



*Assessing paper origin and quality
through large-scale laid lines
density measurements*

Vlad ATANASIU, Ph.D. candidate
École pratique des Hautes Études, IV^e Section
45, rue des Écoles, 75005 Paris, France
atanasiu@excite.com
<http://mywebpage.netscape.com/atanasiuvlad/ad751/>

Abstract

The analysis of the density of papermaking mould screen imprints is instrumental in dating, localizing and classifying papers. By dividing digital reproductions of paper structures into spots, the AD751 software offers a precise measurement tool with rich statistical possibilities, both quantitative and qualitative. AD751 is based on the Fourier transform. The technique of locating the frequency of the laid lines is discussed in this article, with emphasis on the benefits and problems of the 'emboss' high-pass filter.

Keywords

Papermaking, laid lines, image analysis, Fourier transform, 'emboss' filter.

Acknowledgements


AD751 was written by Vlad Atanasiu for the Instituto centrale per la patologia del libro, Rome, which funded the research. The author would like to thank Ezio Ornato (Centre national de la recherche scientifique, Villejuif) for the very late afternoons and Jean-Marie Nicolas (École normale supérieure de télécommunications, Paris) for indulging in my appearances.

Für mein Vater, meine unerhoffte Wiederkehr an die Ingenieuren.

Contents

Background	3
Software characteristics	4
Measurement method	5
Noise filtering	6
The 'emboss' high-pass filter	6
Issues of the 'emboss' filter	6
Period computing	7
AD751 measurement modes	7
Trapping errors	8
Machine vs. human performance	9
Preliminary information	9
Benchmark data source	9
Sources of error	9
Image quality expectancies	9
Measurement precision of AD751 – the results	10
The three methods' precision & speed	11
Image slant influence	11
Influence over the deslanting	11
Influence over the density analysis	12
Methodological considerations on laid lines density measurement	13
Number of measurements per paper	13
Ways to present data	13
Historical information from analysis data	14
Paper technology and mean laid lines density	14
Paper stocks and choice consistency	15
The shark of Leuven: tricks in quality discovered 500 years later	15
References	17

Background



Laid lines spaced by ~1.3 mm

When producing handmade paper, vegetal pulp covers a screen of wires, threads or stems through which the water can drip. These dense, parallel elements, called 'laid lines', leave a relief imprint on the sheet of paper and can be seen when holding it against light.

Developments in the techniques of pulp and mould production, regional particularities, as well as the natural wear of the mould itself, are factors influencing the density of laid lines, offering the historian valuable clues to the origin of a particular paper.

The study of laid lines is still incipient, having for a long time been hampered by a lack of sufficient reproductions and rapid means of measurement. Manual counting is time consuming and prone to error, increasing with the volume of papers to be analysed. Quantities in the thousands of reproductions would probably never have been studied without automation. Reproducing the human analysis process on a computer not only drastically speeds up the research but also allows additional measurements. To take an example, if localising papers from their structure is the traditional motive behind the interest in laid lines density, multiple measurements on different spots of the same paper can tell us about the quality of the mould at the moment when the paper was produced.

Speed, accuracy and a wealth of both quantitative as well as qualitative data about papers is what the AD751 computer program makes possible by measuring the density of laid lines on multiple spots. This paper will present the algorithm, what was learned about the methodology of paper structure analysis and potential ways of interpreting the resulting data. (AD751 stands for the year of the battle of Talas Valley, 400 km NE of Samarqand, Central Asia, between Muslim and Chinese armies, when the legend has that the knowledge of papermaking started its westward journey.)

Software characteristics

AD751 is Open Source software, freely downloadable at the internet address <http://mywebpage.netscape.com/atanasiuvlad/ad751/>.

AD751 was written in Matlab 6.12, a standard of scientific computational research. It is available as a Matlab toolbox and in a standalone Windows 98/2000 version. There are 2300 lines of code with four additional 400-line-long independent m-functions.

The analysis process is displayed on the screen and the resulting data saved: individual spot density, global mean, range and standard deviation. The images can be handled one by one or in batches.

