

MATCHING CANVAS WEAVE PATTERNS FROM PROCESSING X-RAY IMAGES OF MASTER PAINTINGS

Don H. Johnson, Lucia Sun
Elec. & Comp. Engineering
Rice University

C. Richard Johnson, Jr.
Elec. & Comp. Engineering
Cornell University

Ella Hendriks
Conservation Department
Van Gogh Museum

ABSTRACT

Thread counting algorithms seek to determine from x-ray images the vertical and horizontal thread counts (frequencies) of the canvas weave comprising a painting's support. Our spectral-based algorithm employs a variant of short-time Fourier analysis to the image domain that reveals isolated peaks at the proper vertical and horizontal frequencies. Paintings made on canvas sections cut from the same canvas roll have been hypothesized to have similar, distinctive weave characteristics, allowing art historians to more accurately date paintings. Spatial variation of weave frequency measurements across a painting were cross-correlated using a new measure to determine possible common weave patterns between pairs of x-rays. By analyzing a database of x-rays made from 180 paintings by van Gogh, our algorithms confirmed situations where paintings were known to have been made on canvases cut from the same roll and found new ones.

Index Terms— thread counting, art forensics, x-ray image processing, maximal linear correlation

1. INTRODUCTION

To investigate the history and authenticity of paintings by the great masters, signal processing algorithms can provide new insights [1]. Our focus here is on x-ray images that can reveal much about what's below the visible surface [2, 3]. The greater the radiographic-absorbing paint thickness along the beam, the greater the opacity, meaning that x-ray image intensity variations correspond to paint composition and thickness. Letting $i(x, y)$ denote x-ray intensity at a point on the image and z the direction of x-ray propagation,

$$i(x, y) = \exp \left\{ - \int o(x, y, z) dz \right\}. \quad (1)$$

Figure 1 shows an x-ray taken of a painting by Vincent van Gogh. The canvas weave is made visible by the thicker ground (primer) layer of lead-white paint in the grooves between canvas threads as in Figure 2a, the weave pattern can be seen and the vertical and horizontal thread densities can be determined. Thread count data are commonly used as evidence for dating, linking pictures from the same canvas roll, and attribution [2–4].

Considering how a loom works reveals how to think about thread count measurements. The vertical threads mounted in




Fig. 1. X-ray image of van Gogh's *Wild Roses* (catalog number F597 [5]). The painting, the wood stretcher (the lighter border) and canvas-mounting tacks can be clearly seen. The painting itself can also be seen. X-ray image sampled at 600 dpi and provided by the van Gogh Museum.

a loom, known as the *warp*, are usually well aligned with a fairly uniform spacing. The horizontal threads, known as the *weft*, are threaded back and forth through the warp in an interlaced fashion, with the weft compacted occasionally to strengthen the cloth. In most cases, the weft shows more variability than the warp. When the artist cuts a piece of canvas for a painting, he or she will orient the canvas on the stretcher in whatever way seems best: the warp direction may correspond to either the vertical or horizontal threads in the painting. The width of the thread count distribution provides a strong clue as to how the canvas was cut from the roll: one would expect the thread count having the narrower distribution to be the warp direction [3, p. 100]. Thread counts, along with other forensic and historical data, allow the art historian to pose strong hypotheses about how the canvas roll was used for paintings contemporary with each other.

Figure 2a shows a typical test swatch taken from the x-ray image of the van Gogh painting exemplified here. The canvas weave pattern can be clearly seen as well as the artist's brush strokes. Thread counting algorithms seek the weave density, measured in threads/cm, in both the horizontal and vertical directions within a swatch and how these counts vary throughout the painting. The current standards for any measurement technique are manual measurements made with a ruler from a few selected locations in the painting and a human counting of the number of threads in horizontal and vertical directions,

