Improved Methods for Snow-Cloud Separation Using Multi-temporal Meteosat-8 SEVIRI Imagery

Martijn de Ruyter de Wildt

Swiss Federal Institute of Technology, Institute of Geodesy and Photogrammetry, Zürich, Switzerland
and
MeteoSwiss, Zürich, Switzerland

Abstract

Meteosat-8 is the first geostationary satellite that possesses channels at all bandwidths that are of use for snow mapping. It therefore offers new possibilities for multi-temporal snow mapping, as well as for snow mapping in short time intervals, which is for example required for numerical weather prediction models. The spectral capabilities of Meteosat-8 allow an optimal spectral separation of clouds and surface snow cover, whereas the high temporal frequency introduces temporal information that might be used in the classification process. In this work we describe an algorithm that uses a new spectral feature and temporal for snow mapping and cloud detection.

1 Introduction

Snow cover influences several processes that occur at or near the earth's surface. It affects the exchange of energy and moisture between the surface and the atmosphere and is an important aspect of the hydrological cycle. Furthermore, snow cover extent is an indicator of climatic change and affects many human activities. Near real-time information about the surface snow cover is therefore important for studies and applications in many disciplines. This is particularly the case for Numerical Weather Prediction (NWP) models, which are initialised several times per day and require the latest information about the state of the atmosphere and the surface, including snow cover. A valuable tool for detecting snow cover is remote sensing, because it allows us to monitor large areas of the earth at regular time intervals.

A regularly encountered problem in remote sensing of snow is the sometimes similar spectral appearance of snow and clouds. In general, clouds have a similar reflectance as snow, and when they also have the same brightness temperature and phase (i.e. ice clouds), it can be difficult to distinguish them from snow with spectral information alone. Some authors have therefore used the spatial context of satellite pixels to detect clouds, but these methods are based on spatial inhomogeneity and are mainly suited for differentiating between cloud types and for detecting clouds over homogeneous surfaces. To differentiate between clouds and natural land surfaces, which often are quite inhomogeneous themselves, such methods are less suitable.

Another type of contextual information that can be used to classify satellite images is of temporal nature. Image classification that uses temporal information is generally referred to as (digital) change detection. In the literature, this notion generally refers to changes at the earth's surface, but there seems no reason why change detection methods might not be used for detecting clouds. Mostly, change detection involves two images of the same scene acquired at different dates, but in a number of applications temporal series of images are used. To the

latter category, which clearly offers more opportunities for detecting change as it uses more images, belong temporal trajectory analysis and temporal compositing. In temporal trajectory analysis, the temporal trajectory of a pixel is compared with a predefined trajectory, whereas in temporal compositing a composite is made from a series of individual images by retaining those pixels that satisfy a certain criterion. Trajectory analysis of high-frequency images is used for background estimation in video surveillance and photogrammetry. There, the task is to detect and/or remove objects that temporarily obscure the background, which is in fact comparable to detecting moving clouds over a static surface. However, background estimation requires the obscuring object to be present in only a few of a series of images, whereas clouds often cover large areas and single pixels can be cloud-covered during large parts of an observation period.

In remote sensing, change detection is generally used to study processes that occur at rather long time scales of months to years. This is the case for change detection over land, as well as over the oceans (e.g. sea ice). As a consequence, in most of these studies polar orbiting sensors have been used, as these have repeat times of hours to weeks and offer a wide range of spatial and spectral resolutions. In contrast to land surfaces and oceans, clouds often display a dynamic behaviour at time scales of minutes to hours, and only geostationary satellites have a frequency that is high enough to monitor this behaviour. However, unlike many polar orbiting sensors, geostationary platforms did until recently not possess all spectral channels that are required for optimal spectral separation of snow and clouds.

In 2002, the European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT), launched the first of a new series of geostationary satellites, called Meteosat Second Generation (MSG). This new satellite, Meteosat-8 (MSG-1) bridges the gap between polar orbiting sensors with good spectral resolution and geostationary sensors with high temporal frequency. It thus offers an unprecedented dataset of spectral and temporal information, which can be used to detect clouds over cold regions and to map surface snow. Here we describe an algorithm that uses temporal trajectory analysis in conjunction with pixel-based spectral classification to detect clouds and to map surface snow cover. This algorithm is intended for delivering real-time snow cover data to the operational mesoscale NWP model of MeteoSwiss, the Alpine Model (aLMo).

2 Data

Meteosat-8 is currently situated at 3.4° western longitude at an altitude of 36.000 km. It carries the Spinning Enhanced Visible and Infrared Imager (SEVIRI), which has improved spectral, spatial and temporal resolution with respect to its predecessors on board of the previous Meteosat satelites. It continuously monitors the entire earth disk with a frequency of 15 minutes. SEVIRI has twelve spectral channels, most of which measure radiation from the surface. Only the water vapour absorption channels, at 6.25 and 7.35 μm, contain no information about the surface at all. The ozone absorption channel (9.66 μm) measures radiation from the troposphere and the surface and is also sensitive to ozone concentration in the lower stratosphere. The CO_2 absorption channel (13.4 μm) mainly measures radiation from the troposphere with only a small contribution from the surface. Some CO_2 is also detected by the 3.9 μm channel, which slightly overlaps with one of the CO_2 absorption bands. Channel 12 is a high resolution visible (hrv) channel, which has a spatial resolution of 1 km at the sub-satellite point. All of the other channels have a spatial resolution of 3 km at the sub-satellite point. The region of interest in this study is the model domain of aLMo, which corresponds to western and central Europe. The spatial resolution of Meteosat-8 over this area is 1.5 to 2 km for the hrv channel and 5 to 6 km for the other channels.

