


Detection of circular patterns in high-resolution satellite images of agricultural land with CultSearcher



Note no

SAMBA/16/08

Authors

**Øivind Due Trier
Siri Øyen Larsen
Rune Solberg
15 May 2008**

Date

Norsk Regnesentral

Norsk Regnesentral (Norwegian Computing Center, NR) is a private, independent, non-profit foundation established in 1952. NR carries out contract research and development projects in the areas of information and communication technology and applied statistical modelling. The clients are a broad range of industrial, commercial and public service organizations in the national as well as the international market. Our scientific and technical capabilities are further developed in co-operation with The Research Council of Norway and key customers. The results of our projects may take the form of reports, software, prototypes, and short courses. A proof of the confidence and appreciation our clients have for us is given by the fact that most of our new contracts are signed with previous customers.

| | |
|---------------------------|--|
| Title | Detection of circular patterns in high-resolution satellite images of agricultural land with CultSearcher |
| Authors | Øivind Due Trier Siri Øyen Larsen Rune Solberg |
| Date | 15 May 2008 |
| Year | 2008 |
| Publication number | SAMBA/16/08 |

Abstract

Many archaeological sites are discovered during building and road construction work, prompting full excavations and delay in construction. In order to detect more cultural heritage sites in advance of construction work, the Norwegian Directorate for Cultural Heritage has taken an initiative to develop tools for early detection of potential cultural heritage sites. The Norwegian Computing Center has been responsible for developing automatic detection methodology and implementing this into a prototype software system, CultSearcher. The system was previously tailored to detection of amorphous patterns. The current version of the system has been extended to circular patterns, like ring graves.

We have developed an approach with the following steps. First, the image is band-pass filtered in order to remove information that does not contribute to the rings. This step is optional. Next, the image is contrast enhanced, in order to make the weak rings more distinct. Then, the image is convolved with ring templates of varying sizes, giving high absolute values at candidate ring locations. Finally, the ring candidates are presented to a human operator, who may reject some of the detected ring candidates.

In order to test our method, we used Quickbird images from south-east Norway with 0.6m pixel size, from a few years back. Of the rings that were clearly visible in the images, 73% was detected, and of the ones that were fairly visible, 50% was detected. In addition, seven times as many false positives as true positives were detected.

All in all, this demonstrates that our approach is a helpful tool that may assist archaeologists in locating many ring graves and other circular archaeological sites, but that many may also be missed in individual images.

| | |
|------------------------|---|
| Keywords | Ring graves, Quickbird, Ikonos, high resolution satellite images, Fourier transform, band pass filtering, template matching, contrast enhancement, archaeology, remote sensing. |
| Target group | Operators and developers, archaeologists |
| Availability | Open |
| Project number | 236 067 |
| Research field | Earth observation |
| Number of pages | 53 |

Contents

| | |
|---|-----------|
| Executive summary | 6 |
| 1 Introduction | 7 |
| 1.1 Background | 7 |
| 1.2 The CultSearcher prototype system..... | 8 |
| 1.3 Objectives..... | 8 |
| 1.4 Overview..... | 9 |
| 2 Experimental data set | 10 |
| 2.1 Satellite image data | 10 |
| 2.2 Identified rings | 13 |
| 3 Algorithm | 32 |
| 3.1 Approaches tested and associated problems | 32 |
| 3.1.1 Local contrast enhancement..... | 32 |
| 3.1.2 Filtering in the frequency domain..... | 33 |
| 3.1.3 Template matching | 38 |
| 3.1.4 Feature extraction | 40 |
| 3.1.5 Classification based on a decision tree | 43 |
| 3.1.6 Combinations of the above methods | 43 |
| 3.2 Overall description of the best approach..... | 44 |
| 3.3 Detailed description of the best approach..... | 45 |
| 4 Experimental results | 48 |
| 5 Discussion | 50 |
| 6 Conclusion | 51 |
| Acknowledgements | 52 |
| References | 52 |

Executive summary

Many archaeological sites are discovered during building and road construction work, prompting full excavations and delay in construction. In order to detect more cultural heritage sites in advance of construction work, the Norwegian Directorate for Cultural Heritage has taken an initiative to develop tools for early detection of potential cultural heritage sites. The Norwegian Computing Center (Norsk Regnesentral, NR) has been responsible for developing automatic detection methodology and implementing this into a prototype software system, CultSearcher. The system is currently analysing soil-marked and crop-marked patterns.

CultSearcher was previously tailored to detection of amorphous patterns. The last version of the system has been extended to circular patterns. Archaeological sites with circular structures, like ring graves, are sometimes visible in satellite images as soil marks or crop marks. Their visibility tends to vary during the year, and to local variations such as water content in the ground. At best, the rings are distinct, but tend to have less contrast to the background than many other objects in the images. Consequently, reliable automated detection based on pattern recognition is very difficult.

We have developed an approach with the following steps. First, the image is band-pass filtered in order to remove information that does not contribute to the rings. Note that this step is optional. Next, the image is contrast enhanced, in order to make the weak rings more distinct. Then, the image is convolved with ring templates of varying sizes, giving high absolute values at candidate ring locations. Finally, the ring candidates are presented to a human operator, who may reject some of the detected ring candidates.

In order to test our method, we used Quickbird images with 0.6m pixel size, from a few years back. Archaeologists have identified three kinds of potential archaeological site: (1) verified by subsequent excavation, (2) surveyed in the field but not excavated, and (3) ring structures only seen in the images and considered as potential archaeological sites. Of the above rings that were clearly visible in the images, 73% was detected, and of the ones that were fairly visible, 50% was detected. In addition, seven times as many false positives as true positives were detected. The number of false positives can be reduced, at the cost of reducing the number of true positives as well. For example, by reducing the number of false positives from 7 times to 0.5 times the number of true positives, the number of detected rings decreased from 64% to 32%.

All in all, this demonstrates that our approach is a helpful tool that may assist archaeologists in locating potential archaeological sites, but that many may also be missed in individual images. For a thorough search in a limited area, a high number of false positives might be acceptable. On the other hand, for massive search through a large number of images, the number of false positives might be kept at a minimum, as long as *some* sites are detected. Some sites may only be visible from time to time. In order to find these, one may have to process images from, say, a ten year period, and, say, 5-10 images per year. In conclusion, our approach can be used to process large volumes of satellite images that would otherwise not be inspected, thus detecting many new sites.

1 Introduction

This document describes new functionality in the CultSearcher software prototype. The new functionality is for the detection of circular structures in high-resolution satellite images. Such ring structures are potentially ring graves and other archaeological phenomena.

1.1 Background

The increasingly intensive use and modification of the landscape resulting from modern demands for efficient infrastructure and land use (agricultural production, mining, energy sources, leisure/tourism facilities, etc.) exerts growing pressure on cultural heritage in the landscape. In order to match the political intentions of updated and sustainable cultural heritage management, it is necessary to develop a cost-effective method for locating and monitoring cultural heritage sites. Given the enormous costs of surveying the areas in question by traditional fieldwork, alternatives must be sought. The use of modern support technologies is imperative, if such rapid changes are to be balanced against the sustainable management of this resource. One possible approach is through the use of satellite images.

In recognition of this, the Norwegian Directorate for Cultural Heritage (in Norwegian, Riksantikvaren, RA), in collaboration with the Norwegian Computing Center (Norsk Regnesentral, NR), the Norwegian Institute for Cultural Heritage Research (Norsk Institutt for Kulturminneforskning, NIKU), the Museum of Cultural History at the University of Oslo (Kulturhistorisk Museum) and Vestfold County Administration (Vestfold fylkeskommune) started in 2003 a project with the overall aim of developing a cost-effective method for surveying and monitoring cultural heritage sites on a regional and national scale. NR has in this project been responsible for developing the automatic detection methodology and implementing this into a prototype software system. Previous work in this project is described in [3][4][5].

Results obtained in a 2002 pilot project indicated the existence of a correlation between cultural heritage sites and variation in the chemical elements in the soil. A central focus in the early project years was the manual analysis of satellite images followed by chemical profiling of sites observed in these images in order to gain experience as to how cultural heritage sites really manifest themselves in satellite images. The results demonstrated that high-resolution geo-chemical sampling appears to be a promising field for the development of cultural heritage indicators. However, the costs involved demanded a need for funding which was impossible to obtain.

It was then suggested to focus on the development of automated methods, such as pattern recognition, for detecting and locating cultural heritage sites. The working assumption is that cultural heritage sites with no visual apparent manifestations above ground may be detectable in remote sensing images due to alterations in the spectral signature of the bare soil or of uniform vegetation growing there (crops).

The use of aerial imagery for this purpose is quite widespread (e.g., see [7][13]). In addition to airborne multispectral imaging, recent advances in airborne LIDAR also show great promise (e.g., see [15][16]). The collection of airborne imagery and LIDAR is nevertheless both time consuming and costly and this imposes severe limitations on the size of the areas that can be

investigated. In this respect satellite imagery holds great promise. Satellite image based location, surveillance and monitoring of cultural heritage sites has been the subject of some recent research (e.g., see [6][8][10]), but generally this field is not well explored.

Although the costs connected with acquiring and analyzing the satellite data will not be insignificant, and fieldwork will never be replaced entirely by high-technological methods, it seems plausible that an essentially cheaper and possibly even qualitatively better method for the surveying and monitoring of cultural heritage sites can be developed to target fieldwork to a degree not possible today.

1.2 The CultSearcher prototype system

CultSearcher is currently analysing soil-marked and crop-marked patterns. Soil-marked sites are typically the remains of a ditch or pit, buried walls, etc. A ditch or a pit would disturb the local soil profile, and refilled material usually has different characteristics, like density and composition. The refilled material is in most cases not so compact, and it might contain more humus components, making it looking darker. The refilled material may also affect the soil texture with a grain-size distribution that differs from the undisturbed soil (usually larger number of smaller grain sizes). This results in improved water-storage capacity, so the soil would look darker under certain conditions.

Crop marks are an indirect effect of buried archaeological features. Their visibility depends on the soil, climate and vegetation. So-called positive marks are due to more water available that makes plants grow higher and ripen later than the plants around. A colour-tonal contrast may be created because the vegetation stays green for a longer period and/or that the vegetation is darker green. Crop marks may also be due to a vegetation relief. Plants grow higher, enough to throw a shadow in slanting sunlight. So-called negative marks appear when plants grow over buried stones (e.g. walls) and run out of water sooner, ripen earlier and stay shorter. Almost any crop can develop marks, if conditions are well. Cereals react fast on a Soil Moisture Deficit (SMD) and are growing very close, making contrasts clearer.

Various types of remote sensing sensors, airborne and spaceborne, are useful for detecting remains or patterns due to cultural heritage sites. Soil- and crop-marked sites can be measured with high-resolution optical (visible and infrared) sensors. With the optimal selection of observation wavelengths, high contrast can be obtained (in particular appearing from reflectance contrasts due to soil moisture or vegetation density). The spatial resolution of these sensors should be of 1 m or better to be really useful. Therefore, the project has so far focussed on images from Ikonos and QuickBird.

1.3 Objectives

The aim of the software prototype described in this report is to provide computerized assistance to the operator in the analysis of satellite images. In particular, the software identifies circular structures, e.g. potential sites for ring graves, for further inspection by an archaeologist. This means that the archaeologist may concentrate on analyzing the identified sites rather than the entire image.

