Signal Processing 93 (2013) 527–540

Contents lists available at SciVerse ScienceDirect

Signal Processing

journal homepage: www.elsevier.com/locate/sigpro

Weave analysis of paintings on canvas from radiographs

Don H. Johnson [a,*], C. Richard Johnson Jr.[b], Robert G. Erdmann [c]

[a] Department of Electrical and Computer Engineering, Rice University, Houston, TX 77005, United States
[b] School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14853, United States
[c] Department of Materials Science and Engineering, University of Arizona, Tucson, AZ 85721, United States

ARTICLE INFO

Article history:
Received 29 August 2011
Received in revised form
26 April 2012
Accepted 25 May 2012
Available online 7 June 2012

Keywords:
X-ray processing
Canvas weave
Thread counting
Spectral analysis

ABSTRACT

A spectral algorithm is described for measuring from radiographs the weaving densities of the horizontal and vertical threads that comprise a painting's canvas. A framework for relating spectra to canvas weave type is presented. The so-called thread density and angle maps obtained from the algorithm reveal the canvas's distinctive density variations and provide insights into how the canvas was prepared. Applying a two-stage correlation procedure to the density variations allowed determination of which paintings' support could have been cut from the same piece of canvas. The first stage uses a nonparametric test of the similarity of the probability distributions of two painting's thread counts. The second stage is a new correlation procedure more stringent than the usual cross-correlation function. Examples drawn from the paintings by Vincent van Gogh illustrate the algorithms.

1. Introduction

Standard measures of a painting's canvas support are its weave pattern – whether it has a plain, twill, herringbone or diamond weave – and its *thread counts*, the average densities (expressed in threads/cm) of the horizontal and vertical threads. Thread count data are commonly used as evidence for dating, linking pictures from the same canvas roll, and attribution [1–3]. For many years, art historians and conservators manually measured from a painting's radiographs the densities at a few locations and averaged them.[1] The algorithm described here measures from radiographs the thread densities (almost) everywhere, thereby providing much more detail than feasible with manual approaches. As a consequence of how the algorithm works, we also measure the thread angles at each location, providing much more information about the canvas and how it was prepared for painting than was previously available.

Considering how a loom works reveals how to think about thread count measurements. The vertical threads (from the weaver's viewpoint in Fig. 1) are known as the *warp* threads that, because of the way looms are constructed, are usually well aligned with a fairly uniform spacing. The weaver threads the horizontal threads, known as the *weft* threads, back and forth through the interlaced warp threads and compacts the weft threads on each pass to strengthen the cloth. In most cases, the weft threads show more weaving and thickness variability than the warp ones. The artist cuts a piece of canvas for a painting from a larger roll and orients it on the stretcher in whatever way seems best[2]; 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

* Corresponding author. Tel.: +1 713 401 9316; fax: +1 713 348 5686.
E-mail addresses: dhj@rice.edu (D.H. Johnson), johnson@ece.cornell.edu (C. Richard Johnson Jr.), erdmann@arizona.edu (R.G. Erdmann).

[1] Radiographs are used instead of inspecting the painting's verso side because master paintings frequently have had a newer canvas glued to the back to preserve the original canvas.

[2] The *stretcher* is typically a wooden rectangle to which the canvas is attached. The carved or molded wooden frame to which the mounted painting is held for display is known as the *frame*.

0165-1684/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.sigpro.2012.05.029

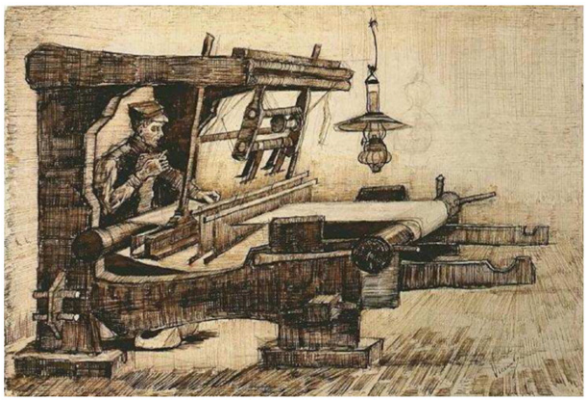

Fig. 1. Vincent van Gogh's drawing *Weaver* catalogued as F1121 [17] and JH453 [18] shows a loom's basic layout. The threads running into the weaver are the warp threads that the apparatus in front of him aligns and separates. The weaver manually slides the weft threads back and forth through the separated warp threads to create the woven canvas, which is rolled onto the take-up drum directly in front of him. Drawing published by permission of the van Gogh Museum.