For testing and validation we selected a three-day period in March 2004. During the selected period, March 8^{th} till March 10^{th}, all mountainous regions of Europe were covered with snow, as well as large parts of Central, Eastern and Northern Europe. The weather was variable, with low pressure activity over Central Europe and the Mediterranean. Clouds, some of them containing ice particles, bare land and snow covered land were therefore all well represented over Europe, making this period suitable for testing a snow mapping algorithm.

3 Pre-processing

The Meteosat-8 data that we use are provided in Level 1.5 Native Format. These consist of raw satellite counts, which need to be calibrated and converted into reflectances (r) and brightness temperatures (BT). We also apply a sea mask to the data, which is based on the SRTM30 Digital Elevation Model (DEM) of NASA and the United States Geological Survey. This global DEM has a horizontal resolution of 1 km, and we resampled it to the Meteosat-8 grid with bicubic interpolation. Because of the large viewing angles of Meteosat-8 over Europe, it is furthermore necessary to ortho-rectify the data.

3.1 Correction for atmospheric effects and anisotropy

Further pre-processing of the measured reflectances involves correcting for the influence of the atmospheric and for anisotropy of reflection at the surface. The atmospheric influence depends on the state of the atmosphere and on the solar and satellite viewing angles, and can be described with Radiative Transfer Models (RTMs). However, state-of-the-art RTMs are not very reliable for solar zenith angles over 70°, whereas snow cover is mainly present during the winter season, when the sun remains rather low above the horizon. Also, RTMs require the atmospheric aerosol content and atmospheric profiles of water vapour, ozone and CO_2, which are generally not available. Water vapour can be provided by NWP models, but always with some degree of uncertainty. We therefore only apply one correction for all angular effects, including those caused by atmospheric radiative transfer and by anisotropic reflection at the surface. For this correction, a semi-empirical model that describes bi-directional surface reflectance is used. The model has five coefficients that can be related to the Normalised Differential Vegetation Index (NDVI). From Meteosat-8 SEVIRI data, the NDVI can be computed for the low resolution channels. For the hrv channel, we use the downscaled low-resolution NDVI values.

The temporal behaviour of the reflectances at two cloud-free locations, illustrating the angular effects, is shown in Fig. 1. The first location was snow-free, whereas the second was partially snow-covered. At both locations the reflectances display a peak in the early afternoon, which corresponds to the hot spot in the BRDF. Also, the angular effects increase with wavelength at both locations. This is most clearly the case over the snow-free location, whereas it is less obvious over snow-cover. Over snow the 1.6 μm reflectance is much lower than the other reflectances, so that the absolute effect of anisotropy is not very large. The relative effect, however, is largest for the 1.6 μm reflectance, and this increase with wavelength corresponds with the observations of other authors.

For each of the solar channels (channel 1, 2, 3 and 12), the coefficients c_i were derived by tuning the bi-directional reflectance model to the cloud-free pixels in the available satellite images. Because the true value of r (the actual hemispherical reflectance of the surface) is unknown, r is set equal to 1 in the tuning procedure. The resulting BRDF's can then be used to bring all observed reflectances to a reference viewing and illumination geometry. Over vegetated surfaces without snow cover the model performs well, but for pixels that contain

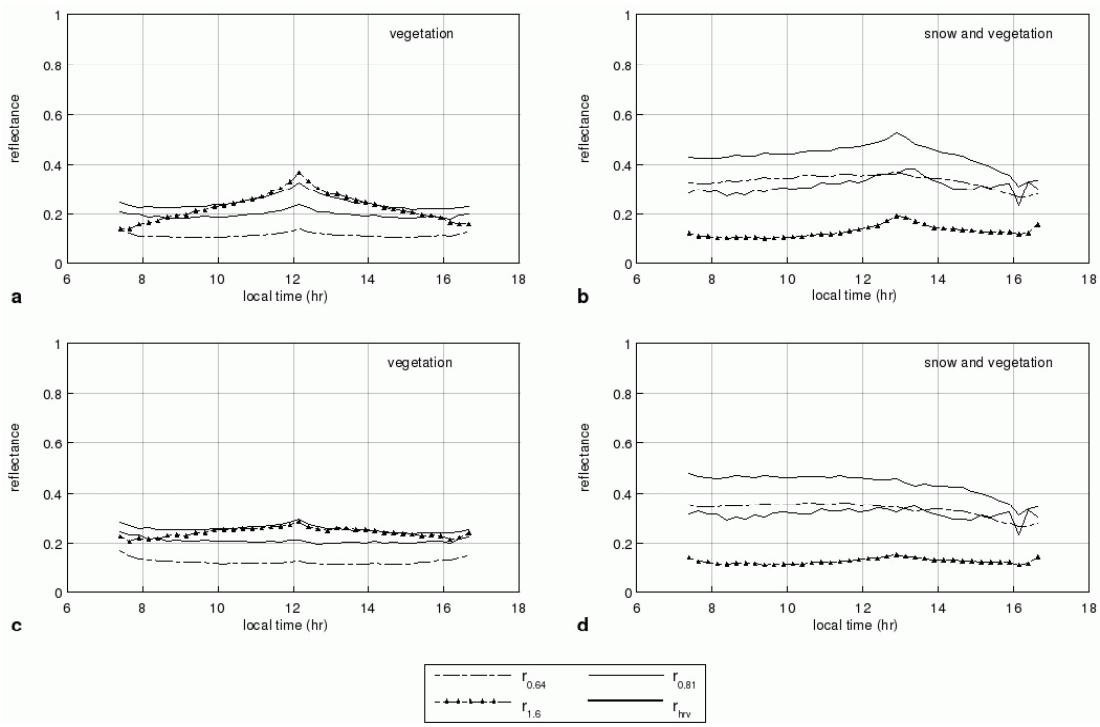

Figure 1: Reflectances as a function of local time at two cloud-free locations on March 10^{th}, 2004. Only data with corresponding solar zenith angles below 75° are shown. Temporal profiles are shown for a vegetated location (a) and for a vegetated location that was partially covered with snow (b). Temporal profiles that have been corrected for angular effects are shown in c and d, respectively.