The interface is multilingual (English, French, German, Italian). For the different settings see AD751's help file.

Measurement method

Reading line quantities directly from the Fourier transform

A reproduction of the paper structure such as can be obtained by radiography, transparency or rubbing, is a series of alternating shades of grey, varying from white to black in a sinusoidal way. Thus the actual variation of the thickness of the paper is translated into a visual 2D image.

The bidimensional discrete fast Fourier transform (Fft) of an image gives the number of evenly spaced points of same grey intensity. When the image is made of stripes, their quantity can be directly read from the transform as the distance from the maximum value of the transform to its constant component. This is an easy and elegant solution to our problem. In practice however several interesting issues arise.




Fig.1 – Region of interest to be analysed – 45 laid lines were found (WM-I-818)

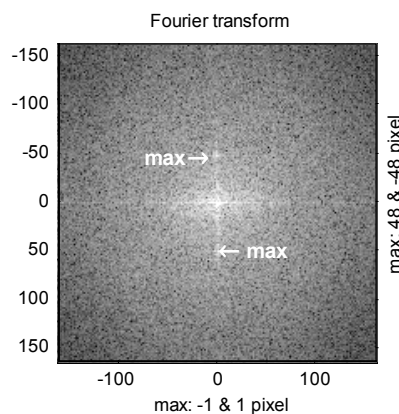


Fig.2 – The maximum correspond to the laid lines frequency. The laid lines have a -1.2° slant.

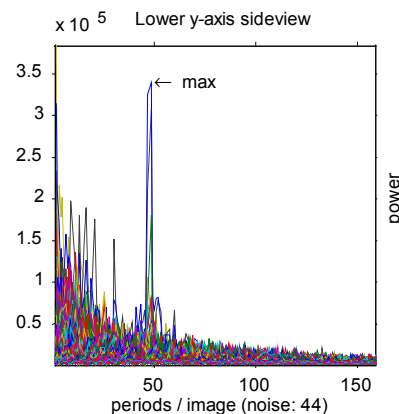


Fig.3 – The lower half of the Y-axis showing the high power of the laid lines frequency

[*Nota bene*: The Fourier transform is a mathematical tool with a very broad spectrum of applications, given its ability to extract basic information from virtually any object or phenomenon. In domains similar to ours that analyse ‘stripes’, the Fft is used at the core of warning systems against giant ocean waves by computing their height and direction from on-board radars or from satellite images; in oceanography it helps track the movement of water masses or in environmental assessment studies that of sand dunes; wood slices for dendrochronology are also successfully examined by the Fourier approach. Even science fiction writers have found it resourceful, imagining a world that would look like a multidimensional Fourier transform, with creatures made out of frequencies, a hardly comprehensible sight unless you wear special glasses that inverse transform it.]

Noise filtering

Noisy images can obscure the laid lines maximum by generating more powerful frequencies. The most commonly encountered on reproductions of paper structure are low frequency noises, generated by poor scanning quality, high image compression or – mostly – by the normally occurring unevenness of the pulp fibres distribution in the body of the paper sheet.

The 'emboss' high-pass filter

A specially designed 'emboss' edge-enhancing filter was used to drastically reduce the low frequencies and clearly 'pop-up' the magnitude of the laid lines frequency on the Fft. 'Emboss' accentuates regions of fast intensity change and 'flattens' even regions. This filter proved to be far more efficient for our image type than the usual Canny, Sobel or Prewitt filters. The algorithm consisted in making a negative of the image, shifting it one pixel down, then subtracting the values of the two overlapping matrixes.



Fig.4 – 'Emboss' filtered image (remark the absence of the vertical line).




Fig.5 – Intensity of a column of pixels before filtering...




Fig.6 – ... and after filtering.

Even highly compressed Jpeg images – with their very particular pixelisation noise – were still analysable (WM-I-688 & 941). Contrary to the norm where 'emboss' was applied twice on the input image, with this particular type of noise the filter had to be used only once.

Another type of noise appears with images digitised from printed material. The rasterising technique used in typography to render grey scales and colour produce very high frequencies, near the total amount of pixels for each dimension of the image. A cut-off low-pass filter with a threshold set at 2/3 of the highest possible frequency was used for 'unrastering' if the user checks this option.