The aim of this document is to describe and discuss the various approaches that have been tried out for circular structure detection, both the methods that didn't work and the ones that did work. The delivered software is documented in [2].

1.4 Overview

The rest of the document is organized as follows.

- Section 2 describes the test images, and highlights the manually identified rings in these images.
- Section 3 first describes and discusses the various approaches that were tried out, then describes the chosen sequence of methods.
- Section 4 describes the experimental results
- Section 5 discusses the experimental results
- Section 6 gives conclusions and recommendations for further work

2 Experimental data set

In this section, the experimental data set is described.

2.1 Satellite image data

The data set consists of two Quickbird images (<http://www.eurimage.com/products/quickbird.html>).

- Lågendalen, a valley between Kongsberg and Larvik. This image was taken on 27 April 2005 at 10:45 AM (Figure 1 a).
- Gardermoen, the area surrounding Oslo's air port. This image was taken on 29 July 2003 at 10:23 AM (Figure 1 b).




Figure 1. The two Quickbird images: (a) Lågendalen, (b) Gardermoen.

Both images consist of a four-band multi-spectral image and a panchromatic (grey scale) image. The panchromatic image has 0.6 m wide pixels, and the single band covers the 450-900 nm wavelengths. The multi-spectral image has 2.4 m wide pixels, and the four bands are: blue (450-520 nm), green (520-600 nm), red (630-690 nm) and near-infrared (760-900 nm).



Figure 2. Detail of the panchromatic image of Lågendalen, with four clearly visible ring graves.

Many circular patterns are clearly visible in the panchromatic images (Figure 2), but can hardly be seen in the multi-spectral images (Figure 3 and Figure 4). Recently, other research groups have used multi-spectral Quickbird [11] or Ikonos [9] images, but the objects they were looking for were much larger than the ring graves in the present work. Since the circular patterns are difficult to spot visually in the multi-spectral images, we chose to use only the panchromatic images.



Figure 3. Detail of the multi-spectral image of Lågendalen. The blue, green and red bands are displayed in their natural colors.




Figure 4. Detail of the near-infrared band of the multi-spectral image of Lågendalen

2.2 Identified rings

35 rings have been identified in the test images. Each of the identified rings is described and shown in a figure. The left half of the figure displays the ring using the standard grey level values, and the right half displays the ring in a contrast enhanced image. In each half, three different zoom factors are used. At the lower left, the entire image is shown. The little red square in this part indicates the part shown in the upper 2/3. And the red square in the centre of the upper 2/3 indicates the part shown in the zoom window at the lower right.

Each ring is characterized as one of

1. Strong ring
2. Fairly strong ring
3. Weak ring

There are 15 strong rings, 10 fairly strong rings, and 10 weak rings. At a few locations, there are two rings located only a few metres apart; these are counted as one ring only, with the understanding that if one is correctly located, then the archaeologist has been directed to an interesting site, and whether one or two rings are marked by the program is not important.

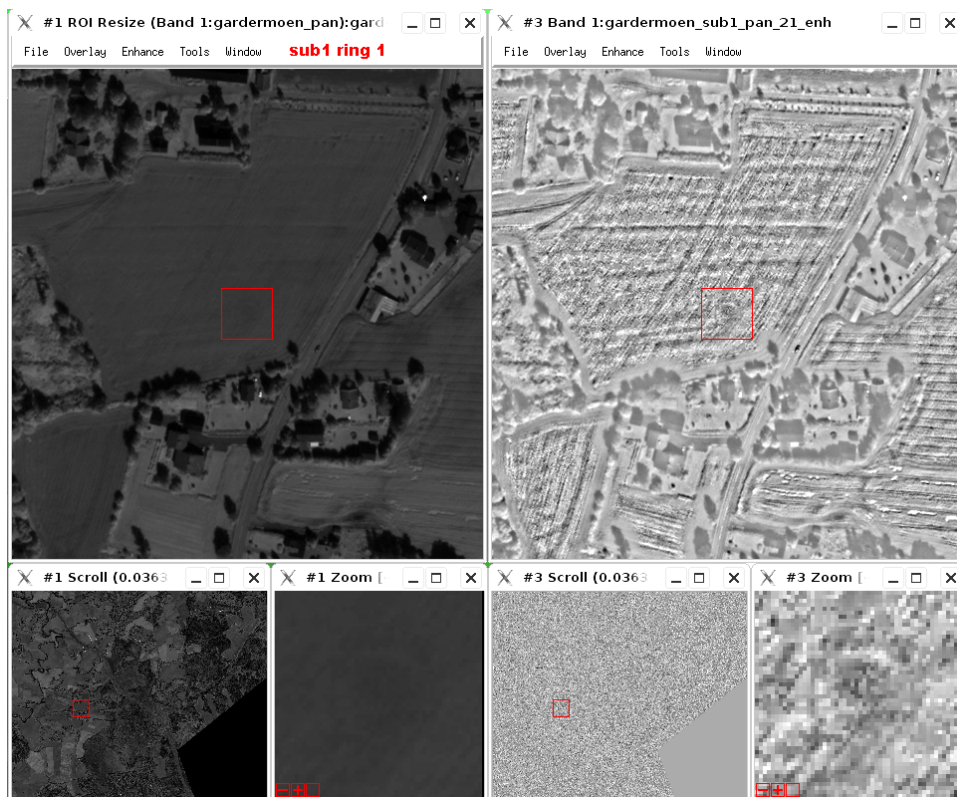


Figure 5. Ring number 1 in the Gardermoen sub1 test image. This ring is weak.




Figure 6. Ring number 2 in the Gardermoen sub1 test image. This ring is in a forested area.




Figure 7. Ring number 3 in the Gardermoen sub1 test image. This ring is fairly strong.

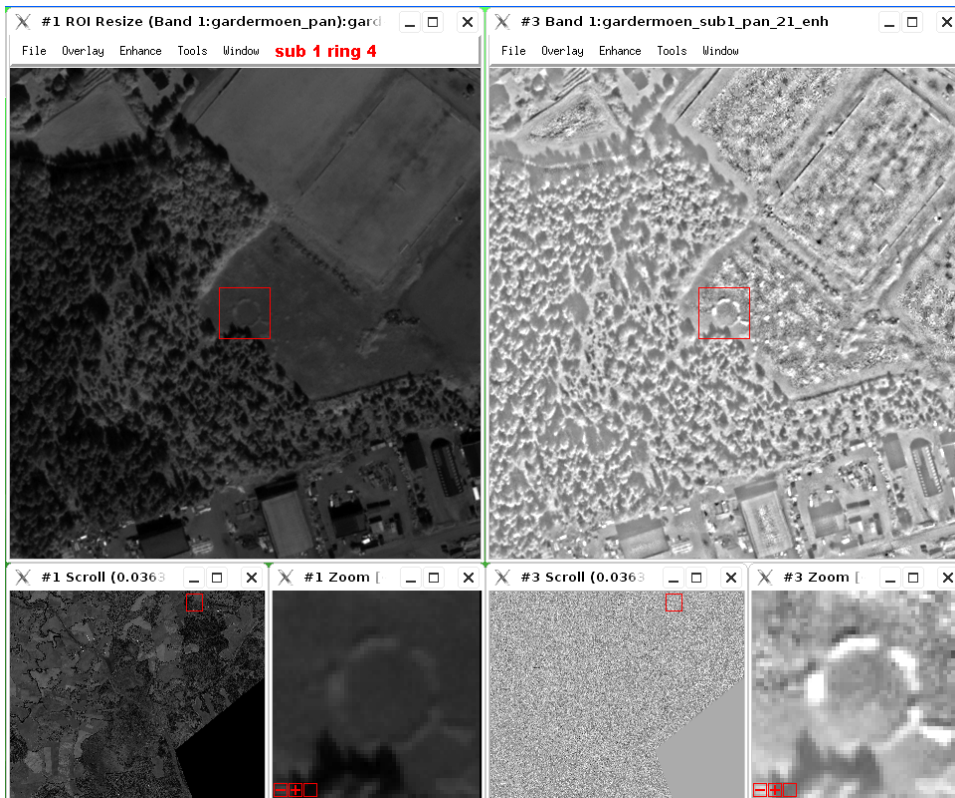


Figure 8. Ring number 4 in the Gardermoen sub1 test image. This ring is strong but fragmented.



Figure 9. Ring number 5 in the Gardermoen sub1 test image. This ring is strong and thick, almost like a disk.




Figure 10. Ring number 6 in the Gardermoen sub1 test image. This is two overlapping, weak rings.




Figure 11. Ring number 1 in the Gardermoen sub2 test image. This ring is fairly strong.

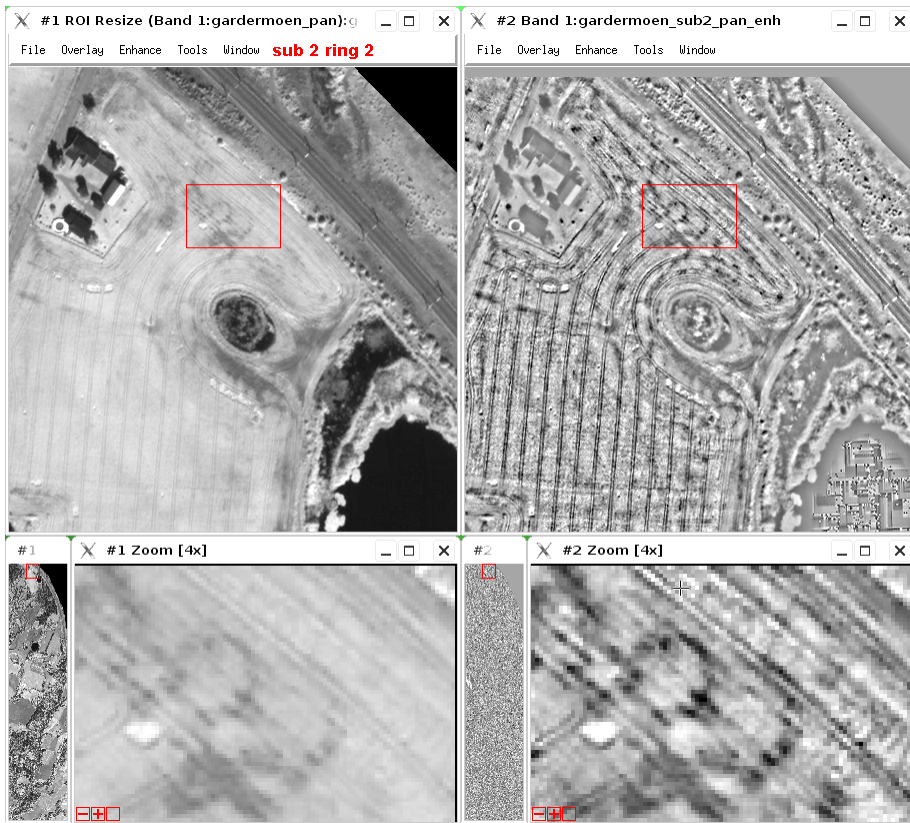


Figure 12. Ring number 2 in the Gardermoen sub2 test image. This is two touching rings; one strong to the left and one fairly strong at the right.




Figure 13. Ring number 3 in the Gardermoen sub2 test image. This ring is weak.




Figure 14. Ring number 4 in the Gardermoen sub2 test image. This ring is fairly strong.