snow cover, mostly mixed with vegetation in our images, it could not be adequately tuned. We therefore determined the tuning coefficients only for pixels without snow, which display a smooth dependence of the coefficients on NDVI. These dependencies can be described by simple polynomial or exponential functions and we let these functions approach zero for very low NDVI (corresponding to pixels with snow). This approach means that over mixed pixels, we only apply BRDF's for vegetation and not for snow. At the locations for which the time series are shown in Fig. 1a and b, the BRDF's that we use remove a large part of the temporal variation (Fig. 1c and d). The most variation remains at the partly snow covered location (Fig. 1d), which may be caused by anisotropic reflectance of the snow. It could also be due to melting of snow and an increasingly lower snow fraction during the afternoon, as suggested by the temporal behaviour of the reflectances. If the snow fraction would remain constant throughout the day and the temporal behaviour was only caused by angular effects, all reflectance channels would follow a similar pattern. Here, however, the visual reflectances are constant in the morning and decrease during the afternoon, whereas the near infrared reflectance is slightly increased during the afternoon. Both effects are well explained by a lower snow fraction.

4 Classification

4.1 Temporal features

Examples of the temporal behaviour of clouds are shown in Fig. 2. The first location (Fig. 2a and b) was covered with ice clouds in the morning, as indicated by the low 1.6 μm reflectance,

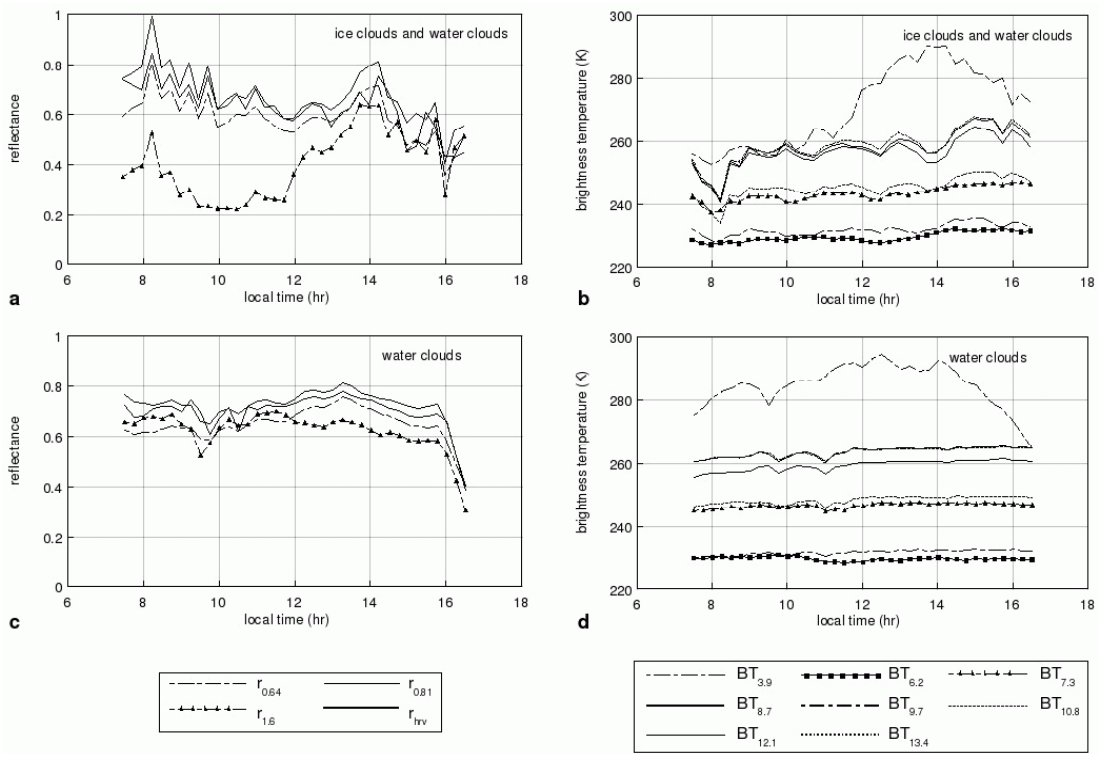

Figure 2: Reflectances and brightness temperatures as a function of local time at two cloudy locations on March 10^{th}, 2004. Only data with corresponding solar zenith angles below 75° are shown. Temporal profiles are shown for a location that was covered with ice clouds (a, b) and for a location that was covered with ice and water clouds (c, d).