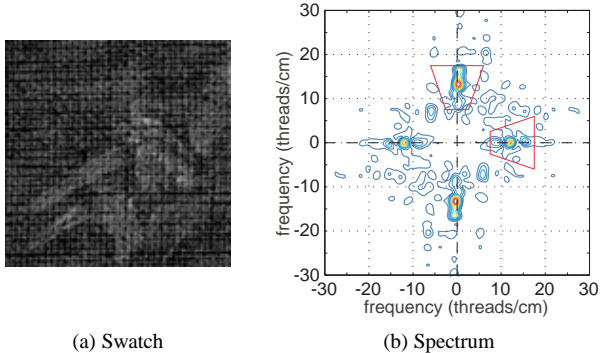


Fig. 2. Panel a shows a 3×3 cm swatch taken from the x-ray of F597. Panel b depicts the spectrum computed from a smaller square (1×1 cm) located in the swatch. The red wedges indicate areas where weave-related spectral peaks can be found.

a tedious process to say the least. Note that the horizontal threads do not run in a precise straight line and that they are not exactly parallel to the x -axis. These effects complicate traditional thread counting algorithms based on direct measurement.

The Thread Count Automation Project seeks to develop signal processing algorithms that can detail the variations in the canvas thread density across a painting [6] and to search for other paintings having a canvas weave pattern that matches it. In his Dutch and late French periods, van Gogh ordered canvas in rolls and, for small to moderate sized paintings, he would cut a rectangular section and mount it on a stretching frame with tacks. The hypothesis is that if two canvas sections share a horizontal or vertical position on the canvas roll, the thread density variations along that axis should agree. Consequently, works associated with the same canvas roll can be presumed to have been painted at about the same time. Since van Gogh worked alone during much of his career, weave-matched paintings could be assumed to be painted by him.

In an internal competition among Project members, a completely automatic spectral algorithm [7], illustrated in Figures 2 and 3, achieved the best performance [8]. Briefly, the spectral algorithm is rooted in the observation that the various layers of paint, including the ground that fills the canvas weave, additively contribute to opacity as expressed by (1). Because the variations of x-ray intensity are relatively small and the tendency of x-ray film exposure to compensate for the exponential, *linear* processing algorithms are most appropriate for extracting thread counts. The algorithm first highpass-filters the x-ray image with a zero-phase FIR filter. It then computes two-dimensional Fourier transforms from raised-cosine-windowed, overlapping sections taken across the entire image. We found that the more computationally demanding high-resolution spectral techniques do not locate spectral peaks more accurately because weave variability cannot be predicted and these methods are not sufficiently robust.

As shown in Figure 2b, peaks located near the vertical and horizontal axes are obviously related to the periodic structure

of the canvas weave. Peak locations are extracted from each spectrum, with a post-processing heuristic applied to resolve cases in which more than one spectral peak emerges because of weave inhomogeneities or “interference” from the artist’s work. Because the sampling grid rarely aligns with thread direction, the radius of the selected spectral peak corresponds to the thread count. We also measured the angle of the spectral peak and discovered that variations of thread angle clearly indicates a phenomenon known as *cupping*: canvas stretching due to nailing the cloth on the stretcher or stretching by the canvas manufacturer to prime the cloth. Typical weave distributions, weave deviation (from average) and thread angle maps are shown in Figure 3. The histograms reveal that the vertical threads correspond to the warp direction on the original canvas roll (the criterion is a smaller spread of the thread count distribution). The weave maps indeed show systematic variations in both the warp and weft directions, a kind of “fingerprint” that can be compared with other paintings. Typical of our investigations, the warp-direction weave maps show a fine, more consistent variability than do the weft-direction maps. The angle maps clearly indicate the presence of cupping of the warp threads along the painting’s left edge and nowhere else, reinforcing the conclusion that this cupping occurred during manufacturing and that the painting’s left edge corresponds to an edge of the canvas roll. To find other paintings that could have come from the same roll, we need an algorithm that finds matching spatial weave variations and locates the paintings’ relative positions in space.

2. WEAVE MAP MATCHING

Using the convention that the warp direction is vertical, paintings made from canvas cut to the left or right of an analyzed painting should share the same variation pattern in weft while one cut from above or below should share the same warp variations. Because of the striping in both the vertical and horizontal weave maps, we averaged the vertical and horizontal counts to create thread count profiles for the vertical and horizontal directions. Thus, for the horizontal thread counts, weave map values in each row were averaged; for the vertical thread counts, columns were averaged. For painting locations where no count was made, no value contributed to the average. We demanded a minimum number of counts contribute to the average; otherwise, no value was provided for the profile at that point. With these one-dimensional summaries of thread density variations, searching for matching x-rays having matching variations can be accomplished with a cross-correlation technique. Because painting orientation cannot be presumed to agree with canvas orientation, taking the various possibilities into account means correlating combinations of profiles and their reversed versions.

