Kristian Ryder Thomsen and Mathias Londin Larsen

Mathematical Image Analysis on Historical Textiles

Introduction to mathematical image analysis
Mathematical image analysis is the field of extracting useful information from images. This could be as simple as reading a barcode or as complex as automatic facial recognition. This is typically done by using mathematical models based on (e.g. multivariate) statistics and evaluated by computer algorithms. Mathematical image analysis is a non-destructive, efficient and precise examination method which makes it ideal for use on historical textiles.

Fig. 1. The dress analysed in this study (Photo: Ulla Mannering, CTR).

Case Study
In a three week project course at the section for Image analysis and Computer Graphics, at DTU Informatics, Department of Informatics and Mathematical Modelling, the Technical University of Denmark, we have examined the possibilities to measure wear of textiles using mathematical image analysis. We have analysed images of a dress which we had borrowed from Danish National Research Foundation's Centre for Textile Research (CTR) at the University of Copenhagen (Fig. 1). The dress is a copy of a garment found in Moselund bog near Viborg in Denmark together with the body of a man (Østergård 2004, 135-140). The original costume is 14C-dated AD 1050-1155. The copy, which we have looked at, was made by the Danish hand weaver Anna Nørgård. It has been used at the Danish Viking Ship Museum in Roskilde by visiting children for about three years. When due to wear and tear it could no longer be used at the museum, it was given to CTR. During its use it has not been washed or cleaned otherwise. This means that the garment has only been subjected to mechanical wear similar to historical use.
The copy is made of an industrially produced 2/2 twill wool fabric. It has clear signs of wear, for instance it has multiple holes, threadbare areas and defects like detached threads in the otherwise homogeneous machine spun yarn.
To analyse the dress we have worked with multispectral image data of various places of the dress. Multispectral imaging is taking the "same image" multiple times with different wavelengths across electromagnetic spectrum. In our case the image has been taken 18-20 times in the spectrum across ultra violet, visible light and near infrared. Multispectral image data have been used because different materials have different reflective properties depending on the wavelength.

This is true for even very small local fluctuations in the material. Images taken with near infrared light have also been used because the colour pigments in the threads are invisible in this light, thereby exposing the pure thread and making it possible to measure thread thickness more accurately than with regular images.

Measuring of visible background by means of canonical discriminant analysis

A logical place to start when you want to characterize wear and tear in textiles only by means of image analysis is to measure and quantify how much background is visible through the fabric. This gives you basic information about the state of the textile with regard to holes and whether or not the textile has been worn thin. In an image all pixels have a value that represents its colour. The basic idea is to count how many pixels have "thread colour" and to get a value of the state of wear and tear in the textile by comparing this to the total number of pixels in the image. There is a complication however. A thread is not uniform single colour but a gradient of similar looking colours partly due to natural variation and partly due to physical aspects of the photography. In order to convert our image into a binary image containing only two colours we threshold the image after using a mathematical orthogonal transformation called normalized Canonical Discriminant Analysis, referred to as nCDA. It is beyond the scope of this article to explain nCDA in detail, however the basic principle is explained briefly (Aasbjerg Nielsen 1999, 8).

We have been working with multispectral image data meaning that the "same image" has been taken 18-20 times with different wavelengths across visible and infrared light. This gives us 19-21 sets of dimensional data. By defining what a pixel in a thread respective to the background looks like in this 19-21 dimensional space, nCDA can be used to make a linear combination of the 18-20 images that separates thread and background, so all pixels of thread get centred around the value 1 and all pixels of background get centred around the value -1. By thresholding this image at zero one gets a binary black and white image in which one can just count the number of black and white pixels respectively. Figures 2-4 illustrate an example of nCDA used in a rather extreme case of the wear and tear in the textile and demonstrate nCDA's capacity to distinguish between thread and background.

Automatic measuring of thread thickness by texture correlation

Visual inspection shows that the thread thickness varies a little throughout the textile. This variation can be automatically measured from a single photograph using a mathematical method based on texture correlation as a function of offset in a Grey Level Co-occurrence Matrix, referred to as a GLCM (Carstensen 2002, 224). The textile consists of parallel threads, with a specific spacing, thereby making the textile periodic. The basic idea is to measure this periodicity. From this information one can with a high degree of precision compute backwards the average thread thickness in a local part of the image. This information can also be used to automatically evaluate yarn evenness, evaluate yarn fineness, calculate the textile's cover factor and thickness group, as well as warp and weft set.