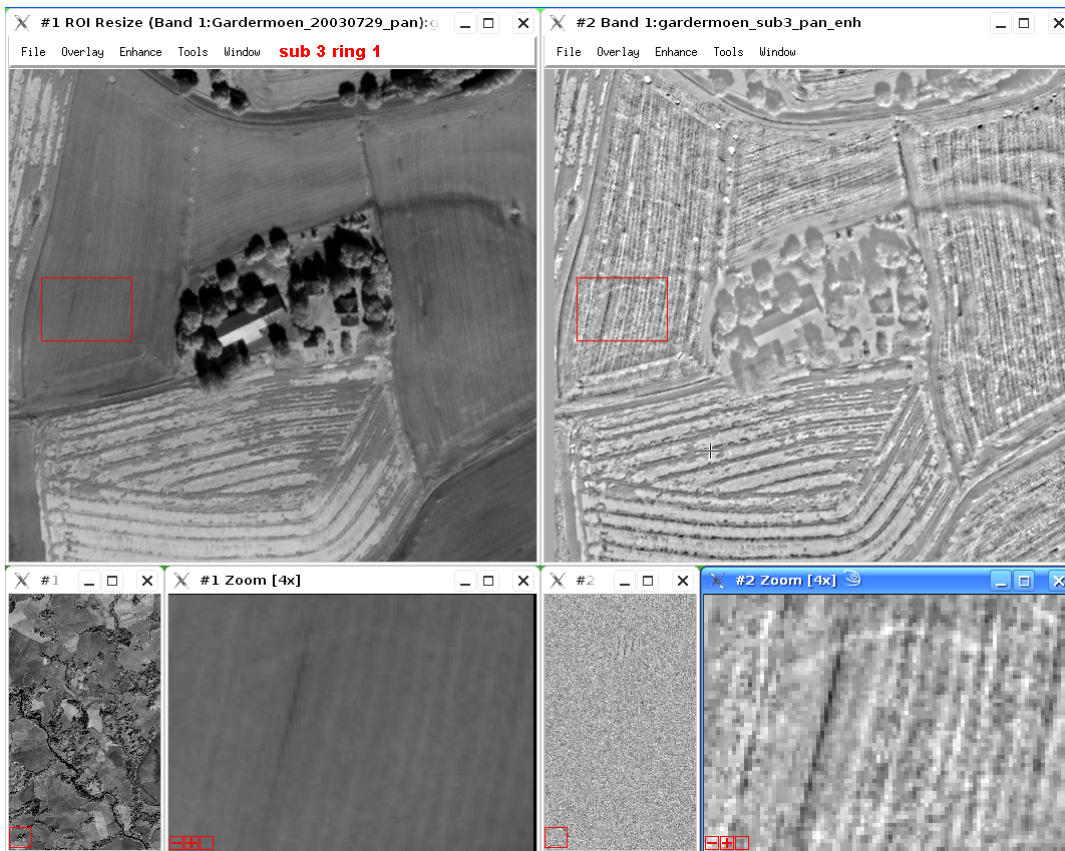


Figure 15. Ring number 1 in the Gardermoen sub3 test image. This ring is weak.

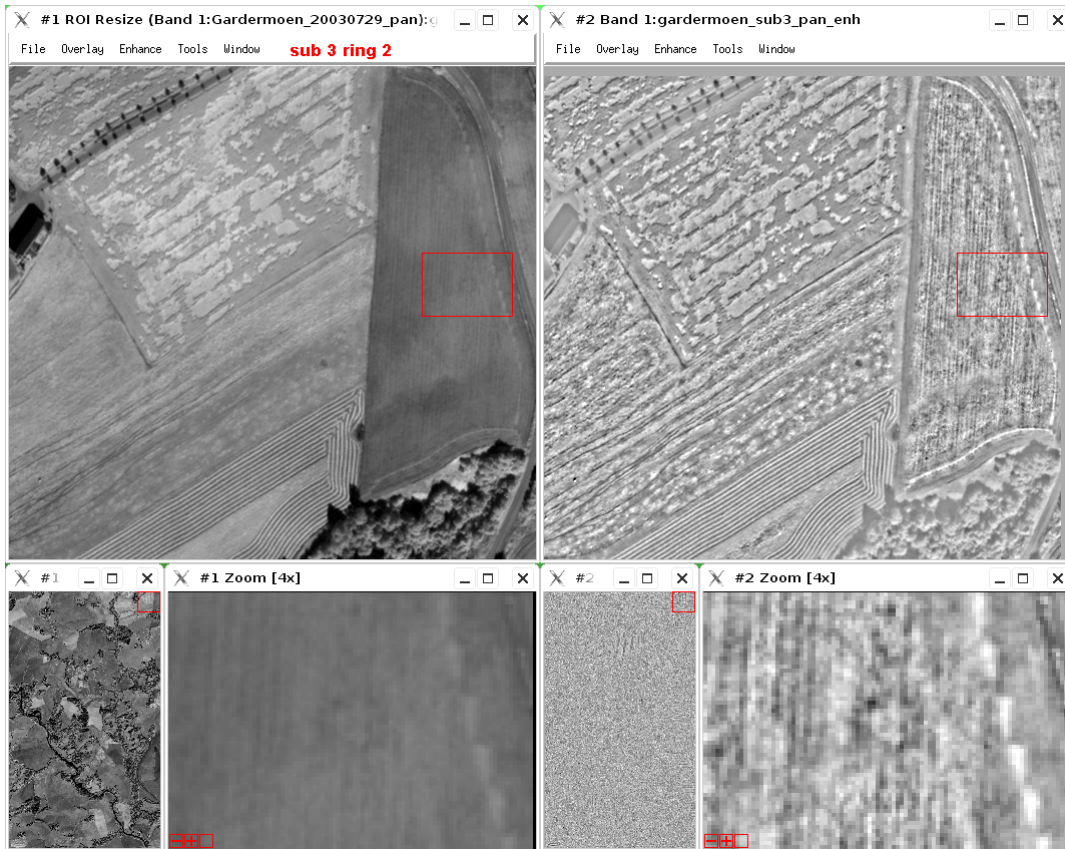


Figure 16. Ring number 2 in the Gardermoen sub3 test image. This ring is strong.




Figure 17. Ring number 3 in the Gardermoen sub3 test image. This ring is strong.




Figure 18. Ring number 4 in the Gardermoen sub3 test image. This ring is fairly strong.

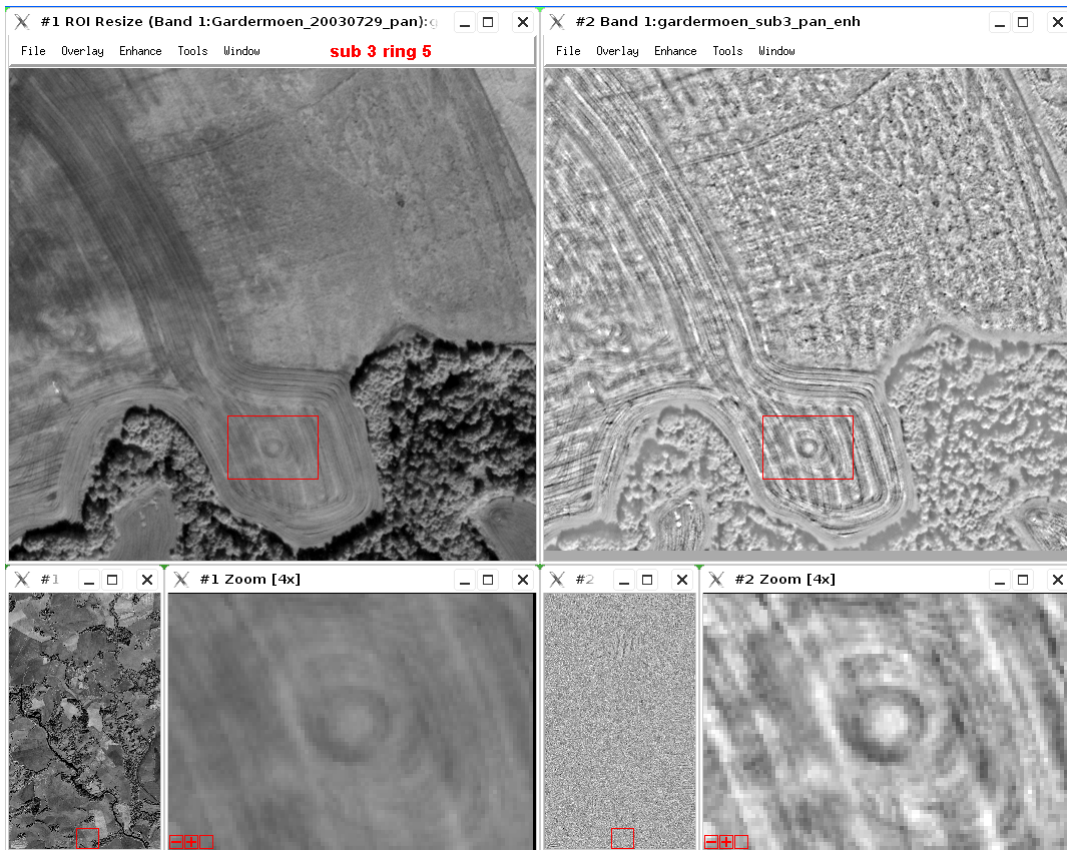


Figure 19. Ring number 5 in the Gardermoen sub3 test image. This ring is strong.




Figure 20. Ring number 1 in the Gardermoen sub4 test image. This ring is fairly strong.




Figure 21. Ring number 2 in the Gardermoen sub4 test image. There are two fairly strong rings.



Figure 22. Ring number 1 in the Laagen sub1 test image. This ring is strong.




Figure 23. Ring number 2 in the Laagen sub1 test image. This ring is weak.

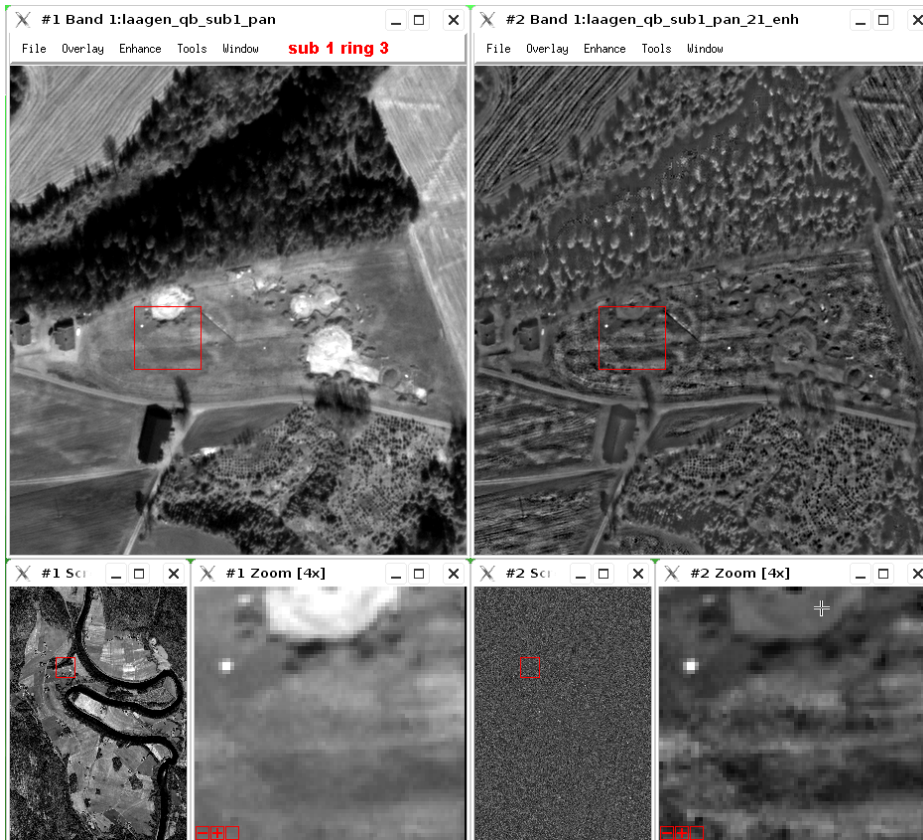


Figure 24. Ring number 3 in the Laagen sub1 test image. This ring is weak.




