edition, 1999.

[65] A. Ryzhkov, P. Zhang, R. Doviak, and C. Kessinger. Discrimination between weather and sea clutter using doppler and dual-polarization weather radars. In *Proceedings of the XXVII General Assembly of the International Union of Radio Science*, Maastricht, The Netherlands, 2002. The International Union of Radio Science (URSI).

[66] A. V. Ryzhkov and D. S. Zrnic. Polarimetric rainfall estimation in the presence of anomalous propagation. *Journal of Atmospheric and Oceanic Technology*, 15(6):1320–1330, 1998.

[67] B. H. Sass, N. W. Nielsen, J. U. Jørgensen, B. Amstrup, and M. Kmit. The operational DMI-HIRLAM system. Danish Meteorological Institute Technical Report 00–26, Danish Meteorological Institute, Copenhagen, Denmark, 2000.

[68] J. Schmetz, P. Pili, S. Tjemkes, D. Just, J. Kerkmann, S. Rota, and A. Ratier. An introduction to Meteosat second generation (MSG). *Bulletin of the American Meteorological Society*, 83(7):977–992, 2002.

[69] M. I. Skolnik. *Introduction to Radar Systems*. McGraw-Hill, New York, NY, USA, third edition, 2001.

[70] M. Steiner and J. A. Smith. Use of three-dimensional reflectivity structure for automated detection and removal of nonprecipitating echoes in radar data. *Journal of Atmospheric and Oceanic Technology*, 19(5):673–686, 2002.

[71] J. Sugier, J. P. d. Chatelet, P. Roquain, and A. Smith. Detection and removal of clutter and anaprop in radar data using a statistical scheme based on echo fluctuation. In *Proceedings of the Second European Conference on Radar Meteorology (ERAD 2002)*, pages 17–24, Delft, Netherlands, November 2002.

[72] H. R. A. Wessels and J. H. Beekhuis. Stepwise procedures for supression of anomalous ground clutter. *COST-75 International Seminar on Advanced Weather Radar Systems*, (EUR 16013 EN):270–277, 1994.

[73] C. Wilson. Review of current methods and tools for verification of numerical forecasts of precipitation. Working Document COST 717 WDF_02_200109_1, 2001.