In order to measure the periodicity in a textile's weave one first needs to calculate a GLCM of the image of the textile. A GLCM is a matrix (two dimensional table of numbers, like a Sudoku) giving a full representation of second order gray level statistics. Loosely speaking a GLCM is just a count of how many times a given pixel value co-occurs in a given distance in the image. In more mathematical terms, a GLCM is defined with respect to a given displacement h, and element (i,j), donated cij, as the number of times a point having gray level j occurs in position h relative to a point having gray level i. The meaning of this definition is more apparent if, as a concrete example, we compute the GLCM of a 4 colour image, with h=(0,1), i.e. one step in the horizontal direction (Figure 5).

By computing the correlation in the GLCM one gets a single number. This number tells you to what extent the original image and the image with an offset h look like each other. By computing many GLCMs with increasing displacement and their corresponding correlation one can get a plot like the one seen in Figure 6. By computing either the average distance between the local maxima marked by green dots or the distance between the first two local maxima one knows the period of periodicity in the textile weave.

In Figure 7, a more intuitive illustration is made of the period in the textile's weave.

It also shows that the relationship between the individual threads is needed to compute the thread thickness.

In our project we did not find an automatic mathematical method for extracting these relationships directly from an image. However we believe it is possible with a little extra work.

By assuming the relation

$$r = 4b \qquad y = \frac{b}{4}$$

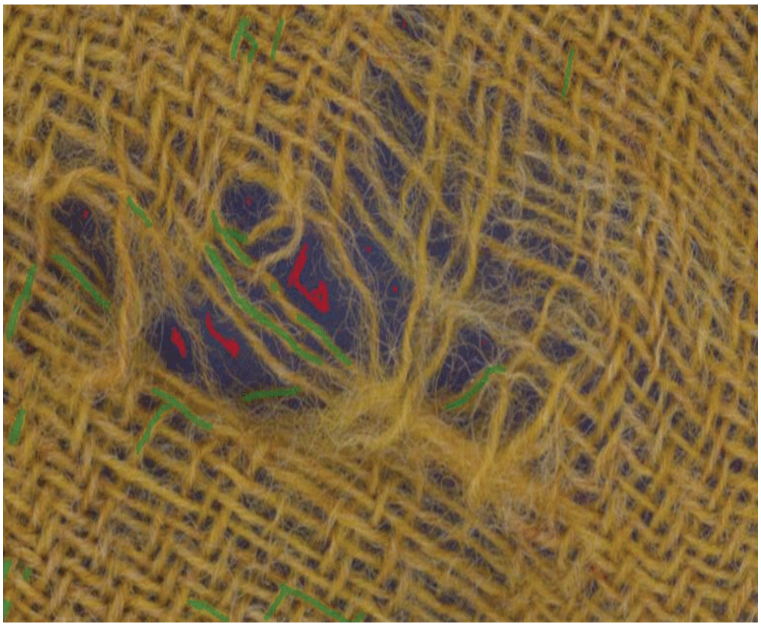

Fig. 2. A multispectral image viewed as one colour image. The red and green markings are manually painted onto the picture to define an example for the computer on what pixels constitute thread and what pixels are background (Photo: authors).

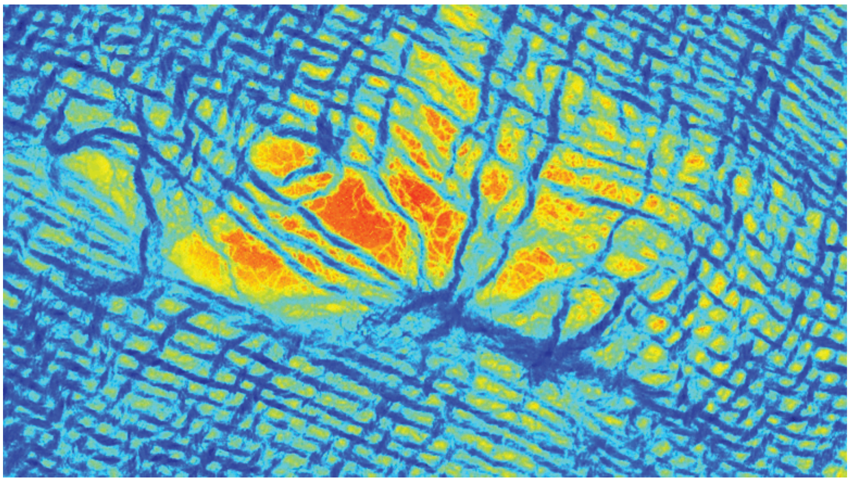

Fig. 3. Output of the nCDA analysis performed on the image seen in Fig. 2. All pixels belonging to a thread are blue and all pixels belonging to the background are orange. Pixels that are not clearly defined to contain either thread only or background only are coloured in a mixture of blue and orange (Photo: authors).