and in the afternoon with water clouds. There is considerable temporal variability in most of the spectral channels. The second location (Fig. 2c and d) was also covered with clouds, but here the cloud cover consisted entirely of water clouds that were fairly homogeneous in space and time. Consequently, the overall temporal variability is lower. At both locations, the infrared absorption channels, in which information from the surface and the lower atmosphere has been (partly) filtered out, display less variation than the other infrared channels. As a measure of temporal variability we use the standard deviation in time. We found that it is also useful to take same temporal information from the eight surrounding pixels into account by averaging the standard deviation in time over each block of nine pixels. Although this classifier considers the eight surrounding pixels, it does not quantify spatial variability, and can therefore be regarded as a quasi three-dimensional classifier. To illustrate the usefulness of the temporal standard deviation, scatter plots of this classifier against the near-infrared reflectance are shown in Fig. 3. Pixels that represent water clouds, which have a high near-infrared reflectance, appear on the right hand sides of these plots. Ice clouds and snow have low infrared reflectances (on the left), whereas mixed clouds and snow-free surfaces display intermediate values. The cluster in the bottom left corners could be identified as surface snow, and the cluster next to it as snow-free surface. These pixels display a low temporal variability. Mixed clouds and ice clouds generally display larger temporal variabilities, whereas water clouds, represented by the cluster on the right-hand side of both plots, display low variabilities. This is especially the case for the brightness temperature (Fig. 3b).

We found the optimal number of time steps for separating clouds from surface snow by calculating a divergence parameter, which indicates the ability of a feature to separate two classes. When the temporal standard deviation at one pixel is used, for most channels the

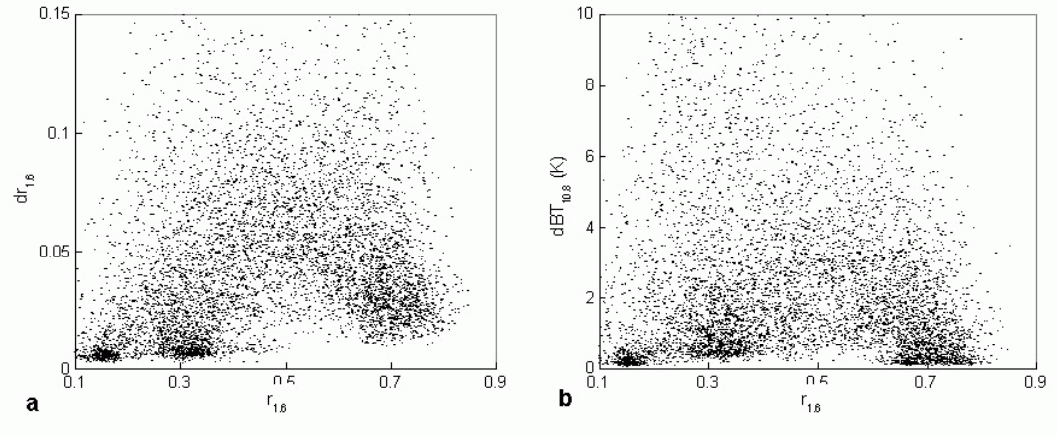

Figure 3: Scatter plots of the quasi 3-D temporal standard deviation against the near-infrared reflectance on March 10^{th}, 2004, 12:12. Displayed are scatter plots for the 1.6 μm reflectance (a) and for the 10.8 μm brightness temperature (b).

divergence is largest when 7 successive images are used. Shorter time series obviously include too little temporal information. On the other hand, the temporal variability of pixels that are cloud-free in the current image and clouded several time steps earlier or later (or vice versa), will increase for longer time series. Such pixels will be classified as clouds when the temporal variability at the beginning and/or at the end of the time series is large. When the temporal standard deviation of the eight surrounding pixels is also taken into account, the divergence improves and only 5 successive image are needed for the best results.

Table 1: The twelve channels of the SEVIRI instrument on board of Meteosat-8.

Channel	Spectral band μm			Description
	centre	min.	max.	
1	0.635	0.56	0.71	visual
2	0.81	0.74	0.88	visual
3	1.64	1.50	1.78	near infrared
4	3.90	3.48	4.36	solar + terrestrial infrared
5	6.25	5.35	7.15	infrared (water vapour absorption)
6	7.35	6.85	7.85	infrared (water vapour absorption)
7	8.70	8.30	9.10	infrared
8	9.66	9.38	9.94	infrared (ozone absorption)
9	10.80	9.80	11.80	infrared
10	12.00	11.00	13.00	infrared
11	13.40	12.40	14.40	infrared (CO_2 absorption)
12	0.75	0.60	0.90	high resolution visual broadband

We found higher divergences for the reflectance channels than for the infra red channels, which is caused by the often static behaviour of water clouds (Fig. 3b). When we omit all cloudy pixels with a near infrared reflectances above 0.5, we find comparable divergences for all channels (Table 1). The highest divergence is found for the reflectance in the hrv channel, which detects the most detailed information. The other three solar channels also display high divergences whereas the infrared channels display somewhat lower divergences. As expected, the lowest values are found for the water vapour absorption channels at 6.2 and 7.3 μm. These channels measure mid-atmospheric water content, and changes in this

quantity can occur independently of cloudiness at lower levels. These channels are therefore not suitable for detecting clouds. Significantly larger divergences are found for the other absorption channels at 9.7 and 13.4 μm, which detect significant amounts of information from low atmospheric levels. Thus, we use all channels except the water vapour absorption channels for temporal detection of clouds.