Figure 25. Ring number 4 in the Laagen sub1 test image. This ring is fairly strong.



Figure 26. Ring number 5 in the Laagen sub1 test image. This is a strong ring.



Figure 27. Ring number 6 in the Laagen sub1 test image. This is a strong disk.

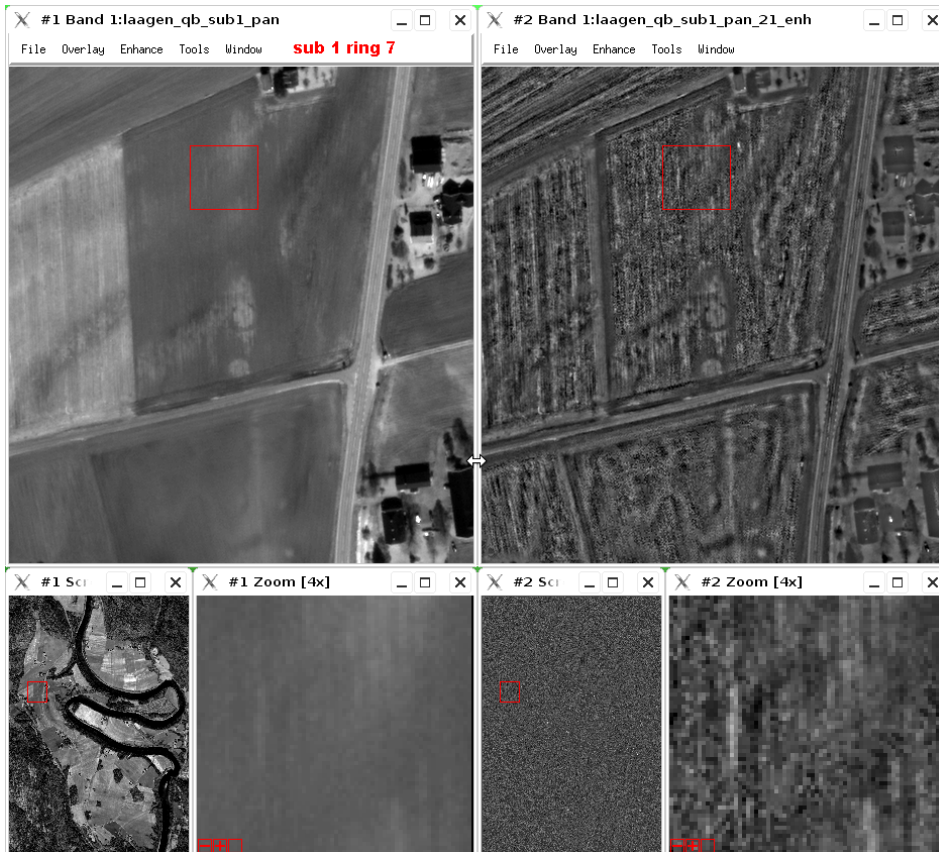


Figure 28. Ring number 7 in the Laagen sub1 test image. This is a weak ring.

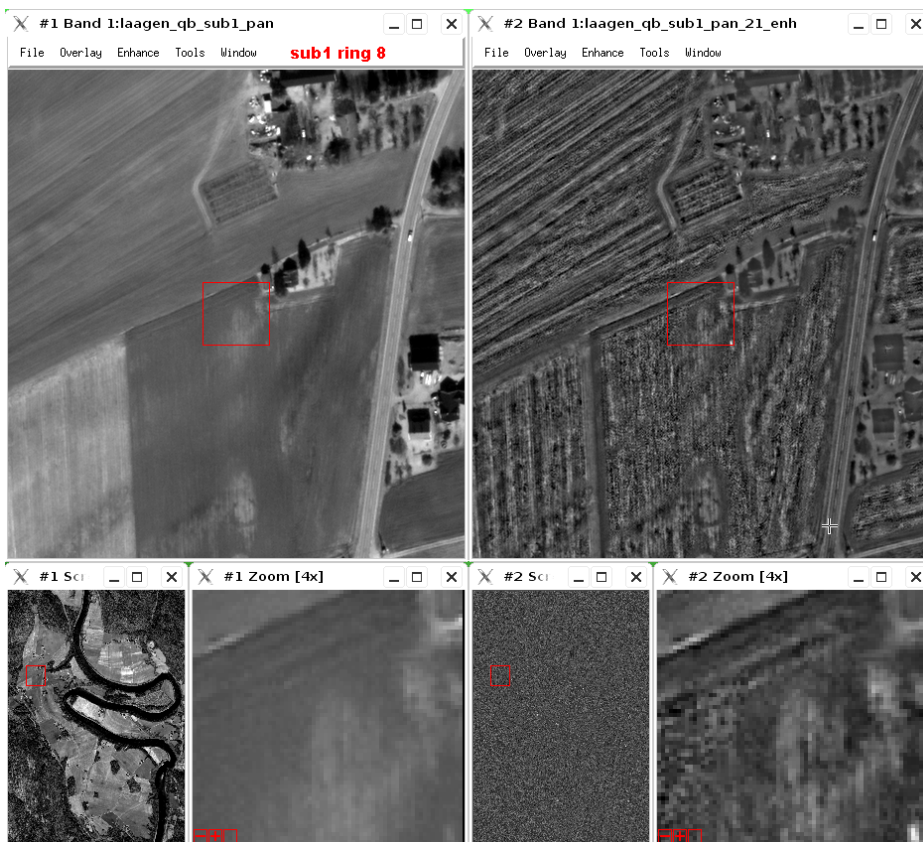


Figure 29. Ring number 8 in the Laagen sub1 test image. This ring is weak.

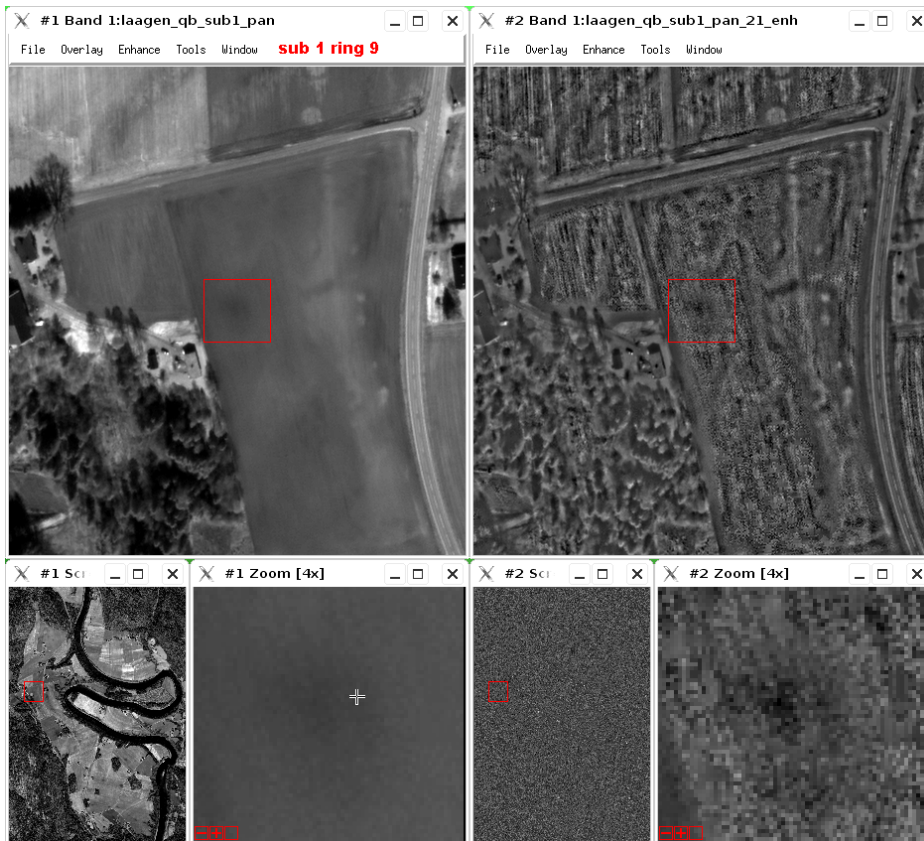


Figure 30. Ring number 9 in the Laagen sub1 test image. This ring is weak.

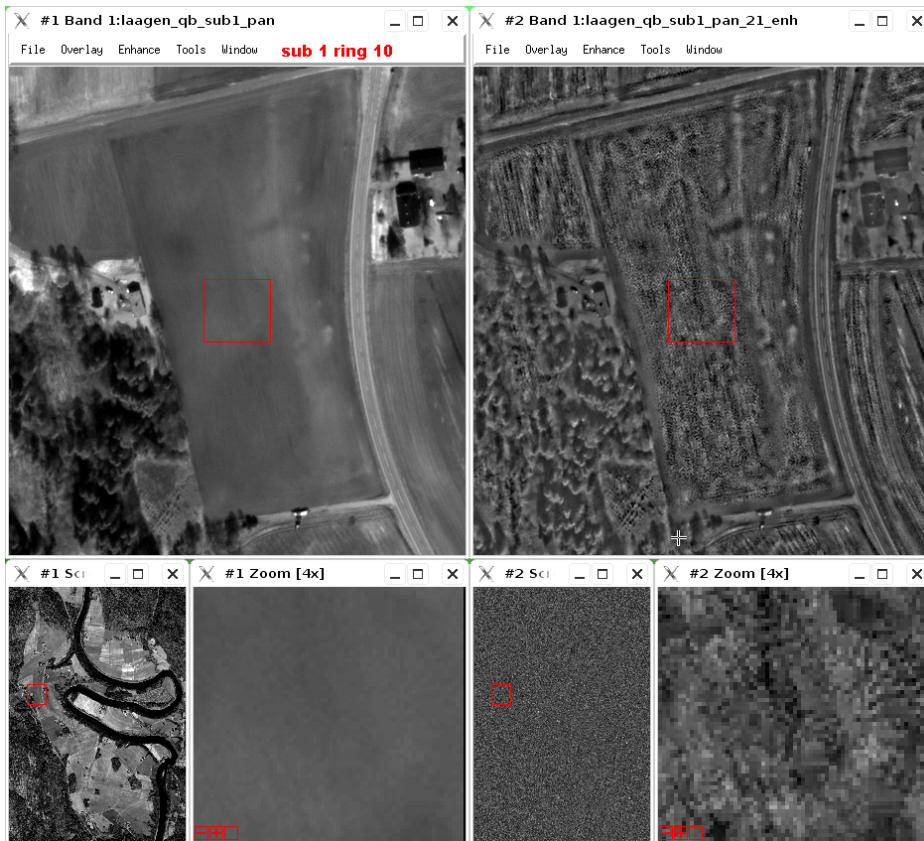


Figure 31. Ring number 10 in the Laagen sub1 test image. This ring is weak.