Fig. 4. Fig. 3 after a threshold. All pixels having a value below zero are coloured black and all pixels having a value above zero are coloured white. By simply counting the white pixels and comparing this number with the total number of pixels in the image one can get a percentage of how much background is visible through the fabric (Photo: authors).

One can express the thread thickness as follows:

$$r + 2b + y = T \cdot p_b$$

$$\Leftrightarrow\ 4b + 2b + \frac{b}{4} = T \cdot p_b$$

$$\Leftrightarrow\ b = \frac{4}{25} T \cdot p_b$$

Where T is either the average distance between all the green peaks or the distance between the first two green peaks in Figure 6 and Pb is the physical size of a pixel in millimetres. This information depends on the camera used to take the image and can be looked up in the specifications for camera.

Our machine-woven textile has a calculated thread thickness of b = 0.23 mm if T = 20 and Pb = 0.0725 mm. This corresponds with visual inspection.

Defect detection

By using mathematical image analysis we found three ways of performing defect detection. In the following we outline independent methods and results.

Defect detection by means of canonical discriminant analysis

The defects we are looking for are detached threads and small knots resulting from these (Fig. 8). These detached thread areas are much fluffier than the rest of the textile and absorb and reflect other wavelengths of light. Therefore, one may use multispectral image data, meaning taking the "same image" 18-20 times with different wavelengths across visible and infrared light. This gives us 19-21 sets of dimensional data. In this multidimensional data, defects and regular threads have very different characteristics, and defects can be easily detected by using a mathematical orthogonal transformation called normalized Canonical Discriminant Analysis, nCDA.

Since the fine structure of the textile is of no interest, we have first smoothed the image using a 21x21 mean filter, hereby replacing each pixel with the average of all 440 neighbouring pixels. This has the same effect on the image as a wet sponge would have on water-painted child's drawing. This leaves only the coarse details in the image, making it easier to visually define exactly which parts of the image are to be considered as defects and which are not. We then literally paint on the image to make a selection of what is to be considered a defect and what is to be considered a regular weave section (Fig. 9).

From this information the computer performs nCDA in the 19-21 dimensional image space and returns an image (linear combination of the 18-20 images) in which defects are centred around the value of -1 and coloured red and regular threads are centred around -1 and coloured blue (Fig. 10). A threshold value of > 0.36 has been experimentally determined as an optimum for separation between defects and regular threads in the nCDA output. By thresholding you get an image containing only two colours: one colour representing regular thread and one colour representing a defect as a direct result of a detached thread or simply the detached thread itself (Fig. 11). A defect can then be automatically detected by simply computing the average pixel value for a given area of the image. If this local average exceeds a user-specified limit, the area is known to contain a defect. The size of each area depends on the size of defect you like to detect and the limit depends on the weave. In our case, a limit of 0.7 gave reliable and consistent results.

Defect detection by means of the Laplacian pyramids

Defects have locally similar colour and smoothed texture, compared locally to textile without defects. This is illustrated in Figure 12. Try for instance to compare Region Of Interest (ROI) 4 and 5. You will see that in ROI 4 the textile weave is clearly visible, whereas in ROI 5 it is not. This property can be used mathematically to locate and detect textile defects. This is done with the use of what is known as a Laplacian pyramid (Carstensen 2002, 22).

In short terms, which will make much more sense in the following, if you extract a small part of the image roughly in the same size as the defects you want to

2	1	1	3	3
3	1	0	2	2
3	3	1	1	2
0	2	1	0	1
0	0	1	0	0

	0	1	2	3
0	2	2	2	0
1	3	2	1	1
2	0	2	1	0
3	0	2	0	2

Fig. 5. Left: A 4 colour image. Right: The GLCM of the left image computed with a h=(0,1), i.e. one step in the horizontal direction. This means that e.g. there are 3 pixels having the value one that has a neighbour value of zero if you take one step to the right in the image (Photo: authors).