4.2 Spectral features

SEVIRI has channels in more spectral bands that can be used for cloud detection and snow mapping than any other currently available sensor, apart from MODIS on board of NASA's Terra and Aqua satellites. An existing cloud mask for the MODIS snow product uses $r_{0.64}$, $r_{1.6}$, $BT_{3.9} - BT_{10.8}$, $BT_{13.4}$ and the Normalised Differential Snow Index (NDSI), which equals $(r_{0.64} - r_{1.6})/(r_{0.64} + r_{1.6})$. In addition to these features we also use the 3.9 - 13.4 μm thermal difference. The latter feature often reveals water clouds, but it does not detect optically thick ice clouds, as illustrated by Fig. 4a. Water clouds have a large $BT_{3.9} - BT_{10.8}$ and appear bright, whereas unclouded regions, which have a small or even negative $BT_{3.9} - BT_{10.8}$, appear dark. Clouds with a high ice content, some of which are indicated in Fig. 4a, appear as dark as or somewhat brighter than snow-covered areas. A similar picture arises when we compute $BT_{3.9} - BT_{13.4}$ (Fig. 4b), but now many ice clouds tend to be darker than snow.

This difference can be attributed to CO_2 absorption, which occurs in the 3.9 and 13.4 μm channels. It reduces the amount of observed radiation, leading to lower observed brightness temperatures. The effect is far larger at 13.4 μm than at 3.9 μm and consequently, $BT_{3.9} - BT_{13.4}$ is always strongly positive. No CO_2 absorption takes place at 10.8 μm so that $BT_{10.8}$ is much higher than $BT_{13.4}$ and $BT_{3.9} - BT_{10.8}$ always smaller than $BT_{3.9} - BT_{13.4}$. Furthermore, more CO_2 absorption takes place when the atmospheric path length is longer, so that in general it has a stronger cooling effect over the surface than over clouds. The difference between $BT_{3.9} - BT_{10.8}$ and $BT_{3.9} - BT_{13.4}$ is therefore smallest for high altitude pixels. A scatter plot of the two brightness temperature differences (Fig. 5) clearly shows two bands of pixels, one corresponding to ice clouds and one corresponding to surface pixels. The use of both features should thus improve the separation of ice clouds and snow, which can be visualised by computing the ration between them. A plot of this ratio (Fig. 4c) clearly reveals many ice clouds that are not detectable with each separate brightness temperature difference (Fig. 4a and 4b).