the larger canvas section. One would expect the thread count having the narrower distribution to be the warp direction [2, p. 100]. Considered jointly, the horizontal and vertical thread count statistics – the averages and the spreads – comprise a crude indicator of whether two paintings' supports could have been cut from the same larger section of canvas. If the thread counts disagree (even when allowing for canvas rotation), then the canvases must have been different. If they agree to some tolerance level, they *could* have had a common origin, but this agreement criterion is very weak. Accurate, comprehensive 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 [4].

To increase the detail of thread counting measurements and enhance the weave matching criterion, we developed a signal processing technique used in [4] for measuring thread counts everywhere across a painting. Our technique is based on "short-space" Fourier analysis, the calculation of the two-dimensional Fourier transforms of overlapping squares (usually 1×1 cm) centered on a sampling grid (usually 0.5 cm spacing). The idea is that a canvas's weave can be considered a doubly, nearly periodic function in two dimensions. Within each square, a thread count measurement amounts to determining the weave's local periodicities. We thus acquire thread count measurements at every grid point, allowing the visualization of what we call the vertical- and horizontal-thread *count maps*.

This paper details our thread counting algorithm and describes its theoretical basis. Once we employed our procedure on paintings, we discovered that the thread counts vary slightly across a painting in a characteristic way that serves as a painting's "fingerprints" [5,8]. Rather than uniquely characteristic of a painting, our results indicate that the "fingerprint" identifies the larger canvas section from which individual canvases were cut. To find paintings that had the same fingerprints, we developed a novel correlation procedure.

2. Theory of weave spectra

Our focus here is on x-ray images that can reveal much about what is below a painting's visible surface [1,2]. The greater the radiographic-absorbing paint and ground[3] thickness along an x-ray beam, the greater the opacity $o(x,y,z)$, meaning that x-ray image intensity variations correspond to paint chemical composition and thickness. Letting $i(x,y)$ denote imaged x-ray intensity at a point and z the direction of x-ray propagation:

$$i(x,y) \propto \exp\left\{-\int o(x,y,z)\, dz\right\} \tag{1}$$

The canvas weave is made visible by the thicker ground (primer) layer of lead-white paint in the interstices between canvas threads. As illustrated in Fig. 2, the weave pattern can be seen and the vertical and horizontal thread densities can be determined from such images. The small variations in opacity due to the ground filling the valleys created by the canvas's threads, coupled with film nonlinearities that tend to reduce the effects of the exponential in this expression, result in image value being proportional to opacity. When the film image is not saturated,[4] the contributions of paint layers, the wood stretcher and tacks superimpose *additively* to form the x-ray image of a painting. Thus, linear processing – measuring spectra in particular – suits this kind of image well.

Appreciating what the spectra of canvas weave patterns should be under ideal conditions allows developing a thread-counting algorithm that exploits the predicted structure. Fig. 2a shows the simplest weave pattern, known as plain weave. The x-ray reveals that the horizontal and vertical threads intertwine in a simple over/under way. Fig. 2b shows a typical twill weave, wherein the threads intertwine in a more complex way. Both examples are periodic but in different ways.

The weave pattern is produced by the interleaved pattern of horizontal and vertical threads. A mathematical description of what the x-ray of a paint-filled ideal canvas-weave surface should be is difficult to determine, to say the least. A more phenomenological model is illustrated in Fig. 2c and d. Here, the vertical and horizontal threads are shown as bars, mimicking the appearance of woven threads. The black rectangles represent unit-height rectangular prisms with the white background corresponding to zero. Because of the measurement process's ultimate linearity, the image's gain (contrast) and offset (overall brightness) do not affect the spectral detail we seek to exploit. The horizontal and vertical thread separations are D_h and D_v, respectively. Each horizontal and vertical thread's thickness and weave are captured by the widths and heights of the bars. For example, the horizontal threads have a thickness H_h and a

[3] *Ground* refers to a special preparatory layer applied to raw canvas to smooth the surface and impart a general color tone to the painting.