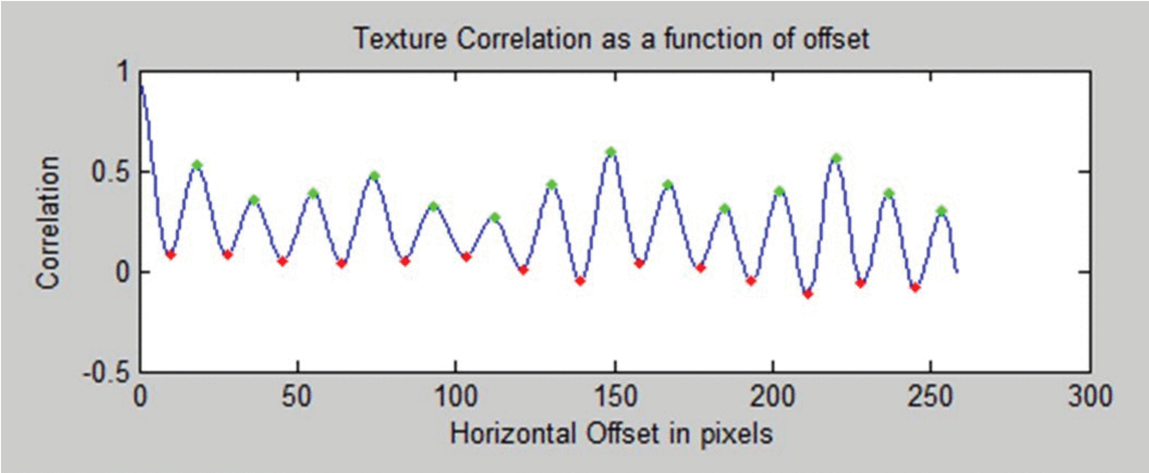

Fig. 6. Texture correlation as a function of offset. As the offset increases the correlation fluctuates periodically due to the natural periodicity in the textile weave. The period can be used to calculate the distance between two threads in the textile and the thread thickness (Photo: authors).

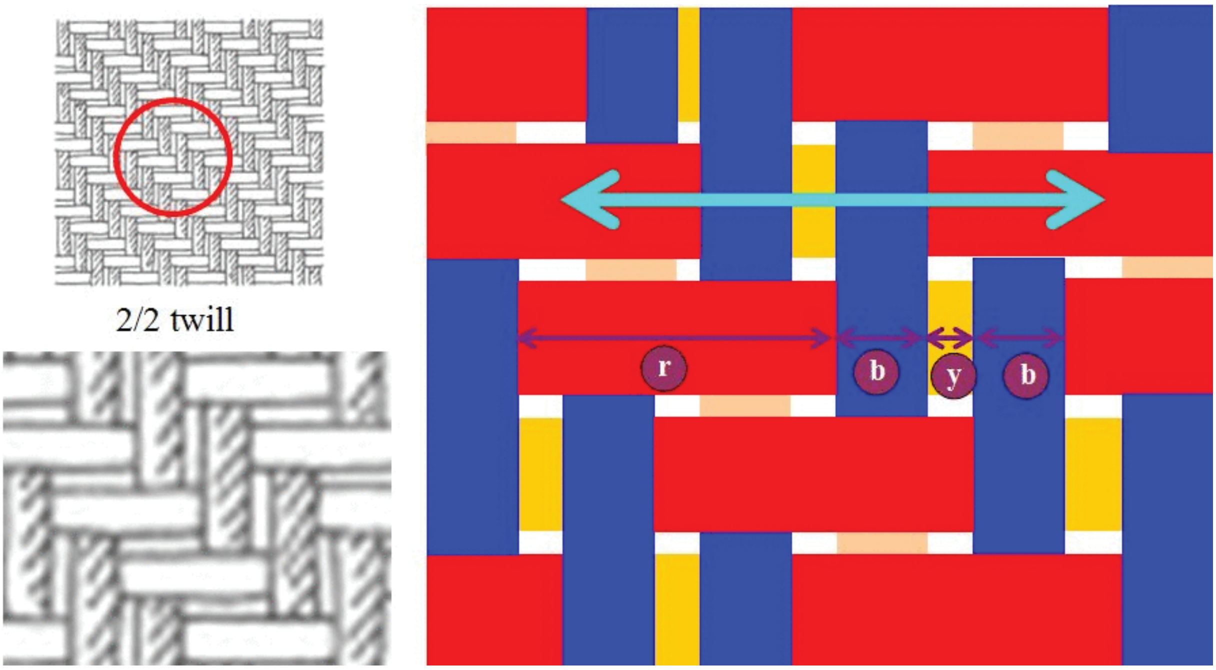

Fig. 7. A 2/2 twill weave. The period computed from Fig. 6 corresponds to the light blue arrow. Depending on the textile's orientation in the image the light blue arrow covers one piece of warp/weft thread, two weft/warp crossings and a little spacing between the weft/warp crossings. These sections are colour coded in red (r), blue (b) and yellow (y) respectively. The relationship between these variables depends on the weave and thread used. By defining the relationships between r, b and y, one can compute the average thread thickness of the entire image (Photo: authors).