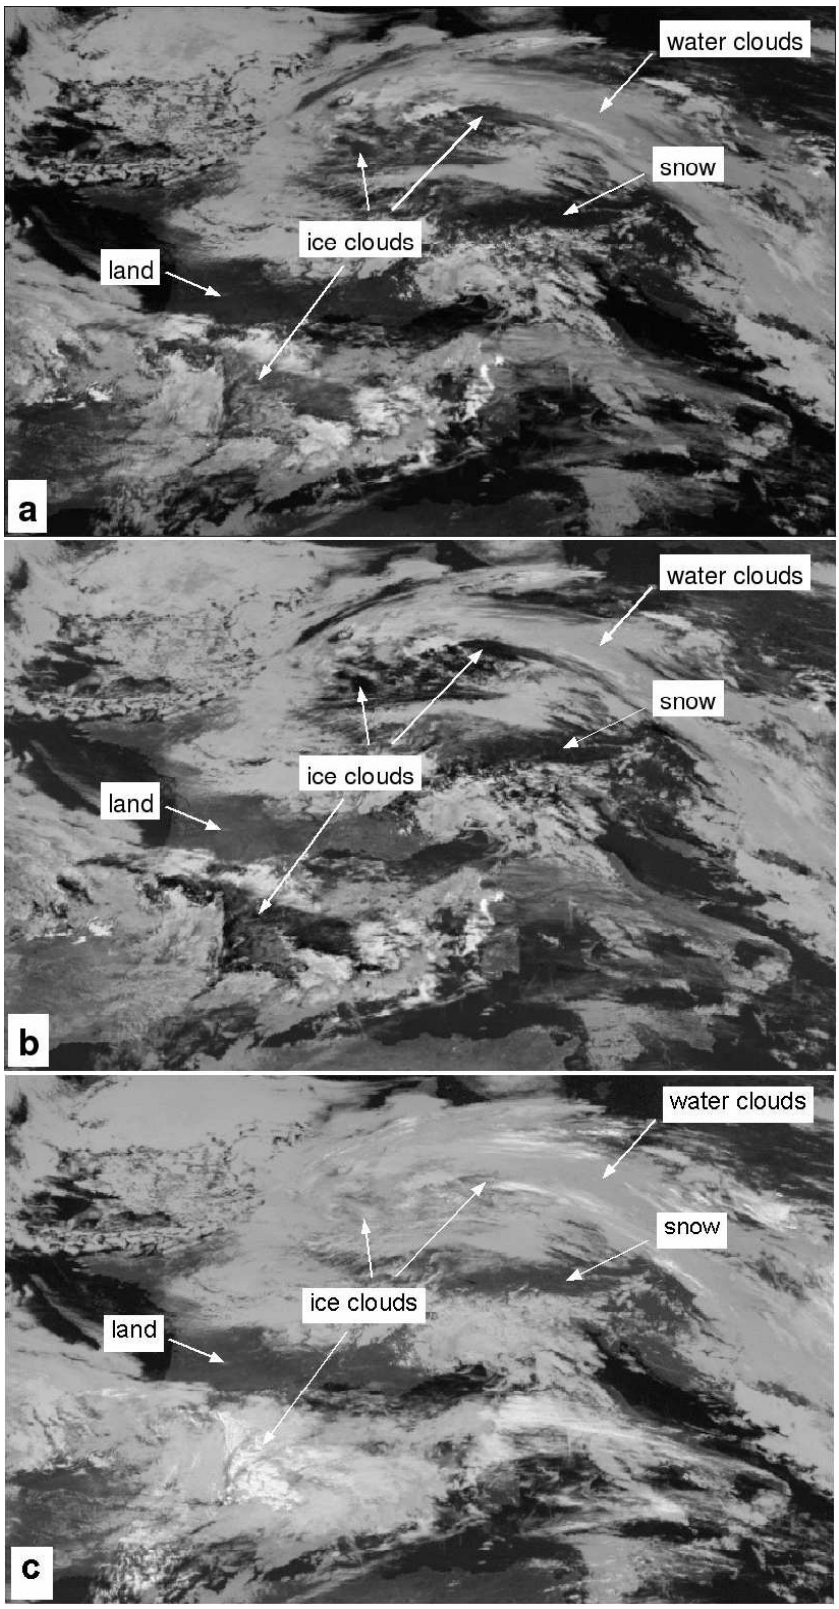

Figure 4: Normalised Meteosat-8 brightness temperature differences over the study area on March 10^{th}, 2004, 12:12 UTC. Shown are $BT_{3.9} - BT_{10.8}$ (a), $BT_{3.9} - BT_{13.4}$ (b) and $(BT_{3.9} - BT_{10.8} - 5)/(BT_{3.9} - BT_{13.4})$ (c). This scene could be visually classified by making use of the multi-spectral and multi-temporal information that is available.

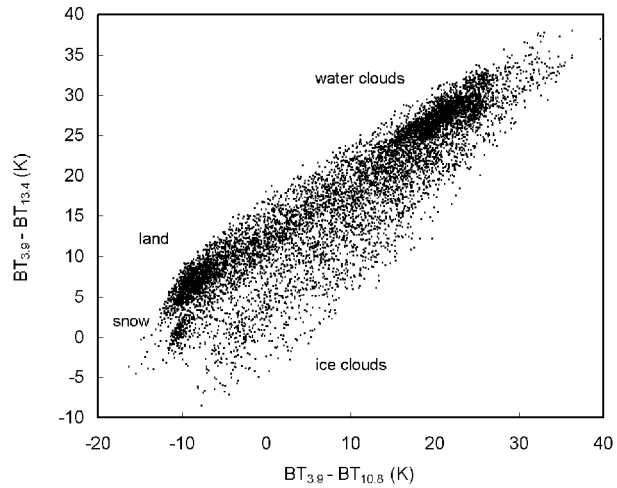

Figure 5: Scatter plot of $BT_{3.9} - BT_{13.4}$ against $BT_{3.9} - BT_{10.8}$ for the same scene as shown in Fig. 3.

4.3 Classification method

In remote sensing of snow cover, often threshold based classification trees are used. The spectral properties of clouds and snow are well known, which makes it straight-forward to choose threshold tests and to set values for the thresholds. This classification method generally gives good results and is easy to implement. Here we choose another standard classification method, namely maximum likelihood classification. With this method each pixel is assigned to the class for which the conditional probability of the pixel is highest. The advantage of this method is that it can adequately classify pixel distributions like the one shown in Fig. 5. Also, maximum likelihood classification gives probabilities in stead of rigid values (e.g. snow or cloud), which can be used for assigning quality flags to the pixels.

When we assume that the features are normally distributed, the conditional Probability Density Functions (PDFs) are given by the multi-variate normal distribution. For each images class, this distribution is described be the mean feature values and the feature covariance matrix. We chose four classes to which pixels can be assigned: snow-free land, snow, ice clouds and water clouds. Although we are not interested here in distinguishing between different cloud types, we do make the division between ice clouds and water clouds in order to improve the separation of clear and cloudy pixels. Clouds containing ice particles differ in appearance from water clouds in several ways: $r_{1.6}$, $BT_{3.9} - BT_{10.8}$ and $BT_{3.9} - BT_{13.4}$ are lower (see Fig. 3 and 5) when ice particles are present. There is also a difference in mean temporal variability between ice clouds and water clouds (Fig. 3). A simple threshold of for $r_{1.6}$ is used for differentiating between training areas for ice clouds and water clouds.

In order to determine for all classes the means of the features and the covariance matrices, we first classified all images with a simple threshold-based classification using only the spectral features. The values of the thresholds were chosen such that the best classification results were obtained, as could be subjectively judge by visual inspection. For the threshold-based classification we use a simple scheme that includes the ratio between $BT_{3.9} - BT_{10.8}$ and $BT_{3.9} - BT_{13.4}$, which we found very suitable for detecting clouds (Fig. 4c). Pixels that are not classified as cloudy in this way, are checked for the presence of snow by a second suite of tests.

The threshold-based classification missed some clouds that were misinterpreted as snow, but

the overall quality was judged acceptable. The classification results of all images were then used as training areas for computing the means and covariances for all features and classes. Then, the images were classified again, now with the maximum likelihood method. In the new results the misclassified clouds were no longer present and these results were therefor used for a final determination of the means and covariances.

5 Results

For analysing the performance of the algorithm, we focus on March 10^{th}, 2004, 12:12 UTC. An RGB image for this image which we found very useful for visual inspection is shown in Fig. 6a. Most image classes are clearly discernible from each other in this RGB combination. Only the colour ranges of snow (red) and ice clouds (red/pink) slightly overlap. For visual discrimination between these two classes one can view animated time-series of this RGB combination, which visualise both the spectral and the temporal component. On the project web-site (`www.photogrammetry.ethz.ch/research/snow/index.html`) examples of such animated time series are available. Before we discuss the classification result of this image, it is worthwhile to have a look at the conditional probabilities. Three of these conditional probabilities can be combined into one RGB image, and this is shown for snow-free surfaces, snow and water clouds in Fig. 6b. The three classes very clearly emerge in different colour groups, and even snow and ice clouds are now clearly distinguishable from each other. Pixels that have comparable conditional probabilities for all three classes appear in grey tones. In this RGB combination, this is the case for water surface and for some ice clouds, which both appear in white.