[4] In reality, radiographs of paintings by van Gogh and others can have segments exhibiting saturation.

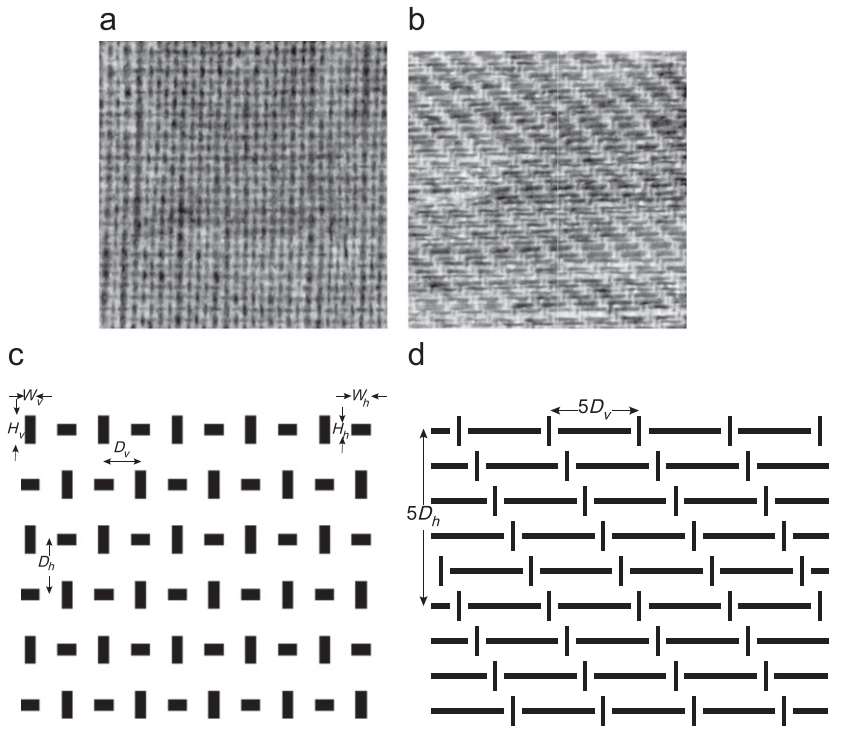

Fig. 2. The top row shows 2×2 cm swatches taken from x-ray images of van Gogh's paintings *Old Man with a Beard* (F205/JH971; left) and *Still Life with Red Cabbages and Onions* (F374/JH1338; right). The left radiograph reveals a plain weave while the right is a twill weave. Idealized woven canvas image models for developing a spectral theory are shown in the bottom row. Our model of the ideal plain weave is expressed in Eq. (2).

width of W_h. The thicknesses and widths must satisfy $D_v \geq (W_h + W_v)/2$ and $D_h \geq (H_h + H_v)/2$. To develop a mathematical expression for this pattern, define $b_h(x,y)$ to be a bar corresponding to a horizontal thread:

$$b_h(x,y) = \begin{cases} 1 & |x| < \frac{W_h}{2}, |y| < \frac{H_h}{2} \\ 0 & \text{otherwise} \end{cases}$$

A similar expression applies to the vertical threads but parameterized by W_v and H_v. The entire weave pattern can be captured as a convolution of the basic thread shapes with a field of impulses that locates them. We thus obtain a general expression for the x-ray image that describes the plain weave pattern and serves as the basis for describing other patterns:

$$\begin{aligned} i(x,y) = & \left[b_v(x,y) \otimes \sum_n \delta(x - 2nD_v) + b_h(x,y) \right. \\ & \left. \otimes \sum_n \delta(x - (2n+1)D_v) \right] \otimes \sum_m \delta(y - 2mD_h) \\ & + \left[b_v(x,y) \otimes \sum_n \delta(x - (2n+1)D_v) + b_h(x,y) \right. \\ & \left. \otimes \sum_n \delta(x - 2nD_v) \right] \otimes \delta(y - (2m+1)D_h) \end{aligned} \quad (2)$$

This complicated expression can be readily deconstructed. To express the top row of Fig. 2c, the first line in square brackets uses convolutions to space the vertical-thread bars $b_v(x,y)$ $2D_v$ apart and the horizontal-thread bars $b_h(x,y)$ by the same amount but shifted to the right by the vertical thread separation (D_v). Continuing the first line, the convolution of this expression with impulses spaced by twice the horizontal thread separation $2D_h$ creates the pattern of every other row. The expression in brackets on the second line shifts the first line's bracketed expression by the vertical thread separation and the outer convolution repeats it, interleaving it with the first line's pattern. A similar expression is developed later for the twill weave.