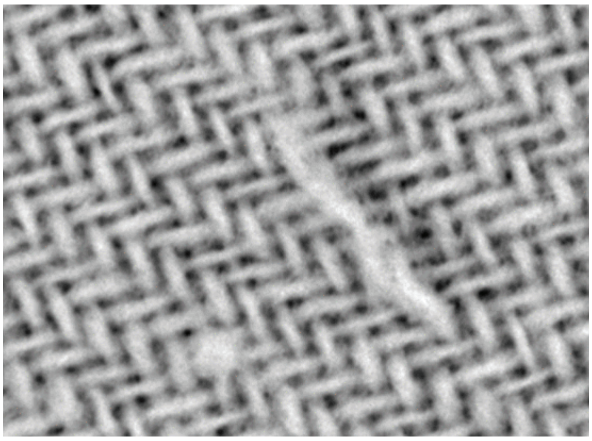

Fig. 8. Textile with a defect due to a detached thread and detached thread itself (Photo: authors).

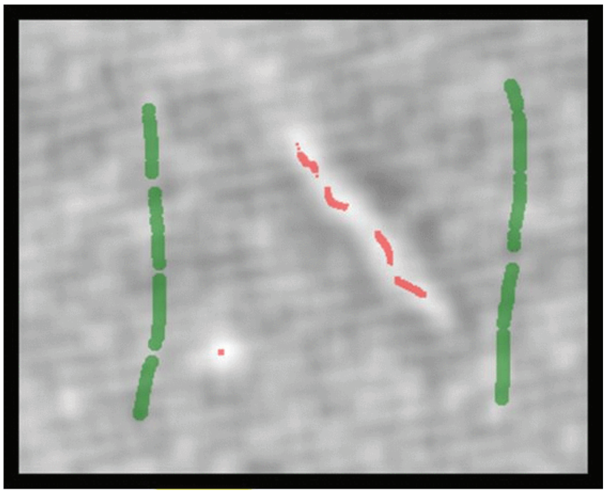

Fig. 9. Fig. 8 in a smoothed out version. The red and green markings are manually painted onto the picture to define an example for the computer on what pixels are regular thread and what pixels are defect (Photo: authors).

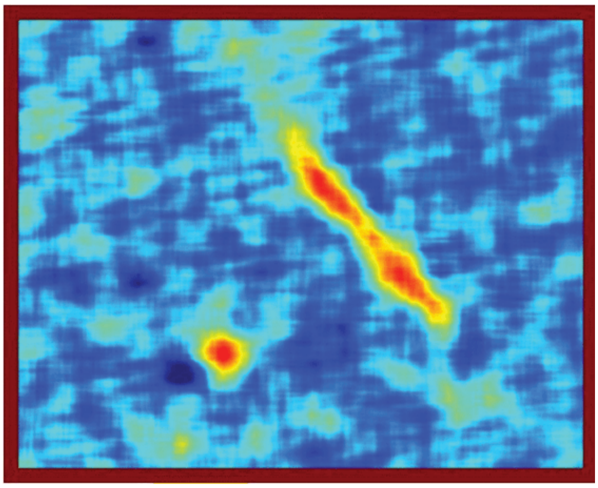

Fig. 10. Output of the nCDA analysis performed on the image seen in Fig. 9. All pixels belonging to a regular thread are blue and all pixels belonging to a defect are red. Pixels that are not clearly defined to contain either regular thread only or defect only are coloured in a mixture of blue and red (Photo: authors).