The result for the full maximum likelihood classification, i.e. including both spectral and temporal features, is shown in Fig. 7a. There are no false positives, i.e. no snow is detected where it is not present. However, a few false negatives occur: these are transparent and/or sub-pixel clouds over snow that are missed. When a binary snow map is requested the latter aspect is an advantage, as more snow is detected, but when fractional snow cover is to be derived pixels should be completely cloud-free. The influence of each type of information, spectral and temporal, upon the classification result can be investigated by using only one of these types of information for the classification. In Fig. 7b the result is shown for the case when only spectral information is used. Now, more snow is detected in some places and less snow in other. The reason for this is that some clouds display very low temporal variability, which lowers the conditional probability for clouds. These clouds may therefor not be detected when temporal information is used, whereas they are detected when no temporal information is used. For clouds with high temporal variability the opposite may be true. When only temporal information is used to mask clouds with the maximum likelihood method, we found that many clouds are missed. The cause of this poor performance is that whereas all temporal features display low values over cloud-free pixels, the opposite is not necessarily true over cloudy pixels. Slightly better results are obtained when we obtain a temporal cloud mask in each single channel and then stack all single channel cloud masks. A more substantial improvement is obtained when in each single channel the decision rule is changed in favour of the conditional probability for clouds. For example, when pixels are masked as cloud when the conditional probability of the temporal classifier is twice as high for clouds as for cloud-free surfaces, we obtain the temporal cloud mask that is shown in Fig. 7c. This cloud mask can be used to check all pixels that were classified as snow in the spectral classification (Fig. 7b) for high temporal variability. As a result (Fig. 7d) many mixed pixels near cloud edged that were previously classified as snow, are now classified as cloud.