Issues of the 'emboss' filter

Relying on 'emboss' however has a drawback. Its force comes from being strictly unidirectional: parallel lines in a given direction are enhanced, but all others are gradually dimmed until those orthogonal are completely suppressed. In other

words 'emboss' is also a phase filter. If edge detection is intended, one has to know beforehand the angle of the edge one seeks to find.

This was the case of our images, where the laid lines were horizontal within a $\pm 3^\circ$ variation. Remarkably, images purposely rotated at angles up to 45° were in many cases – depending on the structure and amount of their individual noise – correctly analysed with the 'emboss' plus Fft combination.

A more subtle problem with the 'emboss' is that it shifts the maximal frequency of the input image with some periods. Thus AD751 uses it only to find the approximate window where the laid lines frequency is to be found. Then the real frequency is computed on the input image with all the frequencies below the window considered as non-relevant and cut off. (The threshold is at 4 periods lower than the maximum found using the 'emboss' filter in the normal case of double embossing and at 15 periods for simple emboss.)

Period computing

Given that the intended result of our computations is the density of laid lines, we need to find the period of the line frequency. This is done by trimming the top, then the bottom of the image, row of pixel after row of pixel, performing the Fft at each step to find out the new number of periods, until we obtain a non-fractional number of periods per image. Dividing the periods by the image height (the lines are horizontal) gives the laid lines density.

In the last twenty years the measurement of laid lines density as the length (in mm) covered by $20\frac{1}{2}$ consecutive lines emerged as a standard practice. More lines would have been difficult to count manually without mistakes (especially on low quality reproductions) and fewer lines yield too much variation across the paper. The additional half line helps avoid the problem of finding the exact position just before the 21st period peak (20 periods span over 21 periods minus one unit), a maximum grey tonality probably being easier to detect by a human eye than the value just below it.

It should be stated that another method to measure the laid lines density by counting how many of them are found on a given length – usually 3 cm – is a loose measurement, with an imprecision equal to the interlineal distance.

AD751 measurement modes

AD751 offers three different ways for measuring densities. The most precise is the hi-fi method, in which the image is screened at a given step for spots encompassing exactly 20 evenly spaced lines, each time the density of the particular spot being computed. The historian is thus provided with a detailed description of the paper structure.

As a much faster alternative, but slightly less accurate, the multi-spot measure can be used. It works by sliding a fixed height spot across the image to calculate the density, independently of the number of lines per spot, inferring what it

should be for 20 ½ lines. The spot height can be manually set as well as automatically determined by dividing the total number of lines per image by its height (adjusted to make it run from peak to peak minus one pixel).

The global method is made available for the sake of its rapidity and its ability to perform measurements of particular regions without piecing them to spots.

The gathered data displayed by AD751 on the computing windows and saved in files are the spot densities, the mean density, the range of variation for each image, the standard deviation from the mean, the number of lines per spot and the faulty spot densities.

Trapping errors

In the sample lot of images used to develop AD751 spots with a variation greater than 5 lines around the measurement unity of 20 lines were always due to blurred regions in the reproduction and not to defects of the mould screen itself. AD751 considers those spots as faulty and does not use them in its statistics. Visual inspection of images declared 'faulty' is however recommended.

If in hi-fi mode a spot is reduced or extended to abnormal dimensions, the aforementioned error trap cannot operate, since the spot always contains 20 lines. By automatically comparing the hi-fi spots mean with that obtained by a multi spot measure, the image is nevertheless detected to contain faulty spots. The value at the extremity of the range is then eliminated, depending on the direction of the skewness (WM-I-80 & 542). Blurs in the image make that sometimes it is not possible to find 20 consecutive laid lines, in which case the multi-spot method is automatically applied (18% of the initial lot of 312 images were rejected).