Several issues arise when using the usual cross-correlation function normalized to produce a correlation coefficient. First of all, the profiles amount to small deviations added to a large constant. For example, the warp variations of F597 shown in

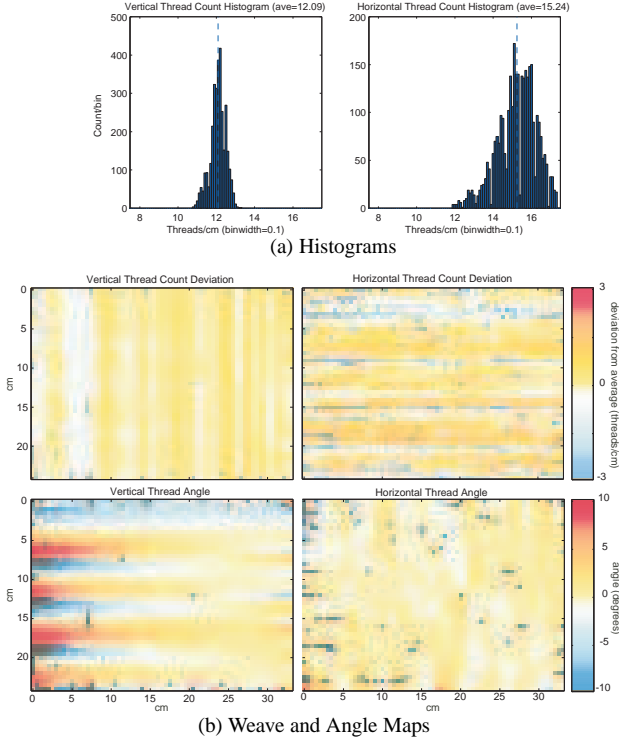


Fig. 3. Thread counts in threads/cm and thread angles in degrees for the van Gogh painting F597 shown in Figure 1. The top row shows the unnormalized histograms of the measured horizontal and vertical thread counts. The middle row shows a heat map of how the thread count measurements deviated from their respective averages: 15.24 threads/cm horizontally and 12.09 threads/cm vertically. Each square corresponds to a spectrum sampled every 1/2 cm across the surface in both x and y . Black pixels indicate where the algorithm made no thread-count estimates. The thread angles are also measured from the spectra and are shown in the bottom row. Manufacturer cusping induces the “rainbow” pattern found in the vertical (warp) thread angles.

Figure 3 is ± 1 thread/cm about an average of 12 threads/cm. Because of the constant, the raw cross-correlation function will be insensitive to thread density variation. Secondly, if each profile’s average is subtracted to remove the constant term, the normalization that is part of computing the correlation coefficient will not take into account the scale of the deviations. Because of these issues, a new cross-correlation method was developed.

The correlation coefficient is rooted in the Cauchy-Schwarz inequality: $|\langle x, y \rangle| \leq \|x\| \cdot \|y\|$. The problem is that equality, equivalent to maximal correlation, occurs when $x \propto y$. We demand maximal correlation when the two quantities are equal, not just proportionality. Simple manipulations lead to what might be called the maximal linear correlation coefficient.¹

$$|\langle x, y \rangle| \leq \|x\| \cdot \|y\| \leq \max\{\|x\|^2, \|y\|^2\}$$

Now, dividing the inner product by the maximum squared

¹The maximal correlation coefficient between two random variables X and Y is defined as the maximal value of $\text{cov}\{\phi(X), \eta(Y)\}$ with respect to all possible functions $\phi(\cdot), \eta(\cdot)$.

norm yields a value of one only when $x = y$. Note that if a constant is subtracted from each the same result applies: $|\langle x - m, y - m \rangle| \leq \max\{\|x - m\|^2, \|y - m\|^2\}$. Removing the average thread count in this way leads to a similarity measure that focuses on the same waveform and amplitude of weave deviations. We take m to be the average of the two profile’s average. Thus, if the two profile’s averages differ, the maximal linear correlation coefficient will be reduced. The resulting maximal cross-correlation function is