As we shall see, the key aspects of the spectra calculated from x-rays of the canvas weave are captured by this model. More accurate models would replace the rectangles with smooth, rounded surfaces that represent interwoven thread profiles. Such models would replace the quantities $b_h(x,y)$ and $b_v(x,y)$ with more accurate descriptions, leaving (2) sufficiently correct for algorithm design. We show subsequently that spectral peaks that indicate the thread densities are due to the weave's periodic structure while their amplitudes are modified by the assumed thread model. Furthermore, we need only consider a continuous-space model for the weave patterns because the digitization issues are well-understood. We require radiographs be sampled at sufficiently high frequencies (300–600 dpi) so that aliasing is never an issue. Furthermore, zero-padded transform lengths are much larger than the size of square sections extracted from each painting's radiograph, which has the effect of over-sampling the spectrum to foster accurate determination of spectral peak locations.

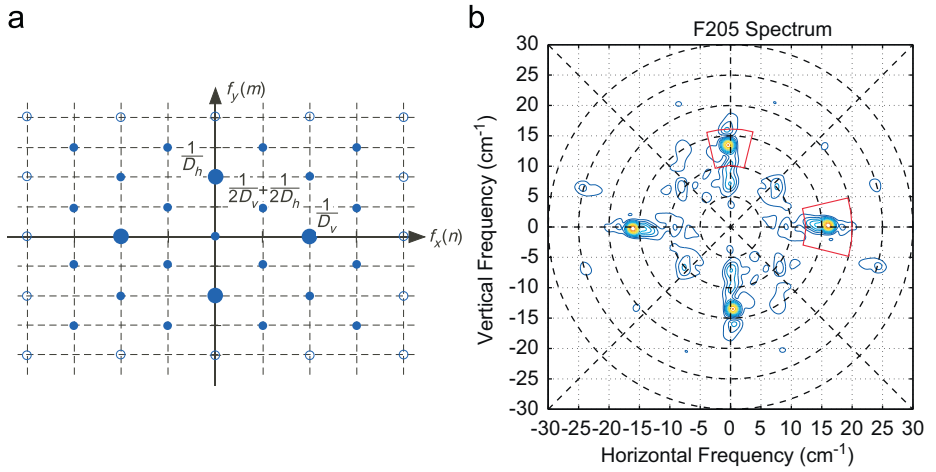

Fig. 3. The frequencies at which the canvas weave spectrum shown in Fig. 2c can be non-zero are shown. Note that $1/D_v$ on the x-frequency corresponds to the vertical thread count and $1/D_h$ on the y-frequency axis to the horizontal thread count. The unfilled circles correspond to frequencies that are identically zero because of the parameter values in the example of Fig. 2c. The spectrum of the swatch shown in Fig. 2a is shown on the right, with annular regions, within which spectral peak locations searches occur, indicated by the red curves. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this article.)

2.1. Model of the plain weave pattern

We denote the Fourier transform of $i(x,y)$ by $I(f_x,f_y)$. Because the images are real-valued, the Fourier transform is conjugate-symmetric, meaning the spectral amplitude values in the third quadrant equal those in the first and those in the fourth quadrant equal those in the second. The spectral plots that follow show all four quadrants for completeness. The Fourier transform of Eq. (2) is easily calculated. The transform of the sums of impulses become sums of impulses located at the separation harmonics. The interleaving shifts become phase terms. Gathering the expression for the spectrum into terms corresponding to the horizontal and vertical bars:

$$\begin{aligned} I(f_x,f_y) = {} & B_v(f_x,f_y)\cdot\left[\sum_{m,n}\delta\left(f_x-\frac{n}{2D_v}\right)\delta\left(f_y-\frac{m}{2D_h}\right)\right. \\ & \left.+e^{-j2\pi f_x D_v}e^{-j2\pi f_y D_h}\sum_{m,n}\delta\left(f_x-\frac{n}{2D_v}\right)\delta\left(f_y-\frac{m}{2D_h}\right)\right] \\ & +B_h(f_x,f_y)\cdot\left[e^{-j2\pi f_x D_v}\sum_{m,n}\delta\left(f_x-\frac{n}{2D_v}\right)\delta\left(f_y-\frac{m}{2D_h}\right)\right. \\ & \left.+e^{-j2\pi f_y D_h}\sum_{m,n}\delta\left(f_x-\frac{n}{2D_v}\right)\delta\left(f_y-\frac{m}{2D_h}\right)\right] \end{aligned}$$