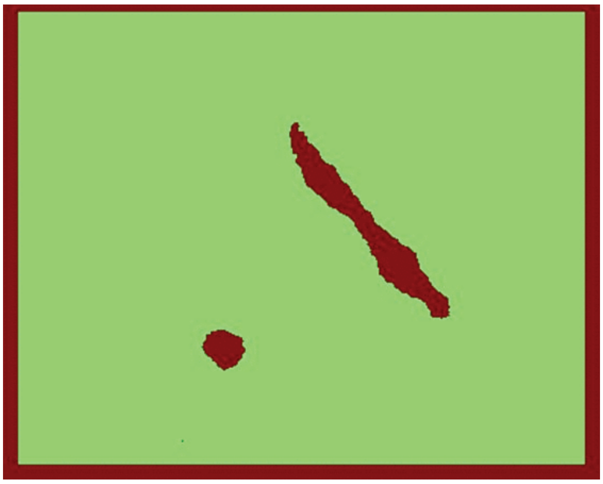

Fig. 11. Fig. 10 after a threshold. All pixels having a value below the threshold are coloured green and all pixels having a value above the threshold are coloured brown. By simply counting the brown pixels and comparing this number with the total number of pixels in the image one can get an indication of whether or not this image contains a defect (Photo: authors).

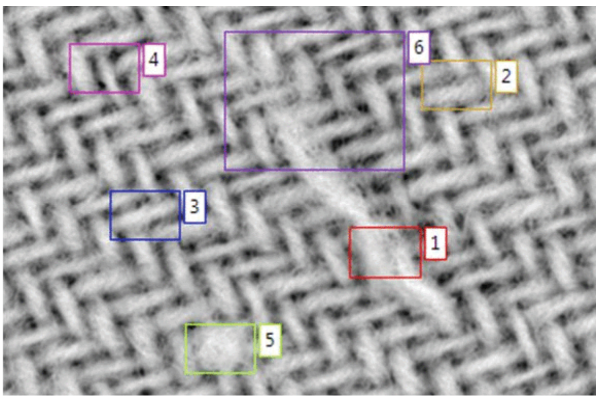

Fig. 12. Image taken at 780nm. ROI 1 and 5 contains defects. ROI 6 contains a partial defect (Photo: authors).

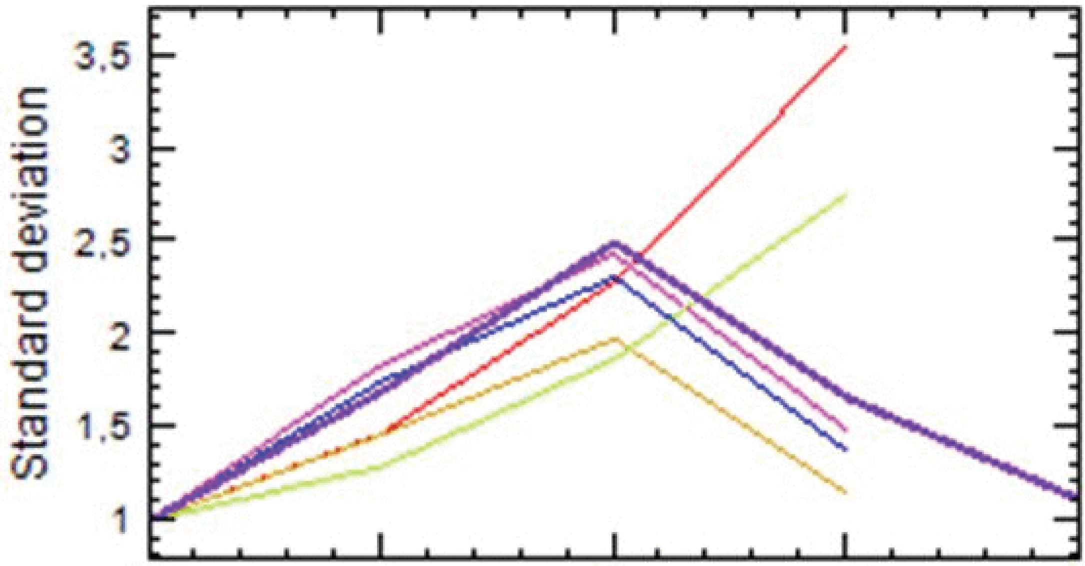

Fig. 13. The standard deviation as a function of scale (level) in a Laplacian pyramid. The curves corresponding to ROI 1 and 5 are peaking, indicating a defect. The curve corresponding to ROI 6 is not peaking, because ROI 6 is significantly larger than the partial defect it contains (Photo: authors).

detect, and plot the standard deviation as a function of scale in the Laplacian pyramid of the image part, you will see a peak in the plot, if and only if, your extracted image part contains a defect.