Machine vs. human performance

Preliminary information

Benchmark data source

The assessment of AD751's performance was based on 312 electronradiography images, digitalised in files of 150dpi medium compressed Jpeg of an average size of 400×500 pixels, downloaded from the National Library of the Netherlands Watermark web database (<http://www.kb.nl/kb/bc/incun/watermerken-en.html>). This particular series – letter P WM-I-1 to 1254 with discontinuities – representing papers printed in the Netherlands in the last quarter of the XVth century, was chosen for benchmarking AD751 because it was an extensive collection of laid lines densities already manually analysed and typed in a computer file, providing the target sample.

Sources of error

In evaluating AD751's precision I found that relying on human measurements as target value is problematic, the difference between their respective returned values being a combination of human and machine error. Thus the final performance estimator will have an additional degree of imprecision. For example in the case of WM-I-421, 17.6 mm density as found by manual counting is obviously a fault, given that this paper stands out by its very large line spacing (AD751 gives 24.9 in hi-fi mode) – but more subtle errors, of 0.1-0.3 mm, remain unrecognisable. I didn't however see any finer solution.

Also, at the core of the laid lines analysis lies the Fourier transform. So, after overcoming programming errors, the performance that we estimate is that of the transform itself, its ability to correctly detect the number of parallel evenly spaced lines on a given type of image. This way we are lead to discuss some of Fft's positive points and shortcomings.

Image quality expectancies

The Fourier transform behaves particularly well in noisy conditions and still correctly counts lines where manually gross errors would be made. If for the human eye contrast is important for visual analysis, it is not an issue for the Fft. This is the very fact enabling so much information to be extracted from paper reproduction: not only part, but most of the sheet can be analysed (WM-I-134).

However there is a limit. If AD751 uses spots of a width equal to that of the image and if borders such as rules are recommended to be trimmed, it's because in practice I found that blurred spots on the image could cause the Fft to misinterpret the actual number of lines. (Wide spots help avoid cuts in lines to

discharge the existence of those lines and by trimming we avoid problems of knowing over how much space the lines are stretching).

Measurement precision of AD751 – the results

Since human measurements rarely provide more than one density per paper sheet, the value they give is considered to be the average density – and accordingly should have been done on the most regular spots. This is why we will compare the manual results with the mean of all the spot densities of each image obtained with AD751.

Performing the Wilcoxon signed rank test of equality of medians of two matched samples shows that the values obtained with AD751 are not significantly different from the manual results (hypothesis H_0 accepted at $\alpha = 0.05$).

The scores given by AD751 are in an interval of ± 0.5 mm far from those obtained manually, with a $\sim 80\%$ confidence. None is farther than 1 mm and $\sim 10\%$ give equal values. (Extreme divergences [> 1 mm] were found to be human miscounting and excluded [$\sim 8\%$ of 257 cases].)

Interval	Hi-fi	Multi-spot	Global	Global unbiased
$\leq \pm 0.25$ mm	46.50 %	43.60	4.08	43.50
$\leq \pm 0.50$	32.00	33.70	8.16	29.10
$> \pm 0.50$	21.60	22.60	87.80	27.40
Means divergence	0.01	0.03	1.17	0.00
Standard deviation of divergences	0.72	0.71	0.74	0.74
Skewness of divergences	3.99	3.86	4.08	4.08
Probability for H_0 test	0.40	0.99	5e-42	0.16

Fig. 7 – Divergence of AD751's results from manual counts.

Remarkably the mean divergence is practically zero, indicating the closeness between machine and human measurements when considered on an overall basis. There seems to be a slight tendency (skewness ~ 4 mm) in manual counts to overestimate the density – i.e. to stop short before the 20th $\frac{1}{2}$ line.

Translated into pixels the difference amounts to less than 5 pixels (our sample is a 150 dpi contact copy at a 1 pixel = 0.16 mm scale). Considering the imprecision of manual measurement and the fact that we used the mean of all spot results given by AD751 which include data not considered by the human operator, it can be concluded that manual and machine measurements are fairly similar.

Noting that different individuals commonly give different density values for the same sheet of paper, the figures given in the specialised literature should rather be understood as the central point of a range. Thus a reliable automatic laid

lines density measurement tool should greatly benefit the historian's work by its speed, constancy of precision and richness of information.