The spectrum contains impulses located on a rectangular grid with centers at $(n/2D_v,m/2D_h)$, the half-harmonics of the basic thread counts (frequencies). Consequently, we need only evaluate the spectrum at these frequencies:

$$\begin{aligned} C\left(\frac{n}{2D_v},\frac{m}{2D_h}\right) &= B_v\left(\frac{n}{2D_v},\frac{m}{2D_h}\right)[1+e^{-j\pi(n+m)}] \\ &\quad +B_h\left(\frac{n}{2D_v},\frac{m}{2D_h}\right)[e^{-j\pi n}+e^{-j\pi m}] \\ &= B_v\left(\frac{n}{2D_v},\frac{m}{2D_h}\right)[1+(-1)^{n+m}] \\ &\quad +B_h\left(\frac{n}{2D_v},\frac{m}{2D_h}\right)[(-1)^n+(-1)^m] \end{aligned} \tag{3}$$

The expressions for the Fourier transforms of the bars are

$$\begin{aligned} B_h\left(\frac{n}{2D_v},\frac{m}{2D_h}\right) &= W_h H_h\,\mathrm{sinc}\left(\pi n\frac{W_h}{2D_v}\right)\mathrm{sinc}\left(\pi m\frac{H_h}{2D_h}\right) \\ B_v\left(\frac{n}{2D_v},\frac{m}{2D_h}\right) &= W_v H_v\,\mathrm{sinc}\left(\pi n\frac{W_v}{2D_v}\right)\mathrm{sinc}\left(\pi m\frac{H_v}{2D_h}\right) \end{aligned} \tag{4}$$

where $\mathrm{sinc}(x)=\sin x/x$.

To interpret the result shown in (3), note that when the sum of the indices n,m is odd, the spectrum is zero. Fig. 3 shows the locations in the spectrum that can be non-zero. The spectra of the bars given by the previous expressions provide the spectral values at these frequencies. In general, the larger the frequency indices, the smaller these spectral values will be. The largest spectral value is located, of course, at the origin. It provides no information about the canvas weave but it must be removed to produce accurate spectral estimates.[5]

Considering (3), the value, even existence of a spectral peak, depends on two factors: the value of the spectra (4) at harmonics of thread count frequencies and the signs of the indexed terms $1+(-1)^{n+m}$ and $(-1)^n+(-1)^m$. The bar spectra depend on horizontal and vertical thread separations and on the heights and widths of the bars. For example, in Fig. 2c, $W_h/D_v=1/2$, $H_h/D_h=1/5$, $W_v/D_v=2/7$ and $H_v/D_h=1/2$. The first zero of $\mathrm{sinc}(x)$ occurs when $x=\pi$. Consequently, $B_h(2/D_v,m/2D_h)$, $B_h(n/2D_v,5/D_h)$, $B_v(7/2D_v,m/2D_h)$ and $B_v(n/2D_v,2/D_h)$ all equal zero. Thus, particular values of thread thicknesses and densities will cause spectral components to disappear. Furthermore, the first indexed term $1+(-1)^{n+m}$ equals either 0

[5] The current spectral thread counting algorithm highpass-filters the image before calculating spectra to remove the spectral spread a peak at the origin would induce. In more detail, we apply a zero-phase, separable (rectangular unit-sample response in each direction) low-pass filter having a cutoff frequency of a few threads/cm to the radiograph, then subtract result from the original radiograph before making spectral calculations.

or 2; the second, $(-1)^n+(-1)^m$, equals -2, 0 or 2. Both terms equal zero when the sum $n+m$ is odd. The primary spectral peaks at $1/D_h$ and $1/D_v$, which correspond to the vertical and horizontal thread counts, respectively, cannot be zeroed this way. The second term equals -2 when both indices are odd, which introduces the possibility that the two terms in (4) could nearly cancel each other. In particular, the diagonal component at $n=1$, $m=1$ falls into this category. In this case,

$$C\left(\frac{1}{2D_v},\frac{1}{2D_h}\right)=2B_v\left(\frac{1}{2D_v},\frac{1}{2D_h}\right)-2B_h\left(\frac{1}{2D_v},\frac{1}{2D_h}\right)$$