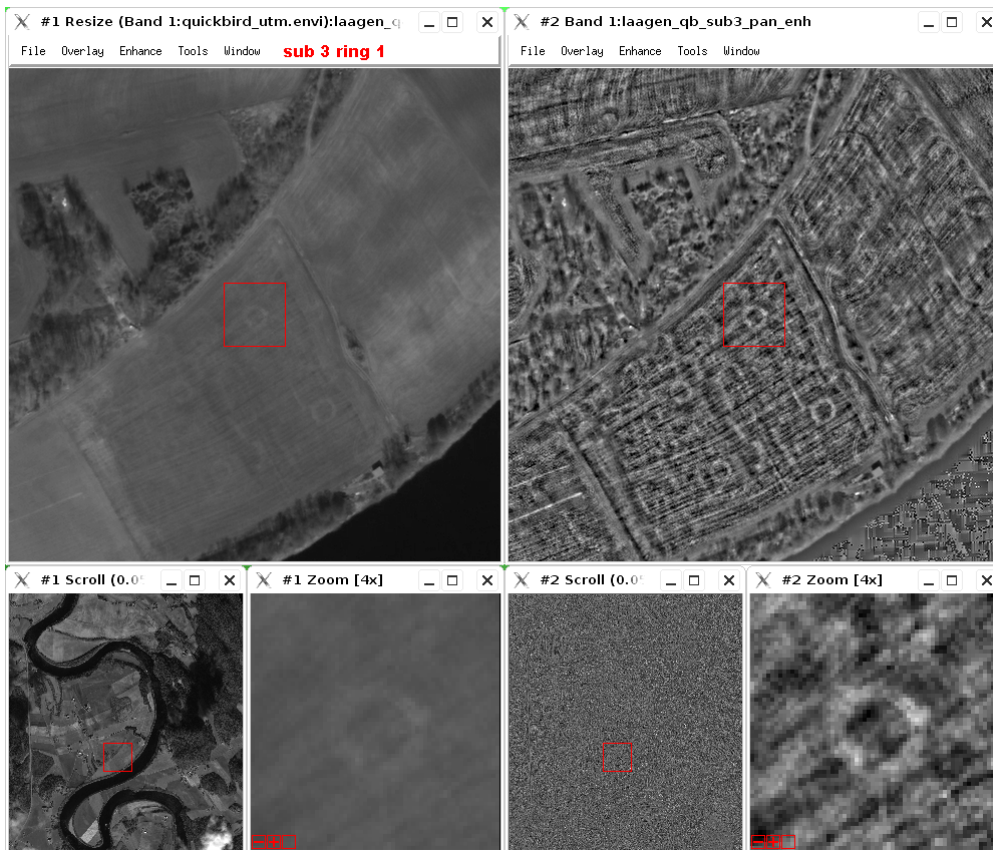


Figure 32. Ring number 1 in the Laagen sub3 test image. This ring is strong.




Figure 33. Ring number 2 in the Laagen sub3 test image. This ring is strong.




Figure 34. Ring number 3 in the Laagen sub3 test image. This ring is strong.

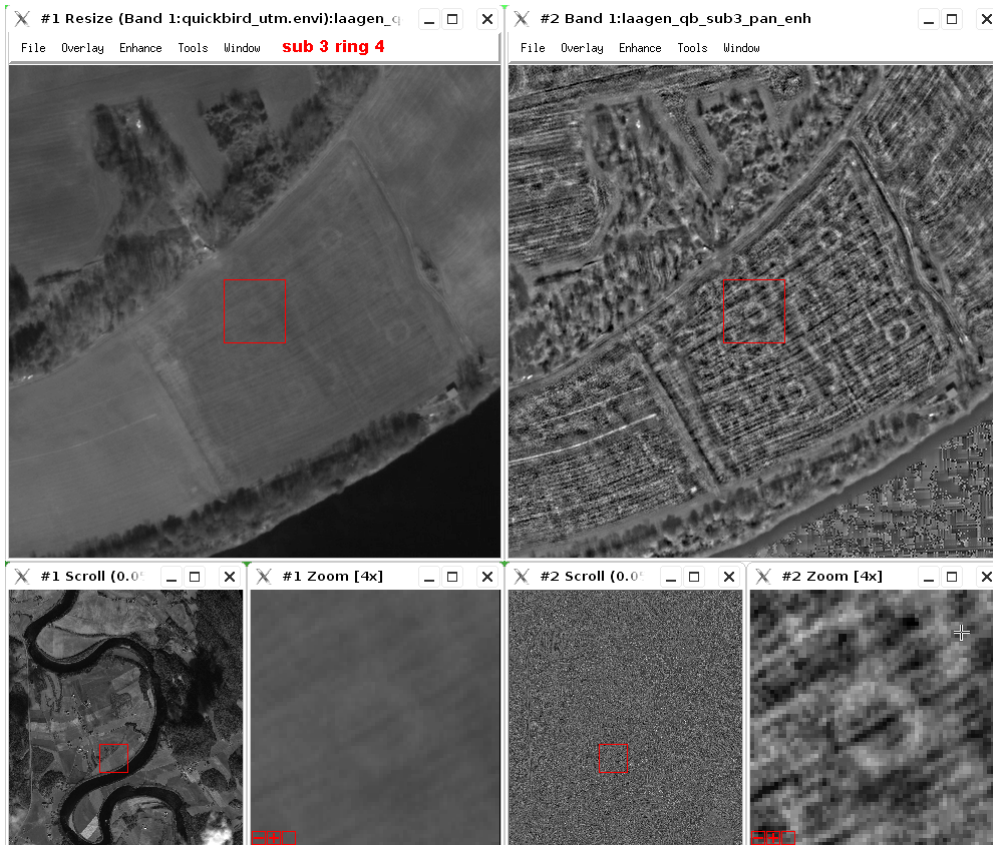


Figure 35. Ring number 4 in the Laagen sub3 test image. This ring is strong.




Figure 36. Ring number 5 in the Laagen sub3 test image. This ring is weak.




Figure 37. Ring number 6 in the Laagen sub3 test image. This ring is weak.




Figure 38. Ring number 7 in the Laagen sub3 test image. This ring is strong.

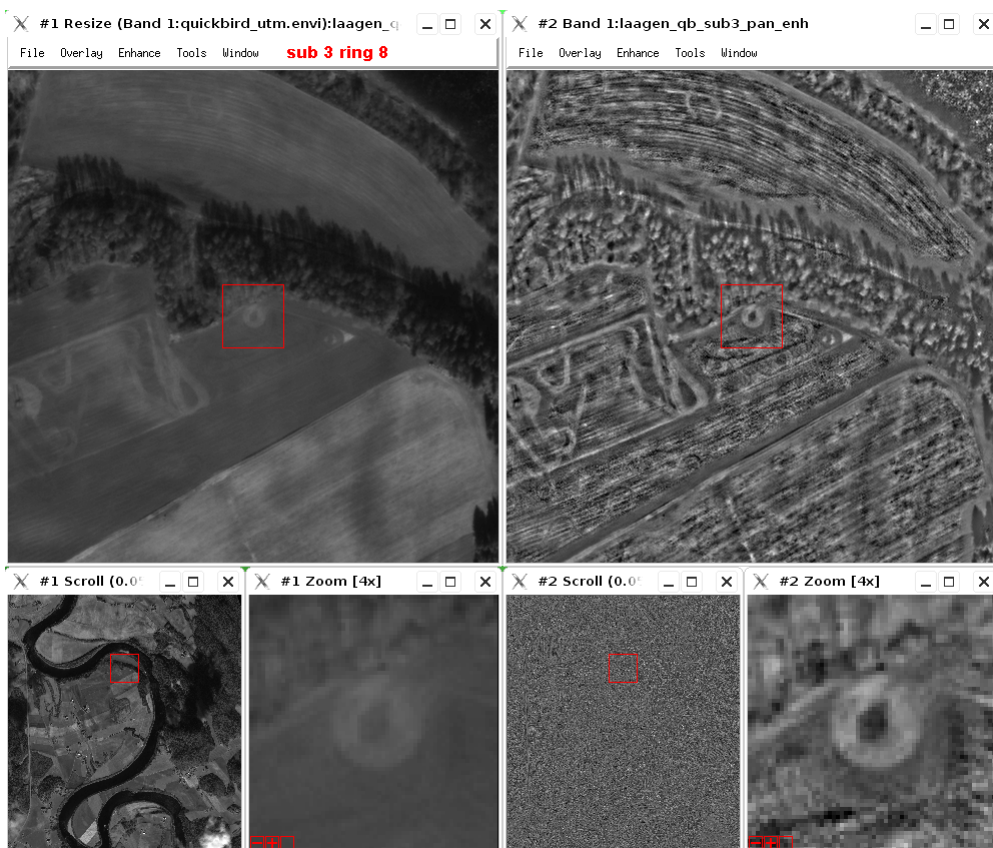


Figure 39. Ring number 8 in the Laagen sub3 test image. This ring is strong.

3 Algorithm

In this section, both the various approaches tested and the chosen methods are described. We have used the Envi image software (<http://www.ittvis.com/envi/>), and written program code in the IDL programming language (<http://www.ittvis.com/idl/>), with calls to Envi functions.

3.1 Approaches tested and associated problems

In order to detect as many rings as possible, while at the same time keeping the number of false positives at a minimum, a number of various methods have been tried out. These are

1. Local contrast enhancement
2. Template matching
3. Feature extraction
4. Decision tree-based classification
5. Low pass-, band pass- and high pass-filtering in the frequency domain

Each of these will be discussed below

3.1.1 Local contrast enhancement

Even the most distinct rings in the test images have relatively low contrast with their surroundings. In order to be able to detect any rings at all, the local contrast has to be more or less constant over the entire image. This can be achieved by, for each pixel, computing the local mean grey level and associated standard deviation in an $N \times N$ neighbourhood centred on the pixel. The pixel value $p_{CE}(x, y)$ in the contrast enhanced image is computed as

$$p_{CE}(x, y) = \frac{p(x, y) - \mu(x, y, N)}{\sigma(x, y, N)}$$

Here, $p(x, y)$ is the grey level value in the input image, $\mu(x, y, N)$ is the mean grey level value in an $N \times N$ neighbourhood centred on (x, y) , and $\sigma(x, y, N)$ is the standard deviation of the grey level values in the same neighbourhood.

The choice of the neighbourhood size, N , doesn't seem to be critical. We have chosen to use $N=21$. However, using, say, $N=15$ or $N=35$ also works quite well. Having a too small value for N may result in too much exaggeration of even small local variations in grey level. Having a too large value for N may suppress the local contrast needed to identify the rings. One undesirable effect that is present when using $N=21$ is that the local contrast is suppressed when the central pixel is less than $N/2$ pixels from a very dark or very bright object in the image. This can be seen a little south of ring number 3 in the "Laagen sub3" test image (Figure 34). Along the row of trees, the plough furrows have almost been suppressed. Similarly, in the river, there are bands of almost homogeneous grey values along its banks. The river is of little concern to us, but the band along tree rows may make it difficult to detect rings near the borders of fields.

3.1.2 Filtering in the frequency domain

The use of the Fourier transform is popular in one dimensional signal processing, where one is often interested in signals within a frequency band, and any signal in other frequencies is regarded as noise.