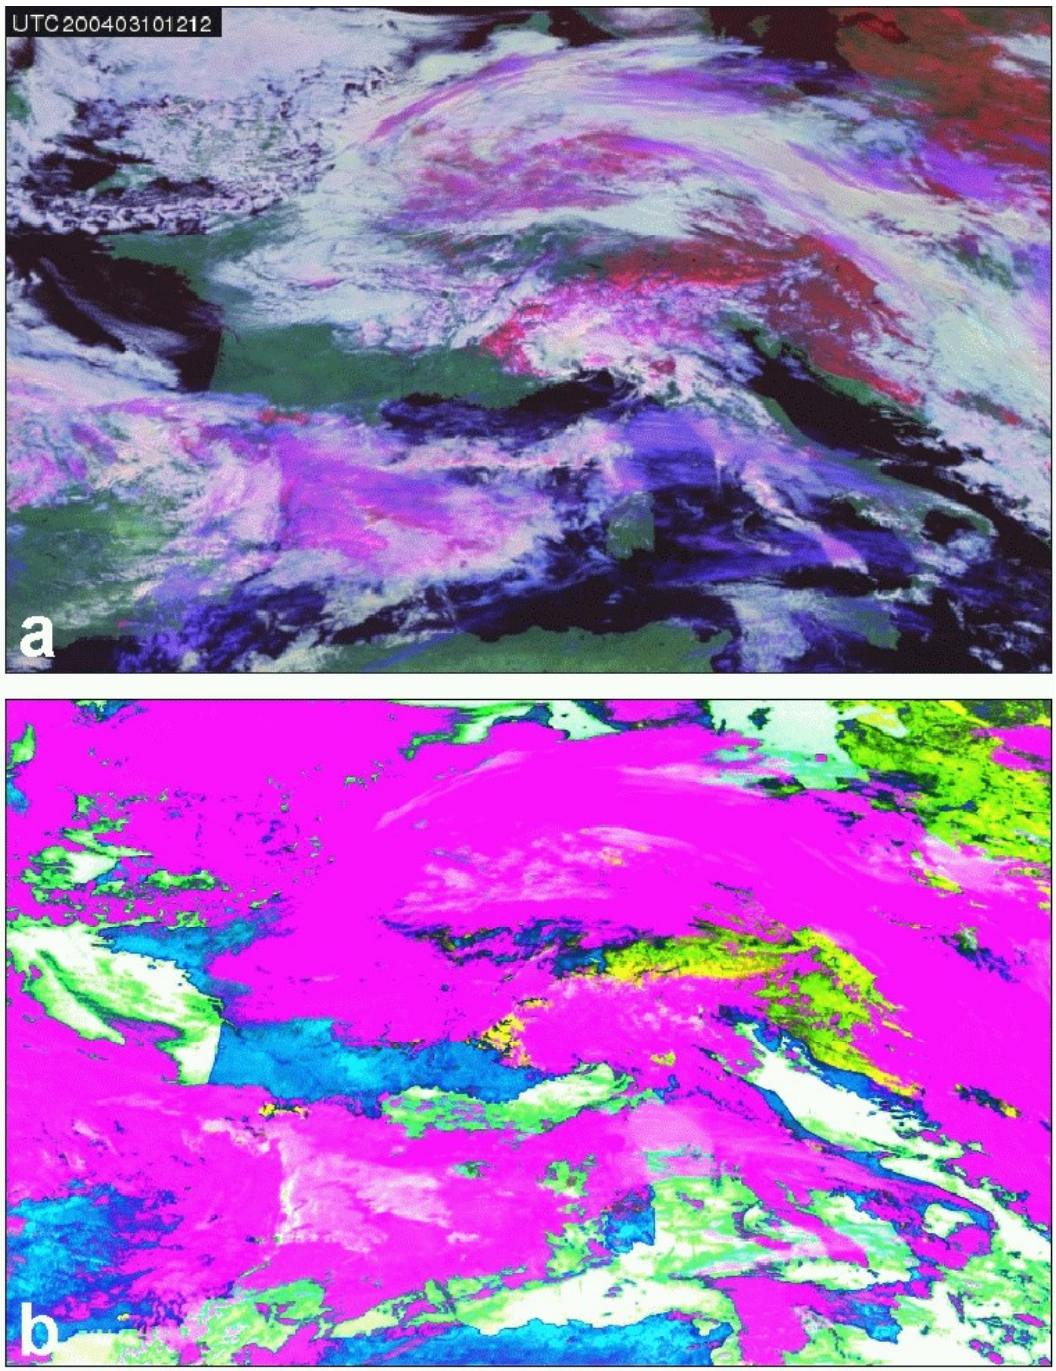

Figure 6: Meteosat-8 RGB images of central Europe, acquired on March 10^{th}, 2004, 12:12 UTC. (a) combination of $r_{0.64}$ (red), $r_{1.6}$ (green) and $(BT_{3.9} - BT_{10.8})/(BT_{3.9} - BT_{13.4})$ (blue). Snow-free surfaces are green, snow is read, water clouds are white, optically thin ice clouds (cirrus) tend to be purple and optically thick ice clouds are pink or red. (b) combination of the conditional probabilities for snow-free surfaces (red), snow (green) and water clouds (blue). Here, snow-free surfaces appear blue, snow appears green or yellow and clouds appear pink. Note that water appears black in (a) and white in (b).

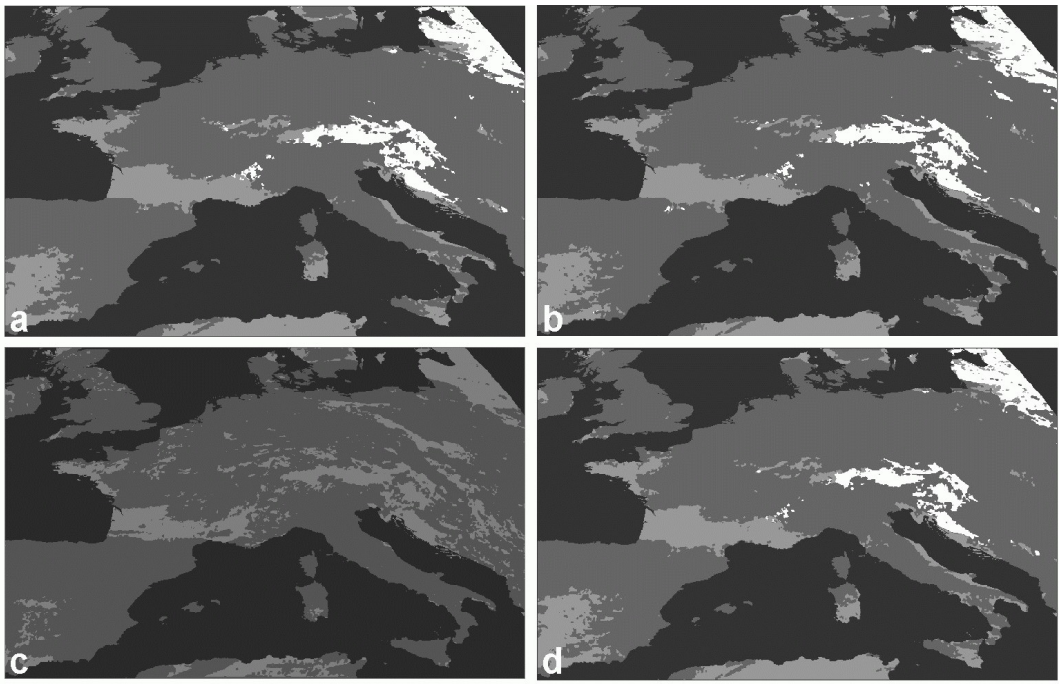

Figure 7: Classification results for March 10^{th}, 2004. (a) maximum likelihood classification with spectral and temporal features; (b) maximum likelihood with only spectral features; (c) temporal cloud mask, obtained from stacking all single channel cloud masks; (d) as (b), but now the temporal cloud mask of plot (c) is used to mask snowy pixels with high temporal variability.

6 Conclusions

With the spectral features that we use for maximum likelihood classification, all pixels that were classified as snow actually contained snow, as judged by visual inspection. Many of these pixels were of mixed type, representing both snow and snow-free land, and sometimes also transparent and/or sub-pixel clouds, mainly near cloud edges. This type of classification thus produces a liberal snow map, in the sense that it detects the highest amount of pixel where snow is to some extent present. No false positives and only some false negatives are present in this snow map. When temporal information is used to filter out pixels with high temporal variability, false negatives are removed and the snow map becomes more conservative (more mixed snow/cloud pixels are classified as clouds). The liberal snow map could be of use when one wants to obtain a binary snow map, i.e. a snow map that simply indicates whether snow is present or not. When a fractional snow map is required, the conservative snow map is more appropriate, because then the detected snow pixels are more likely to be cloud free and to contain only contributions from surface classes.

Apart from improving the detection of clouds during day-light, we anticipate that temporal information can also be used for cloud detection during the night. Of course, the solar channels can not be used during the night, but the temporal variability will still be measurable in the infrared channels. For the purpose of detecting surface snow cover however, which is not possible during the night, this is of no importance.