The three methods' precision & speed

The hi-fi and the automatic multi-spot methods give very close results. The former is preferred if the measurements are operated on individual examples, rather than globally considering a large number of papers.

Using the global method we obtain substantially different values from those given by the other methods. This is due partially to the fact that it relies entirely on the distance of only two lines – the first & the last in the image – and that the surfaces analysed by the three methods are not completely overlapping. Having established its approximate bias (in our test sample approximately equal to the distance between two laid lines), the global mode should be used as a rapid mean to compute the overall density.

In my particular hardware configuration (800MHz, 256RAM) the computing speed was of 400 images / hour in global mode, 300 in automatic multi-spot mode and 50 with the hi-fi method. Manual counting was done at 60 images / hour, for one spot density measured on each paper and physical exhaustion at the end of the third day.

Image slant influence

'Don't touch the image!' is a warning valid in museums as well as in image analysis laboratories. Accordingly, the number of periods per image should be read on the Fft image as the distance between the appropriate maximum and the DC at whatever angle the maximum forms with the vertical axis (laid lines are horizontal). But closer the maximum is to the DC, the more the minimum angle that can be read in the discrete digital space of the image will be higher.

At the frequency of 20 periods / image at which the spots are analysed in hi-fi mode the step is of 2.8° (up to 3.5° in automatic multi-spot that allows for a minimum frequency of 16 periods). Instead if the angle were computed over the whole image, commonly containing more than 60 periods, the precession would be more than 0.9° . This is the reason why we transgressed the rules and rotated the image prior to the analysis.

Influence over the deslanting

The low error slant ($< 3^\circ$) at which the sample images were scanned enables the correct detection of the laid line peaks in the Fft despite the fact that the filter used to establish its probable location is 90° unidirectional. In good quality images, the slant can be even much higher (WM-I-230 manually rotated at 30° is correctly analysed).

However, due to the digital nature of the Fft matrix and to the fact that the slant is inferred solely from the position of the laid line peak, the image can be

rotated only by a minimum angle. Given that in our sample there are between 60 & 100 periods per image, the deslanting can be done by steps of between 0.5 – 0.9° degrees. In fact no images were found to have slants less than 0.9 degrees.

Optical inspection of automatically deslanted images showed no noticeable aberrations. Since the final results were acceptable, we stayed with the simple trigonometric solution for determining the angle of the laid lines peak in the Fft.

Due to the natural wear of the mould, laid lines are not strict parallel lines. This is visible in the Fft: more often than not the laid line peak spills into its neighbourhood and is flanked by a slightly smaller peak.

In the sample lot the slant variation of laid lines within any given image was too low in order to noticeably influence the results. This may not be the case with other types of papers – especially Islamic papers, where laid lines U-shape across the sheet. In this case each spot has to be individually deslanted and the spots with divergent/convergent lines detected.

Experimentally we concluded that in our sample this time consuming analysis would not yield significant improvements and that a global deslanting is sufficient. Moreover, by reducing the analysis spot we would decrease the number of periods per spot and approach the laid line peak to the constant component of the Fft, thus increasing the minimal rotation angle that can be found (to a step of 2.8° in hi-fi mode at 20 periods / image).

A helpful refinement consisted in writing a function that would rotate the selected region of interest in an image without generating black edges. Instead of cutting the region and then rotating it, I increased the region, cut it, rotated and cropped it to the initial dimension. Thus I take advantage of the data surrounding the region of interest.

Influence over the density analysis

Comparing images with 0° slant to slanted images shows a paradoxical situation. There are both more samples of deslanted images closer to values found by manual count and more of them farther away.

Inspection of the most aberrant cases clearly shows human errors. The differences between 0° and non-0° images being low – less than 3% – we can consider that it comes from a mixture of human and image faults. Noteworthy the precision of the analysis does not decrease with the angle of slant: the most slanted image – WM-I-52 at -3° – gives the same result as does the manual count.

In conclusion the slant of laid lines coming from the process of scanning or the actual structure of the mould has minimal impact on the performance of the methods used in AD751.





Fig. 8 – The function 'roirotate' avoids black edges (a) when rotating a region of interest (inner box), by taking advantage of the surrounding data (outer box).