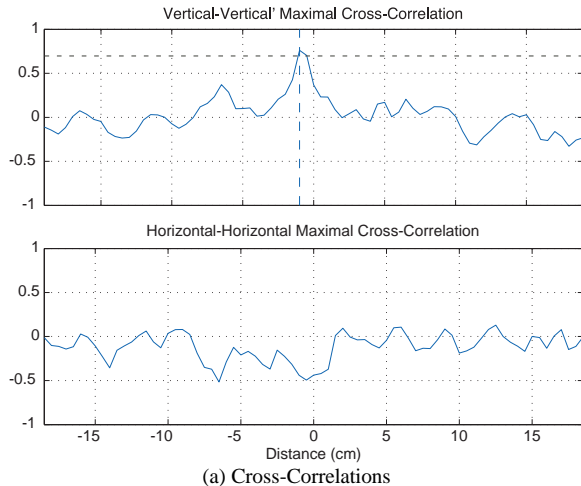
$$r(\ell) = \frac{\sum_m [w_1(m) - \bar{w}] \cdot [w_2(m - \ell) - \bar{w}]}{\max\{\sum_k [w_1(k) - \bar{w}]^2, \sum_l [w_2(l - \ell) - \bar{w}]^2\}}$$

We demanded that at least 10 cm of canvas overlap in the cross-correlation calculations.

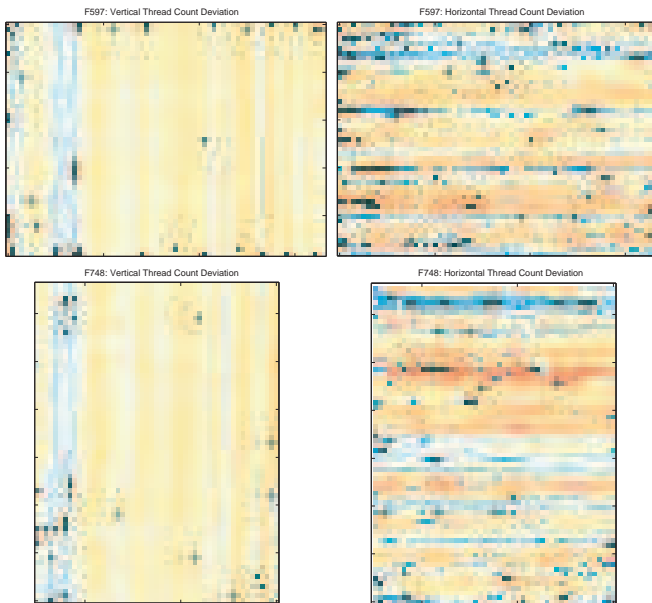
Figure 4a shows the maximal linear cross-correlation between F597 and F748, another painting by van Gogh. The weave maps for these paintings match in warp, the more consistent direction, implying that they were cut from the same canvas roll one above the other, then one was rotated 180°. The maximal linear cross-correlation value for this case was 0.76; the peak is broader than other warp-direction matches we have found. Although not shown here, once rotated to match, F748 also shows primary cusping along the same edge as does F597. In general, we have found that cross-correlation functions for warp-direction matches are far narrower than weft-direction matches and produce larger correlation values (exceeding 0.95 in some cases). We believe that variation in matching correlation values is due to x-ray quality and canvas manufacturing tolerances.

To date, x-rays made from a total of 180 paintings by van Gogh² have been analyzed for thread counts and weave matches. Several cliques of paintings exceed our threshold for declaring a weave match in either warp or weft. We are currently examining these in detail, but one clique of twenty-two paintings stands out. All of these were painted on pieces of “ordinary” quality canvas cut from commercially primed rolls, which van Gogh is known to have customarily ordered from the Paris company Tasset et L’Hôte in the late French period of his production. Painting positions enforced by warp matches span the width of a commercial canvas roll (2m) and extend over a minimum length of 7m (rolls had a maximum length of 10m). The weave match also includes the thread angles: paintings placed along the edges because of the wave match all indicate primary cusping, confirming their putative placement along a canvas roll edge. One painting having an attribution to van Gogh questioned by some experts is part of this clique; this result enhances the possibility van Gogh painted it. The dates of several paintings in this clique can be determined by considering historical information and subject matter. If indeed these paintings were made on sections cut from the same roll of canvas, the datings suggest that the same canvas roll was used over a period of at least eighteen months

²His output is well over 800 paintings. Consequently, less than 20% of van Gogh’s painting output has been examined.