For the plain weave pattern shown in Fig. 2c, the spectral values at the primary thread count frequencies is almost a factor of five greater than the diagonal frequency value. Letting $\tilde{I}(n,m)=I(n/D_v,m/D_h)$, with the numbers in parentheses being normalized values,

$$\tilde{I}(0,2)=0.14(1.00)\quad \tilde{I}(2,0)=0.14(1.00)\quad \tilde{I}(1,1)=0.03(0.21)$$

Consequently, we cannot always expect a diagonal spectra peak located at the vector-average of the horizontal and vertical frequency peaks. Presumably, this situation occurs in the spectrum shown in Fig. 3a and b and is due to similar thread characteristics. When the vertical and horizontal threads are identical, the prediction is no peak on the diagonal. This analysis implies that a peak on the diagonal will occur when the vertical and horizontal threads have very different characteristics.

2.2. Model of twill weave patterns

The plain weave pattern analyzed in the previous section corresponds to a weaving pattern wherein each horizontal thread passes over one vertical thread then under the next one, etc., with the vertical threads following the same pattern. *Twill* patterns vary the plain weave pattern in a simple way. Each horizontal thread passes over T_o vertical threads then under T_u vertical threads. The next horizontal thread starts its interleaving pattern T_s threads from the previous. Fig. 4 shows two such patterns that have $T_s=1$. For such twill patterns, the notation for the pattern is $T_o:T_u$. The first pattern shown in panel 4a shows an over-four/under-one pattern, making it a 4:1 twill. The second (panel 4b) shows an over-four/under-two pattern, making it a 4:2 twill. Note that these different patterns share the same thread counts, differing only in the heights of the vertical bars and the widths of the horizontal ones. Each row in both cases repeats horizontally and vertically at a separation of $(T_u+1)D_v$ and $(T_u+1)D_h$, respectively. The