The images we work with one can mathematically think of as an array of pixels or pixel intensities. The pixel intensity is a number for instance between 1 and 100. This value describes the amount of light reflected at this given point in the image. The amount of light reflected depends on a number of physical properties of the material photographed. The pixel intensities for areas of thread are of course different from the pixel intensities for the space between two threads. This means that if you calculate the standard deviation of the pixel intensities for a region of the image not containing any defects, you get a large number. A region with a defect will not contain as much space between threads, therefore the standard deviation of the pixel intensities in this region will be a smaller number (in practice you need to pre-process your image first, in order to obtain this effect in a statistically stable way).

In order to pre-process your image you first make a Laplacian pyramid in the following way: Take your original image and delete every other pixel both vertically and horizontally. You now have an image a quarter of the original size. Continue doing so a few times, each time making a new level in what is known as a simplified Gaussian pyramid. Mathematically reconstruct each level in the Gaussian pyramid, as well as possible, with the use of interpolation. Pixel-wise subtract each level in the simplified Gaussian pyramid from the interpolated one, and you now have a Laplacian pyramid.

Choosing the right size of the part to extract and compute standard deviation on is very important. If your extracted part is too small the defect detection algorithm will take too much computer time and thereby not be efficient. If the extracted part is too large, the actual defect will not cover a large enough area of the extracted part, and the defect will therefore not be detected. It is hence important to choose the size of part to extract corresponding to the defect you wish to detect.

The algorithm to follow using this version of defect detection is simply to extract a part of the original image and compute the standard deviation for this part in each of the first levels in a Laplacian pyramid. If the standard deviation increases as a function of the level in the Laplacian pyramid, this part of the original image contains a defect no smaller than roughly the size of the extracted image part.

Start by extracting a part in e.g. the top left corner, calculate if it contains a defect, and then simply move

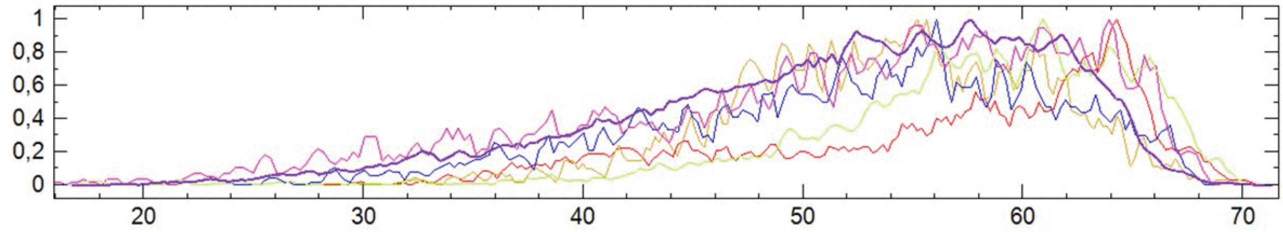

Fig. 14. Histograms for ROI's seen in figure 12 (Photo: authors).

the part to extract a little, calculate if this part of the original image contains a defect, and continue this way until you have checked the entire image. Results of the algorithm can be seen in Figure 13. Depending on the image size and available computer power checking an entire image should take in the neighbourhood of a second or two.

Defect detection by means of histogram analysis

Defects can also be detected with histogram analysis exploiting the same defect characteristic as in the previous section of this article. The fact that defects have different colour characteristics can be used to detect defects with the use of simple statistical tools.

Figure 14 shows the histograms of ROI 1-6. It is clearly seen that these histograms are not identical. Therefore, one just needs a way to differentiate between the ones that represent a defect and the ones that do not. In order to do this, it is necessary to consider the average height of the histograms and the statistical terms skewness and kurtosis. Skewness is a measure of asymmetry of the probability distribution of a real-valued random variable. By considering the histograms as probability distributions, skewness can be used to measure the distribution of dark and light colours in each ROI. A large negative skewness means that the histogram is "tilted" to the right, and thereby indicates that the histogram's left tail is longer than the histogram's right tail. This also implies that most of the pixel values in the ROI are smaller than the mean and possibly the median.

Kurtosis is a measure of how outlier-prone a distribution is. This tells you something about how many values lie far away from the histogram's "centre of mass". If the histogram has a (Gaussian) bell-shaped form it would have a kurtosis of zero. A negative kurtosis indicates a (leptokurtic) shape more like a high thin sandcastle. A positive kurtosis indicates a (platykurtic) shape more like if the top part of the sandcastle has collapsed and now lies around the base of the castle.