(a) Cross-Correlations



(b) Weave Maps

Fig. 4. Panel a shows the cross-correlation functions between the horizontal and vertical directions for van Gogh’s paintings F597 and F748. A clear correlation peak is evident for the vertical (warp) threads once F748 was rotated 180° . The threshold for peak correlation for warp matches is indicated by the horizontal dashed line. Panel b shows both painting’s weave maps, with the vertical maps aligned according to the correlation function peak. F748’s weave maps have been rotated 180° .

and two residence changes. We believe this an unlikely possibility. A chemical analysis of the ground including factors from the weft weave pattern would contribute additional information so that this clique could be objectively separated into sub-cliques based on sharing several measurements.

3. CONCLUSIONS

Signal processing has shown to play an important role in determining authenticity, as well as helping to date works and provide a better understanding of the sequence of artists’ production. Whereas the signal processing tools described in [1] consider colors and brushstrokes evident at the paint surface,

the algorithms outlined here help to fingerprint the different types of canvas picture support used. Automatic weave estimation and matching techniques have provided far more information about the weave than has been possible to date. This subtle, hidden aspect of a painting can contribute quantitative information to help authenticate and understand the artist’s process. Together, these different approaches provide valuable new insights into the artist’s technical and creative processes, complementing traditional types of information gained by chemical analysis of painting materials and study of historical sources.

As the size of the database increases to include artists’ works spanning four centuries, we are learning the variety of weave patterns used in manufacturing artist canvas. These patterns have different spectra [9], most of which are dominated by horizontal and vertical peaks. We have derived the spectra these weave patterns yield and have developed accordingly spectral algorithms that can cope with all that have been seen to date. In this way, we hope to move toward our goal of a truly automatic thread counting algorithm that provides detailed information for weave matching algorithms.

4. REFERENCES

- [1] C.R. Johnson, Jr., E. Hendriks, I.J. Berezhnoy, E. Brevdo, S.M. Hughes, I. Daubechies, J. Li, E. Postma, and J.Z. Wang, “Image processing for artist identification,” *Signal Processing Magazine*, vol. 25, pp. 37–48, July 2008.
- [2] K.H. Lister, C. Peres, and I. Fiedler, “Appendix: Tracing an interaction: Supporting evidence, experimental grounds,” in *Van Gogh and Gauguin: The Studios of the South*, D.W. Druick and P.K. Zegers, Eds., pp. 354–369. Thames & Hudson, 2001.
- [3] E. van de Wetering, *Rembrandt: The Painter at Work*, Amsterdam University Press, Amsterdam, 1997.
- [4] A. Kirsh and R.S. Levenson, *Seeing through Paintings: Physical Examination in Art Historical Studies*, Yale University Press, 2000.
- [5] J.-B. de la Faille, *The Works of Vincent van Gogh: His Paintings and Drawings*, Meulenhoff, Amsterdam, 1970.
- [6] A.G. Klein, D.H. Johnson, W.A. Sethares, H. Lee, C.R. Johnson, Jr., and E. Hendricks, “Algorithms for old master painting canvas thread counting from x-rays,” in *Asilomar Conference on Signals, Systems and Computers*, 2008.
- [7] D.H. Johnson, C.R. Johnson, Jr., A.G. Klein, W.A. Sethares, H. Lee, and E. Hendricks, “A thread counting algorithm for art forensics,” in *DSP Workshop*, Marco Island, FL, 2009.
- [8] J.S. Ng, “A common platform for thread count automation,” Design Project Report, Cornell University, 2008.
- [9] J. Escofet, M.S. Millán, and M. Ralló, “Modeling of woven fabric structures based on Fourier image analysis,” *Applied Optics*, vol. 40, pp. 6170–6176, 2001.