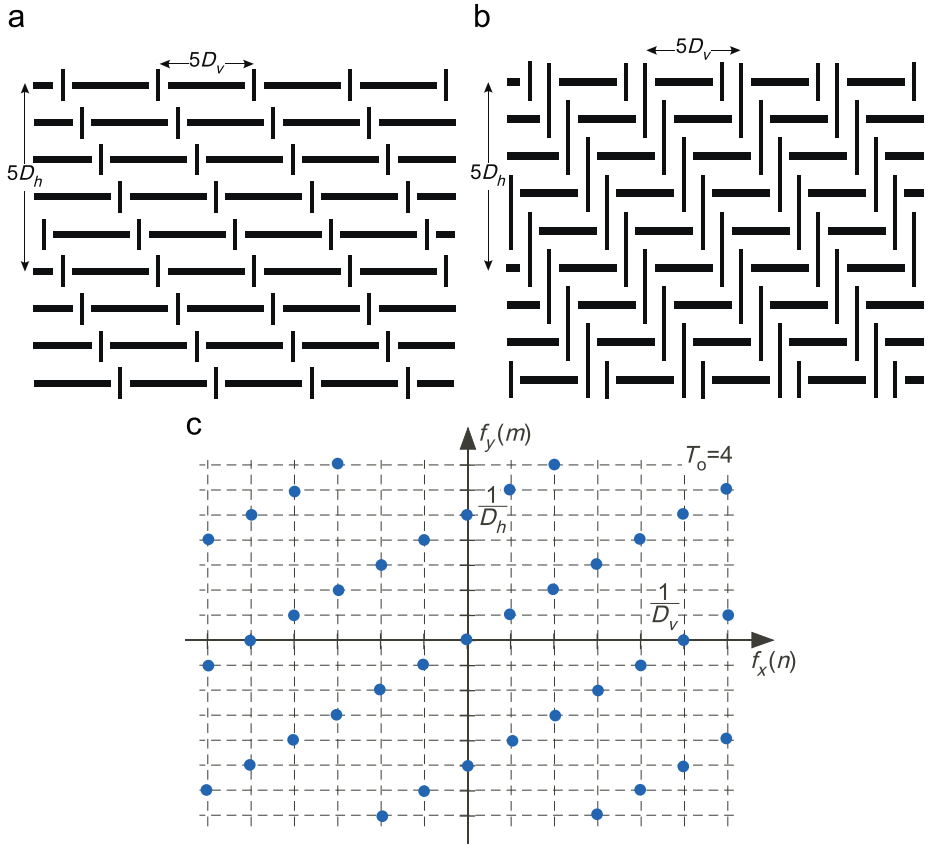

Fig. 4. Panel (a) shows a 4:1 twill pattern, wherein each horizontal thread passes over four vertical threads, then under one. Panel (b) shows a 3:2 twill pattern, in which each horizontal thread passes under two vertical threads. Here, $D_h=1.2D_v$. For panel (a), $W_h=4D_v$, $H_h=D_h/5$, $W_v=D_v/5$ and $H_v=D_h$; in panel (b), $W_h=3D_v$, $H_h=D_h/5$, $W_v=D_v/5$ and $H_v=2D_h$. Panel (c) shows the predicted spectral locations for the 4:1 twill pattern.

northwest-southeast pattern of the vertical threads can run northeast-southwest. For the derivation here, we consider the former; handling the second amounts to changing a few signs.

To describe mathematically twill patterns, we simply extend the description of plain weave patterns:

$$\begin{aligned} i_{T_o:1}(x,y) = b_v(x,y) \otimes \Bigg[& \sum_{k=0}^{T_o} \sum_{m,n} \delta(x-(n(T_o+1) \pm k)D_v) \\ & \delta(y-(m(T_o+1) \mp k)D_h) \Bigg] \\ & + b_h(x,y) \otimes \Bigg[\sum_{k=0}^{T_o} \sum_{m,n} \delta\left(x-\left(n(T_o+1) \right.\right. \\ & \left.\left. \pm k \pm \frac{T_o+1}{2}\right)D_v\right) \delta(y-(m(T_o+1) \mp k)D_h) \Bigg] \end{aligned} \tag{5}$$

$$\begin{aligned} i_{T_o:2}(x,y) = b_v(x,y) \otimes \Bigg[& \sum_{k=0}^{T_o} \sum_{m,n} \delta(x-(n(T_o+1) \pm k)D_v) \\ & \delta(y-(m(T_o+1) \mp k)D_h)] \\ & + b_h(x,y) \otimes \Bigg[\sum_{k=0}^{T_o} \sum_{m,n} \delta\left(x-\left(n(T_o+1) \pm k \pm \frac{T_o}{2}\right)D_v\right) \\ & \delta\left(y-\left(m(T_o+1) \mp k-\frac{1}{2}\right)D_h\right) \Bigg] \end{aligned} \tag{6}$$

The choice of sign in the $\pm$ terms depends on the direction of the diagonal pattern. For the patterns shown in Fig. 4, we choose the plus sign (and minus sign for the $\mp$ term). Concentrating on this choice and converting to the frequency domain simplifies this expression considerably. Looking first at the 4:1 twill,

$$\begin{aligned} I_{T_o:1}(f_x,f_y) = & B_v(f_x,f_y) \sum_{m,n} \delta\left(f_x-\frac{n}{(T_o+1)D_v}\right) \\ & \delta\left(f_y-\frac{m}{(T_o+1)D_h}\right) \frac{\sin \pi(T_o+1)(D_v f_x - D_h f_y)}{\sin \pi(D_v f_x - D_h f_y)} \\ & + e^{-j\pi(T_o+1)D_v f_x} B_h(f_x,f_y) \sum_{m,n} \delta\left(f_x-\frac{n}{(T_o+1)D_v}\right) \\ & \delta\left(f_y-\frac{m}{(T_o+1)D_h}\right) \frac{\sin \pi(T_o+1)(D_v f_x - D_h f_y)}{\sin \pi(D_v f_x - D_h f_y)} \end{aligned}$$