(a)

(b)

Methodological considerations on laid lines density measurement

Number of measurements per paper

Measuring the laid lines density on just one spot of the paper gives too little information compared to what is possible, and it is prone to error. But in manual counting we hardly take the trouble to analyse multiple spots, especially if there are hundreds of papers to analyse. Evaluating the performance of AD751 gave an insightful observation on this matter.

When using the hi-fi AD751 method the mean of deviation from the one-spot-measure manual results is -0.01 mm and 70% of all results fall in a 1 mm range in both cases. This means most of the time the human analyser has chosen the median density spot of a paper. Indeed, this is the way to determine laid lines density with only one measure, if automatic counting is not used. (That he did not chose a spot at random is demonstrated by the fact that if he had, the score would be different, around 90%.)

Ways to present data

In the case of manual measurements of just one spot of the paper structure, the most regular should be chosen and that is considered to be of a median value. Otherwise as many spots as possible should be analysed.

When using an automated procedure it is usually more efficient to present synthesised data, in the following order of relevance: mean density, range and standard deviation. Additionally – mostly in case expertise – the different spot densities can be enumerated in their order of appearance. In extreme cases, the entire mould screen structure could be reconstructed by measuring the distance between every two laid lines (AD751 does not support this functionality in its actual version).

Historical information from analysis data

Paper technology and mean laid lines density

The most used statistics to give information about the paper's laid lines structure is their mean density (X). If spots usually have different values, the mean will correct the imperfections of the screen. The screen will thus look more evenly spaced than before its first usage. The mean produces the ideal density sought by the mouldmaker (supposing however that his intention was not to produce an uneven screen).

As with all measurements, the mean density can be used to expertise two papers, to appreciate their similarity or see if they were produced by the same mould – or more precisely, by the same mould in short laps of time.

When analysing a large number of papers, the mean will reveal an interesting feature, related to advancements in the technology of papermaking. On the basis of the sample used to benchmark AD751, printers from the Netherlands gradually switch between 1474 & 1500 from using papers with a large laid lines spacing to a trend towards dense mould screens. The reason behind this development is the production of finer paper pulps, which need denser screens to be sustained.