A two dimensional version of the Fourier transform exists, which transforms an image to the frequency domain as an image of the same size. In the frequency domain, one may apply a low pass, band pass or high pass filter, followed by the inverse two dimensional Fourier transform. The result is an image where some frequencies have been removed.

If the image has width = height, and the width is a power of 2, then the fast Fourier transform, FFT, will preserve all the information of the original image in the transformed one, so that the inverse transform brings back the original image. We have used 4096×4096 pixels as the image size.

3.1.2.1 Experiments

We tested the below filters. Band pass 20 – 40 means that in the frequency domain, a mask with inner radius 20 pixels and outer radius 40 pixels was used, and, broadly speaking, the pixels in the frequency image that were not within these radii were set to zero. High pass 10 means that a mask with radius 10 pixels was used, and the pixels in the frequency image that were inside this radius were set to zero. No low pass filters were tested, since we were interested in suppressing the general grey level distribution in the image and enhance local contrast.

- Band pass 20 – 40.
- Band pass 20 – 200.
- Band pass 100 – 500.
- Band pass 100 – 800.
- Band pass 200 – 600.
- Band pass 200 – 800.
- Band pass 200 – 2000.
- Band pass 400 – 800.
- Band pass 1000 – 2000.
- High pass 10.
- High pass 100
- High pass 200

The Envi functions for creating band pass and high pass filters were used. (In Envi, high pass is called circular cut.) These filters have smooth edges.

3.1.2.2 Results

The filter that seemed to preserve the rings best and at the same time suppress as much other information as possible, was the band pass filter with cut off radii 100 and 800.

To illustrate what information was removed and what information was kept, a number of filters were run on one test image. The test image itself has quite low contrast (Figure 41), so to be able to see anything at all in this printed document, Envi's "Enhance" functionality (Figure 40) was used to improve the global contrast visually. Note that the original image was not changed.




Figure 40 The "Enhance" menu in Envi.

The band pass filter with radii 100 and 800 seems to preserve all the information about the ring (Figure 42, bottom half), while at the same time, there is no trace of the ring in the information that was removed (Figure 42, top half).

To further visualize the different information that was removed, two more filters were applied on the test image. The low frequency information that was removed can be obtained by using a low pass filter with radius 100 pixels on the original image, and similarly, the high frequency information that was removed can be obtained by using a high pass filter with radius 800 pixels (Figure 43).




Figure 41 Original image (Gardermoen sub3) at the bottom half, and the globally contrast enhanced version at the top half.




Figure 42. Result of band pass filtering. At the bottom half, the resulting image. At the top half, the information that was removed.




Figure 43. The information that was removed. At the bottom half, the low frequency information that was removed. At the top half, the high frequency information that was removed.

3.1.3 Template matching

The principle in template matching is to have some predefined “ideal” images that we slide across the image, and for each template and each location, we compute some similarity measurement. The locations with the highest similarity values are regarded as detections.

We used ring shaped templates with radii in the range from 4.5 m to 9.0 m, with 0.5 m increment in radius, giving ten different radii in total. We tried different ring thicknesses, a disk instead of a ring, and also whether the filter boundary should be a square or a circle.

Technically speaking, each ring filter was convolved with the image, producing a new correlation image, where the value at each pixel indicated how well the ring filter, when centred on that pixel location, agreed with the image. A high positive value then indicated a bright ring, and a high negative value a dark ring.

In order to extract ring candidates, a threshold value T is used twice on the correlation image. First, bright rings are identified at regions with correlation $> T$. Next, dark rings are identified at regions with correlation $< -T$. By selecting a high T , few ring candidates will be extracted. This will reduce the number of false detections, but may also reduce the number of true detections. By selecting a low T , many ring candidates will be detected. This will increase the number of false detections, but may also increase the number of true detections.

The threshold value T is the single most sensitive parameter that the user may adjust. What value to use depends on the following factors.

- If good features can be extracted, resulting in good classification performance, a fairly low value of T may be used, since most of the false detections will be removed in the classification step.
- If, on the other hand, the classification performance is poor, then a higher value of T might be necessary.
- If the goal is to locate some rings in a large number of images, then a high T value is desired.
- If the goal is to locate as many true rings as possible in a limited number of images, and thorough manual inspection is acceptable, then a low value of T is desired.
- If the number of false detections is too high, then the usefulness of the automatic recognition is questionable, since it would be just as time-consuming to inspect a large number of ring candidates as to inspect the images manually without the aid of CultSearcher.

3.1.3.1 Ring template shapes

We have experimented with various ring template shapes. A common principle for all these templates is that they have a mean value of zero.

Envi has a function to create band pass filters for use in the frequency domain, but this can be used to create ring templates, for use in the image domain, as well. By setting the two cut-off radii to the same value, a fairly thick ring is obtained (Figure 44 a). This template has a square-shaped boundary, meaning that some pixels quite far from the centre will contribute, while some pixels at the same distance will not. To eliminate this problem, the template boundary was made circular (Figure 44 b), and the pixel values adjusted so that the average value was still zero.




Figure 44. Different template boundaries: (a) square boundary, and (b) circular boundary. In (b), the cyan pixels are outside the template boundary.




Figure 45. Different ring shapes. (a) ring type 1: normal ring, (b) ring type 2: thin ring, (c) ring type 3: thick ring, (d) ring type 4: disk. Note that pixels outside the template boundary are cyan.

Experiments on some of the test data indicated the following.

- A circular template boundary was much better than a square-shaped boundary.
- Two different circular boundary radii were tested: 1.5 times the ring radius, and 2 times the ring radius, and 2 times the ring radius was better.
- Three different ring thicknesses were tested: 2 pixels wide (Figure 45 a), 1 pixel wide (Figure 45 b), and the ring obtained by using the band pass filter function in Envi (Figure 45 c).

- A disk template (Figure 45 d) was not better than the best ring template for detecting circular patterns that looked more like a disk than a ring. However, with only three disks in the test data, this is hardly a conclusive result.

3.1.4 Feature extraction

The purpose of feature extraction is to measure some qualities of each object. The idea is that each object is described by its features. Good features are features that will make the subsequent classification of the objects easy. It is common to use the term *feature vector* to mean the collection of features that were extracted from one object.

Prior to feature extraction, each segmented ring candidate carries the following information:

- X- and y-coordinates of the ring centre in the image
- Ring radius – r
- Ring type – dark or bright.

To be able to extract feature vectors, sub images of each ring were extracted as follows. For each ring candidate, a $4r \times 4r$ sub image was extracted, centred on the ring candidate's centre, where r is the radius found in the template matching. The corresponding sub image from the local contrast enhanced image was extracted, and two thresholded versions of it were created. For bright rings, the two thresholds were 0.5 and 1.0, and each pixel with a higher value than the threshold resulted in a white pixel in the corresponding location, otherwise black. For dark rings, the thresholds were, -0.5 and -1.0, and each pixel with a lower value than the threshold resulted in a white pixel, otherwise black. In addition, all pixels outside a $2r$ radius from the sub image centre were set to black.

This resulted in four sub images for each ring candidate – two grey level sub images and two binary sub images

A number of feature extraction methods have been suggested in the literature [14][17]. We have tried out the following on the extracted sub images. Some were used only on the binary sub images

- Ring cover, that is, the amount of overlap between a binary sub image and a binary version of the ring filter. This is measured as the number of pixels in the intersection between the two binary images.
- Mean value of binary image. The mean x- and y-coordinates of the binary pixels, normalized to a value between 0.0 and 1.0. A symmetric image would give 0.5 and 0.5 for the mean x- and y-coordinates, respectively.
- Hu moment invariants – these were computed for both grey level and binary sub images
- Real weighted Fourier moments – as for Hu moment invariants, these were computed for both gray level and binary sub images

In order to identify the most promising features, scatter plots were made for all features, two features at a time. In the scatter plots, all “true” rings from the test images were included, along with “false” rings from one or two test images. Even for the most promising features, they still seemed to be drawn from the same population.

3.1.4.1 Computation of Hu moment invariants

The following description is in part taken from [17]. Given a grey scale sub-image f containing a segmented ring candidate (or other object of interest), the regular moments of order $(p + q)$ are defined as

$$m_{pq} = \sum_{i=1}^M f(x_i, y_i)(x_i)^p (y_i)^q$$

Here, the sum is taken over all the M pixels in the sub-image. The translation invariant *central* moments of order $(p + q)$ are obtained by placing the origin at the centre of gravity:

$$\mu_{pq} = \sum_{i=1}^M f(x_i, y_i)(x_i - \bar{x})^p (y_i - \bar{y})^q$$

where

$$\bar{x} = \frac{m_{10}}{m_{00}} \quad \bar{y} = \frac{m_{01}}{m_{00}}$$

Scale invariant moments [12] are obtained as

$$v_{pq} = \frac{\mu_{pq}}{\mu_{00}^{\frac{p+q}{2}}} \quad p + q \geq 2$$

From these, rotation invariant features can be constructed. There are two second order invariants ($p + q = 2$) and five third order invariants ($p + q = 3$):

$$\phi_0 = v_{20} + v_{02}$$

$$\phi_1 = (v_{20} - v_{02})^2 + 4v_{11}^2$$

$$\phi_2 = (v_{30} - 3v_{12})^2 + (3v_{21} - v_{03})^2$$

$$\phi_3 = (v_{30} - v_{12})^2 + (v_{21} - v_{03})^2$$

$$\phi_4 = (v_{30} - 3v_{12})(v_{30} + v_{12})((v_{30} + v_{12})^2 - 3(v_{21} + v_{03})^2) \\ + (3v_{21} - v_{03})(v_{21} + v_{03})(3(v_{30} + v_{12})^2 - (v_{21} + v_{03})^2)$$

$$\phi_5 = (v_{20} - v_{02})((v_{30} + v_{12})^2 - (v_{21} + v_{03})^2) + 4v_{11}(v_{30} + v_{12})(v_{21} + v_{03})$$

$$\begin{aligned} \phi_6 &= (3v_{21} - v_{03})(v_{30} + v_{12})((v_{30} + v_{12})^2 - 3(v_{21} + v_{03})^2) \\ &\quad + (v_{30} - 3v_{12})(v_{21} + v_{03})(3(v_{30} + v_{12})^2 - (v_{21} + v_{03})^2) \end{aligned}$$

3.1.4.2 Computation of real weighted Fourier moments

The real weighted Fourier moments (page 18 in [14]) are defined on the unit circle as follows

$$h_{p1}(r) = \frac{1}{\sqrt{r}} \cos 2\pi pr$$

$$h_{p2}(r) = \frac{1}{\sqrt{r}} \sin 2\pi pr$$

$$v_{pq1} = \frac{1}{2\pi} h_{p1}(r) e^{jq\theta}$$

$$v_{pq2} = \frac{1}{2\pi} h_{p2}(r) e^{jq\theta}$$

$$RWF_{pq1} = \int_0^{2\pi} \int_0^1 v_{pq1}^*(r, \theta) f(r, \theta) r dr d\theta$$

$$RWF_{pq2} = \int_0^{2\pi} \int_0^1 v_{pq2}^*(r, \theta) f(r, \theta) r dr d\theta$$

or, written out as

$$RWF_{pq1} = \int_0^{2\pi} \int_0^1 \frac{\sqrt{r}}{2\pi} \cos 2\pi pr \cdot (\cos q\theta - j \sin q\theta) \cdot f(r, \theta) dr d\theta$$

$$RWF_{pq2} = \int_0^{2\pi} \int_0^1 \frac{\sqrt{r}}{2\pi} \sin 2\pi pr \cdot (\cos q\theta - j \sin q\theta) \cdot f(r, \theta) dr d\theta$$

where p and q are integer.