By computing skewness, kurtosis and average height of the histograms for ROI 1-5, a clear pattern emerges. ROI 1 and 5, which contains defects, have significantly higher values than the ROI's not containing any defect. From a defect detection point of view, a ROI is interesting if one more of the following conditions are fulfilled:

• The skewness of the ROI's histogram is larger than average.
• The ROI's histogram has a smaller average height than the average of all histograms for all ROIs.
• The kurtosis of the ROI's histogram is deviating significantly from the mean of all ROIs' kurtosis.

A ROI is defined to be containing a defect if two or more conditions are fulfilled.

By means of this method a defect is detected in ROI 1 and 5 seen in Figure 12.

Conclusion

Many archaeological and historic textiles have been used in contemporary reconstructions. The analysis of these fabrics shows a high degree of wear, but when doing reconstructions of these fabrics we use the data directly because we have no formulas for systematic analyses of wear that can characterize the change that happens with a piece of textile and its threads during use. It is therefore important to systematically analyze the wear, since it can refine the analysis of the archaeological or historic textiles, and the description of their wear. In our work we have found several ways to do systematic analysis of wear and tear in textiles.

Using nCDA analysis we obtained a concrete value of the state of wear and tear in the textile. This provides basic information about the state of the textile with regard to holes and whether or not the textile has been worn thin, which may give information about the social status of the owner. If the fabric is worn thin you can imagine that it is from a lower social status because one could not afford to buy a new garment when the old clothes were worn thin. The method also enables us to compare wear patterns of two textiles in relative terms.

Using texture analysis we have automated the measure of mean thread thickness. This has been done by measuring texture correlation in the horizontal direction. From this we have plotted the texture correlation as a function of offset. This illustrates the periodicity in the textile's weave. Based on the basic assumption about the aspect ratio between the warp and weft threads, it is possible to compute the average thread thickness of the entire image. This can be applied to measure both warp and weft thread thickness by simply rotating the image 90 degrees. As of now, an educated guess is needed to know the aspect ratio but it is possible to develop an automatic method to extract the information directly from the image.

In order to perform an automatic defect detection we have used three different methods. It is important to know that these methods can be used both to detect defects and to count the number of defects in a textile. Firstly, we used canonical discriminant analysis to automatically detect defects. This gives us an image containing only two colours. One colour represents the regular fabric and the other colour represents a defect as a direct result of a detached thread or the detached thread itself. A defect can then be automatically detected by simply computing the average pixel value for a given area of the image. If this local average exceeds a user-specified limit, the area is known to contain a defect. The size of each area depends on the size of defect to be detected and the limit depends on the weave. In our case, the limit of 0.7 gave reliable and stable results.

Secondly, we have worked with defect detection using a Laplacian pyramid. Laplacian pyramids have been used because defects have locally similar colour and smoothed texture, compared locally to area without defects. We have used this property to mathematically locate and detect textile defects.

Thirdly, we have detected defects with histogram analysis exploiting the same characteristics as in the second method. We have illustrated the histograms of six different regions of an image with defects. It could clearly be seen that the histograms were not identical. Therefore one just needed a way to differentiate between the ones representing a defect and the ones that did not. In order to do this, we considered the average height of the histograms and the statistical tools skewness and kurtosis. Based on these statistical characteristics, a pattern emerged and it was possible to clearly detect defects.

All mathematical and image analysis methods used in this project can be adapted to other textiles and weaves. We have developed these methods for archaeological textiles, but they might also have an industrial application. Automatic counting of defects or thread spacing can for instance have applications in automatic quality control.

Acknowledgements

We thank our supervisor Jens Michael Carstensen, Department of Informatics and Mathematical Modelling, Technical University of Denmark and Ulla Mannering, Centre for Textile Research for their advices.

Bibliography

Østergård, E. (2004) *Woven into the Earth. Textiles from Norse Greenland*. Aarhus.

Aasbjerg Nielsen, A. (1999) *Orthogonal Transformations*. Department of Informatics and Mathematical Modelling, Technical University of Denmark.

Carstensen, J. M. (2002) I*mage Analysis, vision and computer graphics*. Department of Informatics and Mathematical Modelling, Technical University of Denmark.

Corresponding Authors:
Kristian Ryder Thomsen: s103253@student.dtu.dk
Mathias Londin Larsen: s103456@student.dtu.dk