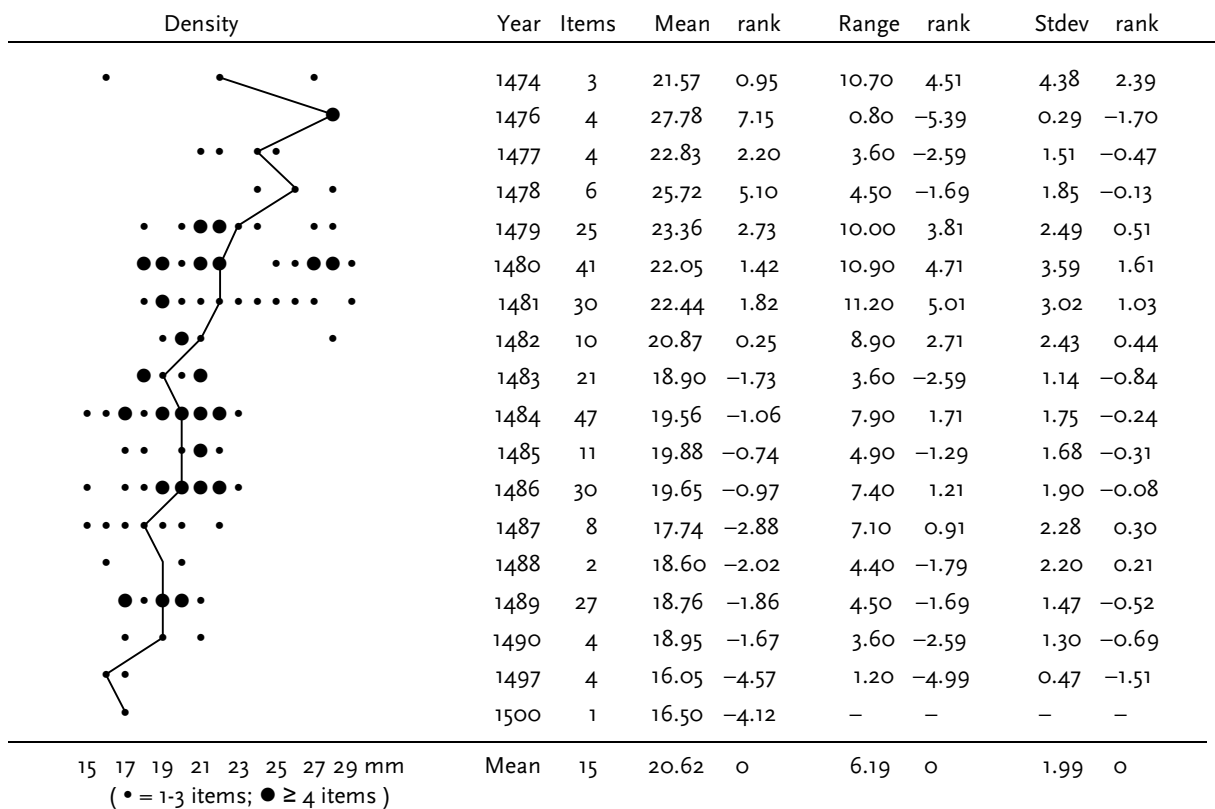


Fig. 9 – Variation of laid lines density in dated papers from the Netherlands.

Once densities are classified in time by using explicitly dated material, it becomes possible to statistically localise papers.

Paper stocks and choice consistency

Not surprisingly, printers who have been in activity for just one year seem to have bought stocks of closely related paper types. This is suggested by the tight range of variation (R) between the density means of paper used by them and the low standard deviation (s) of means.

On the other hand, a long running enterprise such as that of J. de Westfalia from Leuven, used papers coming from a broad range of mould types, the high standard deviation of their density means attesting that they were not interested in harmonizing the paper types they used (at least in what concerns the laid lines).

Printer	Place	Period	Items	Mean	rank	Range	rank	Stdev	rank
Westfalia. J. de	Louvain. B	1474-1486	84	0.83	-0.03	3.10	1.64	0.47	0.04
Fratres	Brussels. B	1476-1481	19	0.88	0.05	1.90	0.44	0.47	0.04
Pafraet. Richardus	Deventer. NL	1477-1479	5	0.74	-0.09	0.50	-0.96	0.22	-0.21
Leeu. Gheraert	Antwerp. B	1477-1484	44	0.88	0.05	1.80	0.34	0.43	0.00
Leeu. Gheraert	Gouda. NL	1478	2	1.40	0.57	1.00	-0.46	0.50	0.07
Vollenhoe. Johannes de	Zwolle. NL	1479	17	0.85	0.02	1.50	0.04	0.36	-0.07
Veldener. J.	Utrecht & Culemborg. NL	1480	9	0.86	0.02	0.70	-0.76	0.26	-0.17
Bellaert. Jacob	Haarlem. NL	1484-1486	27	0.59	-0.24	1.10	-0.36	0.30	-0.13
Os van Breda. Peter van	Zwolle. NL	1485	4	0.48	-0.36	0.30	-1.16	0.11	-0.32
Heerstraten. A. Van der	Louvain. B	1486-1487	29	1.07	0.24	2.60	1.14	0.63	0.20
Meer. Van der or Snellaert	Delft. NL	1488	28	0.88	0.05	2.60	1.14	0.60	0.17
Martens. Theodoricus	Alost. B	1490	4	0.83	-0.01	1.90	0.44	0.70	0.27
Snellaert	Delft. NL	1488-1489	28	0.88	0.05	2.60	1.14	0.60	0.17
Collaciebroeders	Gouda. NL	1497	3	0.73	-0.10	0.90	-0.56	0.37	-0.06
Bac. Govaert	Antwerp. B	1500	1	0.60	-0.23	-	-	-	-
		Mean	(304)	0.85	0	1.61	0	0.43	0

Fig. 10 – Variation of spot densities on the same paper classified by printer.