For a discrete image, we need to replace the integrals by summations:

$$RWF_{pqn} = \sum_{0 < (x^2 + y^2) \leq 1} v_{pqn}^*(r, \theta) f(x, y)$$

or, written out as

$$RWF_{pq1} = \sum_{0 < (x^2 + y^2) \leq 1} \frac{1}{2\pi\sqrt{r}} \cos(2\pi pr) \cdot (\cos q\theta - j \sin q\theta) \cdot f(x, y)$$

$$RWF_{pq2} = \sum_{0 < (x^2 + y^2) \leq 1} \frac{1}{2\pi\sqrt{r}} \sin(2\pi pr) \cdot (\cos q\theta - j \sin q\theta) \cdot f(x, y)$$

Here, the image $f(x, y)$ is scaled so that $-1 \leq x \leq 1$ and $-1 \leq y \leq 1$, and r and θ are implicit functions of x and y . We have also used that $rdrd\theta = dxdy$ when changing from integrating over polar coordinates to integrating over Cartesian coordinates, prior to replacing integrals by summation. In Cartesian coordinates, the integral is undefined for $r = 0$, so we have simply omitted the image value at the centre pixel.

3.1.5 Classification based on a decision tree

Since we had a very limited test data set, containing only 35 identified rings, we could not train a statistical classifier and get any reliable estimate for a covariance matrix. Instead, a classifier based on simple if-tests was used. For each feature that was used, a lower and an upper bound for the acceptable values were set. These values were determined from the scatter plots, thereby training on the test data and, possibly, giving too optimistic estimates for the classification results.

Even these classification results were not too promising. Depending on the threshold value on the correlation in the template matching in the segmentation step, 10 to 100 times as many false positives as true positives were detected.

3.1.6 Combinations of the above methods

We wanted to investigate what combination of the above methods should be used.

First, we considered the case when no band pass filtering in the frequency domain was done. If local contrast enhancement was skipped, or, alternatively, replaced by a global contrast enhancement, then none of the true rings were detected in the subsequent template matching. Only false detections were made. Global contrast enhancement means that the mean μ and standard deviation σ for the entire image is used in the below equation

$$p_{CE}(x, y) = \frac{p(x, y) - \mu}{\sigma}$$

instead of re-computing $\mu(x, y, N)$ and $\sigma(x, y, N)$ for each pixel location. The latter is done in local contrast enhancement.

We next considered if band pass filtering in the frequency domain could replace local contrast enhancement. This was not the case. Again, none of the true rings were detected. Also, if global contrast enhancement was applied, then none of the true rings were detected. This demonstrated that if band pass filtering followed by inverse fast Fourier transform was applied, then local contrast enhancement still had to be applied.

We next investigated feature extraction and classification. The idea is that template matching could produce a lot of ring candidates, and that the false rings could be removed by a classifier, based on features with high discriminative power. Although the individual scatter plots did not identify any features that were clearly able to separate the true rings from the false rings, some of them seemed to be able to remove a few false rings. For each such feature, we could set an interval which the feature had to be within. The hope was that by including all these features and intervals in a decision tree, different rings would be excluded by different features. Experiments demonstrated that this was indeed the case, but still many false rings remained. It should be noted that the intervals were determined from the test set, thereby training on the test set, so the less than promising results are in fact too optimistic.

Based on this, we simply decided to skip the decision tree in the current version of CultSearcher, meaning that all feature vectors will be classified as rings.

3.2 Overall description of the best approach

The algorithm for ring structure detection can be summarized as follows.

1. Define masks of agricultural fields. This can be done by drawing regions of interest, or by importing a vector file.
2. Extract sub-images of size 4096×4096 pixels. This is done manually using ENVI tools. (Note that the requirement that the sub-image be 4096×4096 can be relaxed if the optional band pass filtering is skipped.)
3. Perform image processing on sub-image:
 - a. Optionally, do band pass filtering in the frequency domain, as follows.
 - i. Apply the fast Fourier transform on the image. This is a way of sorting the information content in the image based on frequencies, that is, how quickly the grey level intensities change in the original image.
 - ii. Apply a band pass filter to remove low and high frequencies that don't contribute to rings.
 - iii. Apply inverse fast Fourier transform.
 - b. Apply local contrast enhancement
4. Search for rings:

- a. Construct ring templates of increasing sizes
 - b. Convolve image with a ring template
 - c. Threshold convolution result to find bright rings
 - d. Threshold convolution result to find dark rings
 - e. Repeat b – d for all ring template sizes.
5. Validate ring detection by letting the operator delete false detections.

3.3 Detailed description of the best approach

Look at the code and describe the algorithm in detail, with formulas and everything.

The algorithm for ring structure detection can be summarized as follows.

1. Define masks of agricultural fields. This can be done by drawing regions of interest, or by importing a vector file.
2. Extract sub-images of size 4096×4096 pixels. This is done manually using ENVI tools. (Note that the requirement that the sub-image be 4096×4096 can be relaxed if the optional band pass filtering is skipped.)
3. Perform image processing on sub-image:
 - a. Optionally, do band pass filtering in the frequency domain, as follows.
 - i. Apply the fast Fourier transform on the image. This is a way of sorting the information content in the image based on frequencies, that is, how quickly the grey level intensities change in the original image. Since its size is a power of 2, this can be done lossless. Envi has built-in functions for applying the fast Fourier transform, creating a band pass filter, applying the band pass filter and doing the inverse Fourier transform.
 - ii. Apply a band pass filter to remove low and high frequencies that don't contribute to rings. The cut-off radii were 100 and 800 pixels.
 - iii. Apply inverse fast Fourier transform.
 - b. Apply local contrast enhancement. This is described in section 3.1.1.
4. Search for rings:
 - a. Construct ring templates of increasing sizes. The system suggests to use rings of radii from 4,5 to 9.0 m, with 0.5 m increment between successive ring templates, and to use "ring type 1" (Figure 45). Also, the system suggests a

value for the correlation threshold needed in step c below. The user may change these five parameters. For each ring template to create, the radius in meters is converted to pixels as a floating point value. The ring template of the selected type is drawn inside an image of size $s = 4 \times \text{round}(\text{ring_radius}) + 1$. Initially, the values are in the range 0.0 ... 1.0. The template boundary has radius $2 \times \text{round}(\text{ring_radius})$, meaning that pixels with a longer distance from the centre than the boundary radius are set to zero (cyan in Figure 45). For all other pixels, the mean filter value is subtracted, so that the final template has a mean value of zero, small negative values outside the ring (black in Figure 45), and positive values inside the ring (white and light grey in Figure 45).

- b. Convolve image with a ring template. Envi has a built-in function $B = \text{CONVOL}(A, K, S)$, which can take an $n \times m$ input image A , a $k \times k$ convolution kernel K , and a scale factor S , to compute a convolution result B

$$B(t, u) = \frac{1}{S} \sum_{i=0}^{k-1} \sum_{j=0}^{k-1} A\left(t + i - \frac{k}{2}, u + j - \frac{k}{2}\right) K(i, j)$$

for $k/2 \leq t \leq n-1-k/2$ and $k/2 \leq u \leq m-1-k/2$, otherwise, $B(t, u) = 0$.

The scale factor is

$$S = \sqrt{\frac{1}{n_R} \sum_{i, j \in R} K(i, j)^2}$$

Here, R is the region inside the circular template boundary, and n_R is the number of pixels in this region.

- c. Threshold the convolution result to find rings. In the thresholded image $B_T(x, y) = 1$ if $|B(x, y)| > T$, otherwise 0. For each region in B_T , do the following
- i. Find the local extreme, that is, the maximum value of $|B(x, y)|$, and the corresponding location (x, y) , within the region
 - ii. Check if there is an existing ring less than 5 pixels away. If there is, then merge the two rings by keeping the higher local extreme, and corresponding pixel location, of the two. Else, create a new ring with the local extreme and pixel location found in step i above, and ring type. Ring type = "bright ring" if $B(x, y) > 0$, otherwise "dark ring".

The value of T greatly affects the result of the search for rings. If a high value for T is used, very few rings will be detected, whereas if a low value for T is used, many rings will be detected.

5. Extract features – no features are extracted at the moment. However, this step is used to save the complete result of the template matching.
6. Classify segments – all input features are classified as rings at the moment. However, this step is used to produce the input file needed by the validation step.
7. Validate ring detection by letting the operator delete false detections. The operator is led through the rings one by one, and has the opportunity to delete false detections.

4 Experimental results

The algorithm, as described above in section 3.2-3.3, was applied to the entire data set, that is, a collection of sub images of sizes 4096×4096 pixels, which together covered all the identified true rings. Three parameters were varied:

- The correlation threshold
- Whether band pass filtering in the frequency domain was used or not.
- Whether a normal ring (Figure 45a) or a thin ring (Figure 45b) template was used in template matching.

The number of detected false rings varies dramatically with the correlation threshold (Table 1). If band pass filtering in the frequency domain is used (PBF=1 in Table 1), then the correlation threshold should be around 0.04 higher to obtain similar results as without band pass filtering. With this adjustment, the number of false detections is in the same order of magnitude regardless of whether band pass filtering is used or not.

If a template with a thinner ring was used, more false detections were made. Note that in this case, the ring template had the value zero outside the ring, one on the ring, and thus a positive mean value. This ring template closely resembles the initial experiments done with code written in the Matlab programming language (<http://www.mathworks.com/products/matlab/>).

When the number of detected true rings is similar, we observe that a few rings were detected with band pass filtering but not without, and also vice versa.

A reasonable compromise between not detecting too many false rings and at the same time detect as many true rings as possible, might be when the number of false detections is approximately seven times the number of true detections (yellow lines in Table 1). In this case, 11 out of 15, or 73%, of the strong rings were detected, and 5 out of 10, or 50%, of the fairly strong rings were detected. This is 16 out of 25 of the strong and fairly strong rings, or 64%.

The number of false positives can be reduced, at the cost of reducing the number of true positives as well. For example, by reducing the number of false positives from 7 times to less than half the number of true positives, the number of detected strong and fair rings decreased from 64% to 32%. On the other hand, even if the correlation threshold is set so low that almost 30 times as many false rings as true rings are detected, many of the strong and fairly strong rings are not detected. Further, none of the weak rings are detected.