The shark of Leuven: tricks in quality discovered 500 years later

Out of the pulp evolution arises the question of the quality of paper used in printed books. Looser screens are generally associated with coarser paper surfaces, whereby printers prefer smooth papers that can better retain the details of the printing characters and their legibility.

When considering papers individually, the amount of variation between the spot densities is a measure of the mould's regularity and implicitly of its quality at the moment when the particular paper was produced. A high range of variation will probably mean that the mould has served extensively or that it was loosely

sewn (long: WM-I-57 ; short : WM-I-192). The standard deviation of the densities shows the extent of the paper affected by the irregularity: localised damages give low values and inversely, high values are a sign that the defects are scattered all across the paper.

Classifying papers by printer can order them according to the quality of paper they used. The range of mean variation and its standard deviation are low for the printer Gheraert Leeu, active in Antwerp, when taking into account that he is one of the most prolific printers. The fact that he used papers with relatively minor screen deformations can be a hint to the good quality of his press. On the other hand, the higher production rate of J. de Westfalia – the double of Leeu at roughly the same period – could mean that he was forced to accept whatever paper quality he was presented with in order to feed his machines.

I would like to stress here that the above statements do not come from a specialist of the Dutch incunabula and that they were made by using only a fraction of the available laid lines reproductions connected to the subject. By doing so I wanted to show what wealth of information it is possible to obtain by modern means in a short time – even by a non-specialist. *À bon entendeur salut!*

References

There is no previous work on **computer automating** laid lines measurements, but laid lines are dealt with in watermark extraction – although the objective is here to remove the laid lines from the paper image, similar techniques to ours are used and some of the problems are the same. The Fourier and the morphological approach for line detection are compared in P.F. WHELAN and P. SOILLE (1998), “Automated extraction of paper watermarks”, *Proc. Optical Engineering Society of Ireland & Irish Machine Vision and Image Processing Joint Conference 1998*, National University of Ireland, Maynooth (Ireland) (1998), 287-299 (a subsequent paper on the subject is available on-line at http://www.eeng.dcu.ie/~whelanp/vsg/papers/IMVIP99_Watermark.pdf). See also Christian RAUBER’s web page: <http://cui.unige.ch/~rauber/watermark/filintro.html>.

A good manual to **Image Analysis**, with an up-to-date 2002 version & computational implementation routines on CD, is Bernd JÄHNE, *Digital Image Processing – Concepts, Algorithms, and Scientific Applications* (Berlin: Springer Verlag), 2002 (5th ed.).

His ‘Edges’ chapter (p. 314-337) discusses the characteristics of the ‘**emboss**’ **filter** and compares it to others.

JÄHNE makes also a wide usage of the **Fourier transform**, about which the synthetical presentation in MATLAB’s on-line manual is recommendable: <http://www.mathworks.com/access/helpdesk/help/toolbox/images/transfo2.shtml#17064>.

The “Fourier world” appears in Greg Egan, *Diaspora* (London: Gollancz), 2001, p. 233.

An introduction to **statistics** accessible to humanities scientists is Joseph F. HEALEY, *Statistics – A Tool for Social Research* (Belmont, CA: Wadworth), 2001 (6th ed.).

On-line databases of laid lines reproductions can be found at:

– National Library of the Netherlands Watermark web database
<http://www.kb.nl/kb/bc/incun/watermerken-en.html>

– database of the Austrian National Library in Vienna
<http://www.oeaw.ac.at/ksbm/wz/wzma/index2.htm>

– Archive of Watermarks and Papers in Greek Manuscripts, by Robert W. ALLISON, Bates College, Lewinston, Maine, United States
<http://abacus.bates.edu/Faculty/wmarchive/index.html>

– A Digital Catalogue of Watermarks and Type Ornaments Used by William Stansby in the Printing of The Workes of Benjamin Jonson (London: 1616), by David GANTS

<http://jefferson.village.virginia.edu/gants/>

– The Thomas L. Gravell Watermark Archive, at the Center of Applied Technology, Virginia, United States

<http://ada.cath.vt.edu:591/dbs/gravell/>

Commentaries from a **historical** point of view are: ***.