Table 1. Detection results.

| BPF | filter | corr | Gardermoen | | | | | | | | | | | | Lågen | | | | | | | | | | | | sum | | | | | | | | | | | | | | | |
|-----|--------|------|------------|----|-------|-------|----|-------|----|----|----|----|---|---|-------|---|---|---|---|---|----|---|---|---|---|---|------|---|---|------|-----|-----|-----|----|---|---|--|--|-----|---|---|----|
| | | | Image | | | | | | | | | | | | sub1 | | | | | | | | | | | | sub3 | | | | | | | | | | | | all | S | F | IW |
| | | | 1a | 1b | sub1c | sub2a | 2b | sub3a | 3b | 4a | 4b | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | | | | | | | | | | | | | |
| 0 | ring | 0.3 | 1 | 0 | 0 | 1 | 0 | 0 | 1 | 0 | 1 | 0 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 | 0 | 0 | 1 | 1 | 16 | 11 | 5 | 0 | | | | | | |
| 0 | ring | 0.3 | 32 | 26 | 6 | 24 | 26 | 26 | 33 | 12 | 25 | 63 | | | | | | | | | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 | 0 | 0 | 1 | 1 | 450 | | | | | | | | | |
| 0 | ring | 0.33 | 1 | 0 | 0 | 1 | 0 | 0 | 1 | 0 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 | 0 | 0 | 1 | 1 | 16 | 11 | 5 | 0 | | | | | | |
| 0 | ring | 0.33 | 11 | 4 | 3 | 9 | 9 | 9 | 4 | 4 | 6 | 14 | | | | | | | | | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 109 | | | | | | | | | | |
| 0 | ring | 0.35 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 12 | 10 | 2 | 0 | | | | | | | |
| 0 | ring | 0.35 | 1 | 0 | 1 | 1 | 5 | 5 | 1 | 0 | 4 | 5 | | | | | | | | | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 39 | | | | | | | | | | |
| 0 | ring | 0.4 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 8 | 8 | 0 | 0 | | | | | | | | |
| 0 | ring | 0.4 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 3 | | | | | | | | | | | |
| 1 | ring | 0.35 | 0 | 0 | 0 | 1 | 0 | 0 | 1 | 0 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 15 | 12 | 3 | 0 | | | | | | | | |
| 1 | ring | 0.35 | 11 | 9 | 1 | 8 | 12 | 12 | 10 | 7 | 8 | 20 | | | | | | | | | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 | 0 | 0 | 174 | | | | | | | | | | | |
| 1 | ring | 0.38 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 13 | 11 | 2 | 0 | | | | | | | | |
| 1 | ring | 0.38 | 5 | 2 | 0 | 3 | 4 | 4 | 1 | 1 | 3 | 5 | | | | | | | | | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 48 | | | | | | | | | | | |
| 1 | ring | 0.39 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 12 | 10 | 2 | 0 | | | | | | | | | |
| 1 | ring | 0.39 | 1 | 1 | 0 | 1 | 2 | 2 | 1 | 1 | 1 | 4 | | | | | | | | | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 31 | | | | | | | | | | | | |
| 1 | ring | 0.4 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 10 | 9 | 1 | 0 | | | | | | | | | |
| 1 | ring | 0.4 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 12 | | | | | | | | | | | | |
| 0 | thin r | 0.4 | 1 | 0 | 0 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 21 | 14 | 6 | 0 | | | | | | | | | |
| 0 | thin r | 0.4 | 62 | 88 | 14 | 45 | 58 | 58 | 70 | 38 | 42 | 74 | | | | | | | | | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1314 | | | | | | | | | | | | |
| 0 | thin r | 0.45 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 14 | 13 | 1 | 0 | | | | | | | | | |
| 0 | thin r | 0.45 | 11 | 7 | 8 | 7 | 12 | 12 | 12 | 4 | 12 | 13 | | | | | | | | | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 198 | | | | | | | | | | | | |
| 0 | thin r | 0.5 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 9 | 9 | 0 | 0 | | | | | | | | | | |
| 0 | thin r | 0.5 | 1 | 0 | 1 | 0 | 0 | 3 | 0 | 0 | 0 | 2 | | | | | | | | | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 21 | | | | | | | | | | | | |

5 Discussion

One may argue that it is rather disappointing that feature extraction followed by classification was not able to separate the false rings from the true rings. On the other hand, we are trying to recognize a rather simple geometrical shape, which should be well suited for template matching. The better the template matching result, the less there is to gain on improving feature extraction and classification. So, in order to gain anything, one probably has to lower the threshold on correlation used in the template matching to give more false positives at that stage. Hopefully, a few more true positives may appear as well. Then, the feature extraction and classification steps must be so good that the final recognition results are improved.

Another issue is what information one can hope to extract from rings. Since the ring shape is used in template matching, one could argue that this shape information is already used, so features describing the ring shape itself would not represent any new information. Features that could be useful might describe acceptable deviations from the circular shape.

One approach we haven't tried is to include multi-spectral information as features, either from the four individual bands, the normalized difference vegetation index, or a pan-sharpened image.

At the moment, statistical classification is not being used, mainly because we lack a sufficiently large training set. If a large training set could be produced, then we could investigate if a statistical classifier could obtain acceptable recognition rates.

However, classification performance is highly dependent on whether features with high discriminatory power have been extracted from the data. Perhaps the multi-spectral bands or pan-sharpened images could provide suitable features.

If a large training set can not be produced, we could still attempt to use a decision tree on the new features extracted from the multi-spectral information.

At present, only one parameter is varied for the different ring templates, namely the radius. One could also vary the thickness of the ring, and see if that enables us to use a high correlation threshold, thus eliminating many of the false detections while at the same time detecting more true rings.

6 Conclusion

We have presented a method for detection of circular patterns in the panchromatic band of Quickbird images of agricultural land. These circular patterns are potentially ring graves or other circular archaeological sites. The method is based on image processing to enhance the appearance of low contrast rings, followed by template matching. The method has been tested on images from two different areas in Norway, containing a total of 35 manually identified circular patterns.

The experiments demonstrate that the proposed algorithm is able to detect many circular patterns. Still, many are also missed by the algorithm. If the goal is to detect each and every circular pattern, then the algorithm needs to be improved to be really useful.

For a thorough search in a limited area, a high number of false positives might be acceptable. On the other hand, for massive search through a large number of images, the number of false positives might be kept at a minimum, as long as *some* sites are detected. Some circular patterns may only be visible from time to time. In order to find these, one may have to process images from, say, a ten year period, and, say, 5-10 images per year. In this perspective, our approach can be used to process large volumes of satellite images that would otherwise not be inspected, thus detecting many new sites.

The multispectral images are not used in the present work. Further work could investigate whether

- the NDVI could be used to locate crop marks in areas with green vegetation, as suggested by Lasaponara and Masini [11]
- the near-infrared band could be used to locate crop marks in dry vegetation, also suggested by Lasaponara and Masini [11], or even to locate soil marks.
- the crop marks would be easier to detect in pan-sharpened images. Usually, the pan-sharpened image is created by taking the three visible bands red, green and blue, together with the panchromatic image, to produce a colour image with the same resolution as the panchromatic image. Alternatively, the near-infrared band could be used as well, giving a pan-sharpened image with false colours. Pan-sharpened Ikonos images were used by De Laet et al [9], and several algorithms to produce pan-sharpened images exist [1].

In this work, we have only used two satellite images, containing in total 35 identified rings. Many more satellite images and identified rings are needed to evaluate the current version of the system, spot weaknesses and experiment with possible enhancements.

All in all, the proposed computer software is a helpful tool for archaeologists in locating circular patterns in high resolution satellite images of agricultural land.

Acknowledgements

We thank Ragnar Bang Huseby, Hans Koren, Marit Holden and Jostein Amlien for fruitful discussions.

The work presented in this report was funded by The Norwegian Directorate for Cultural Heritage (Riksantikvaren) and The Norwegian Space Centre.

References

- [1] Alparone, L., Wald, L., Chanussot, J., Thomas, C., Gamba, P. and Bruce, L. M. Comparisons of pansharpening algorithms: Outcome of the 2006 GRS-S data-fusion contest, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 45, pp. 3012-3021, October 2007.
- [2] Amlien, J., Holden, M., Larsen, S. Ø., Trier, Ø. D. and Solberg, R. *CultSearcher – Computer-assisted detection of potential cultural heritage sites. Software Guide, version 2*, Norwegian Computing Center, Note no SAMBA/15/08, 2008.
- [3] Aurdal, L., Solberg, R. and Lous, J. *Computer assisted detection and classification of possible cultural heritage sites: User requirements and conceptual design*. Norwegian Computing Center, Note no SAMBA/11/05, 2005.
- [4] Aurdal, L., Eikvil, L. and Koren, H. *Cultural heritage search: Software prototype description*. Norwegian Computing Center, Note no SAMBA/13/06, 2006.
- [5] Aurdal, L., Eikvil, L. and Koren, H. *Cultural heritage search: Software prototype evolution and status report*. Norwegian Computing Center, Note no SAMBA/10/07, 2007.
- [6] Aurdal, L., Eikvil, L., Koren, H. and Loska, A. Semi-automatic search for cultural heritage sites in satellite images. *Proceedings of “From Space to Place”, 2nd International Conference on Remote Sensing in Archaeology*, Dec. 4-7 2006, Rome, Italy, BAR International Series 1568, pp 1-6
- [7] Campana, S., Frankovich, R., Corsi, M. and Pericci, F. Aerial survey project: seven years of flights over Tuscany. *Proceedings of “From Space to Place”, 2nd International Conference on Remote Sensing in Archaeology*, Dec. 4-7 2006, Rome, Italy, BAR International Series 1568, pp 497-503.
- [8] Campana, S., Piro, S., Felici, C. and Ghisleni, M. From space to place: the Aiali project (Tuscany-Italy). *Proceedings of “From Space to Place”, 2nd International Conference on Remote Sensing in Archaeology*, Dec. 4-7 2006, Rome, Italy, BAR International Series 1568, pp 131-136.
- [9] De Laet, V., Paulissen, E. and Waelkens, M. Methods for the extraction of archaeological features from very high-resolution Ikonos-2 remote sensing imagery, Hisar (southwest Turkey). *Journal of Archaeological Science*, vol. 34, pp. 830-841, May 2007.
- [10] Lasaponara, R. and Masini, N. Performance evaluation of data fusion algorithms for the detection of archaeological features by using satellite Quickbird data. *Proceedings of “From Space to Place”, 2nd International Conference on Remote*

- Sensing in Archaeology*, Dec. 4-7 2006, Rome, Italy, BAR International Series 1568, pp 13-21.
- [11] Lasaponara, R. and Masini, N. Detection of archaeological crop marks by using satellite QuickBird multispectral imagery. *Journal of Archaeological Science*, vol. 34, pp. 214-221, February 2007.
- [12] Maitra, S. Moment invariants. *Proceedings of the IEEE*, vol. 67, pp. 697-699, April 1979.
- [13] Musson, C., Driver, T. and Pert, T. Air photo applications in Wales, UK. exploration, landscape analysis, conservation and public presentation. *Proceedings of "From Space to Place", 2nd International Conference on Remote Sensing in Archaeology*, Dec. 4-7 2006, Rome, Italy, BAR International Series 1568, pp. 55-60.
- [14] Reiss, T. H., *Recognizing planar objects using invariant image features*. Lecture Notes in Computer Science, vol 676, Springer-Verlag, 1993.
- [15] Risbøl, O., Gjertsen, A. K. and Skare, K., 2006. Airborne laser scanning of cultural remains in forests: some preliminary results from a Norwegian project. *Proceedings of "From Space to Place", 2nd International Conference on Remote Sensing in Archaeology*, Dec. 4-7 2006, Rome, Italy, BAR International Series 1568, pp 107-112.
- [16] Sittler, B. and Schellberg, S., 2006. The Potential of LIDAR in Assessing Elements of Cultural Heritage Hidden Under Woodland Canopies. Possibilities and Limits in Detecting Microrelief Structures for Archaeological Surveys. *Proceedings of "From Space to Place", 2nd International Conference on Remote Sensing in Archaeology*, Dec. 4-7 2006, Rome, Italy, BAR International Series 1568, pp 117-122.
- [17] Trier, Ø. D., Jain, A. K. and Taxt, T. Feature extraction methods for character recognition – a survey. *Pattern Recognition*, vol. 29, no. 4, pp. 641–662, April 1996.