The spectrum can be non-zero only when $f_x = n/(T_o+1)D_v$, $f_y = m/(T_o+1)D_h$:

$$\begin{aligned} & I_{T_o:1}\left(\frac{n}{(T_o+1)D_v}, \frac{m}{(T_o+1)D_h}\right) \\ & = \frac{\sin \pi(n-m)}{\sin \frac{\pi(n-m)}{(T_o+1)}} \left[B_v\left(\frac{n}{(T_o+1)D_v}, \frac{m}{(T_o+1)D_h}\right) \right. \\ & \left. + (-1)^n B_h\left(\frac{n}{(T_o+1)D_v}, \frac{m}{(T_o+1)D_h}\right) \right] \end{aligned} \tag{7}$$

First of all, note that when $T_o = 1$, this expression is exactly the same as the one for the plain weave. The ratio of sines term is non-zero only when $n-m = k(T_o+1)$, k any integer, which corresponds to the grid indicated by the dots shown in Fig. 4c. Depending on the spectral characteristics of the vertical and horizontal bars, the spectrum at some grid locations can be very small compared to others. Furthermore, the value of n determines whether the bar spectra add or subtract. These spectra are given by the same formula as (4) but are evaluated at different frequencies:

$$\begin{aligned} & B_h\left(\frac{n}{(T_o+1)D_v}, \frac{m}{(T_o+1)D_h}\right) \\ & = W_h H_h \operatorname{sinc}\left(\pi n \frac{W_h}{(T_o+1)D_v}\right) \operatorname{sinc}\left(\pi m \frac{H_h}{(T_o+1)D_h}\right) \end{aligned}$$

$$\begin{aligned} & B_v\left(\frac{n}{(T_o+1)D_v}, \frac{m}{(T_o+1)D_h}\right) \\ & = W_v H_v \operatorname{sinc}\left(\pi n \frac{W_v}{(T_o+1)D_v}\right) \operatorname{sinc}\left(\pi m \frac{H_v}{(T_o+1)D_h}\right) \end{aligned}$$

However, the twill pattern constrains the values of two bar parameters, $T_o D_v > W_h \geq (T_o-1)D_v$ and $2D_h > H_v \geq D_h$.

Returning to the case of the 4:2 twill, its spectrum is also non-zero only when $n-m = k(T_o+1)$. The expression differs only in the phase relationship between the bar spectra:

$$\begin{aligned} & I_{T_o:2}\left(\frac{n}{(T_o+1)D_v}, \frac{m}{(T_o+1)D_h}\right) \\ & = \frac{\sin \pi(n-m)}{\sin \frac{\pi(n-m)}{(T_o+1)}} \left[B_v\left(\frac{n}{(T_o+1)D_v}, \frac{m}{(T_o+1)D_h}\right) \right. \\ & \left. + \exp\left\{-j\pi \frac{nT_o-m}{T_o+1}\right\} B_h\left(\frac{n}{(T_o+1)D_v}, \frac{m}{(T_o+1)D_h}\right) \right] \end{aligned} \tag{8}$$

Furthermore, the constraints on the bar characteristics are different—$(T_o-1)D_v > W_h \geq (T_o-2)D_v$ and $3D_h > H_v \geq 2D_h$—making the bar spectra different as well. A comparison of the spectra for the 4:1 and 4:2 twill patterns (Fig. 4) reveals that the difference is the values of the spectra when evaluated on the grid shown in Fig. 4c. The primary difference between the two is the appearance of large spectra values on the diagonal $n=m$ for the 4:2 twill, which are absent for the 4:1 twill. In both cases, the prediction is that spectral peaks will appear on the horizontal and vertical frequency axes that correspond to the thread counts we seek. However, the vertical-frequency peak is predicted to be bigger than the horizontal-frequency peak, especially for the 4:1 twill. Also note the appearance of larger spectral values near the on-axis grid points for the 4:2 twill case. These spectral peaks make it difficult to find the vertical-frequency peak *and* allow for spectral rotation due to deviation of the weave pattern from an ideal horizontal/vertical orientation. The spectrum shown in Fig. 5c of a twill swatch demonstrates the expected diagonal spectral locations in the second and fourth quadrants. The presence of strong spectral values on these diagonals indicates that the twill factor is 4:2, an observation not easily discerned from visual inspection of the x-ray, at least for the ones we have examined.

3. From spectra to algorithm

The two-dimensional regularity of the canvas weave patterns shown here directly point to a spectrally based algorithm for estimating the two thread counts. As the theory shows, the peak on the horizontal axis corresponds to the vertical-thread density, the peak on the vertical axis

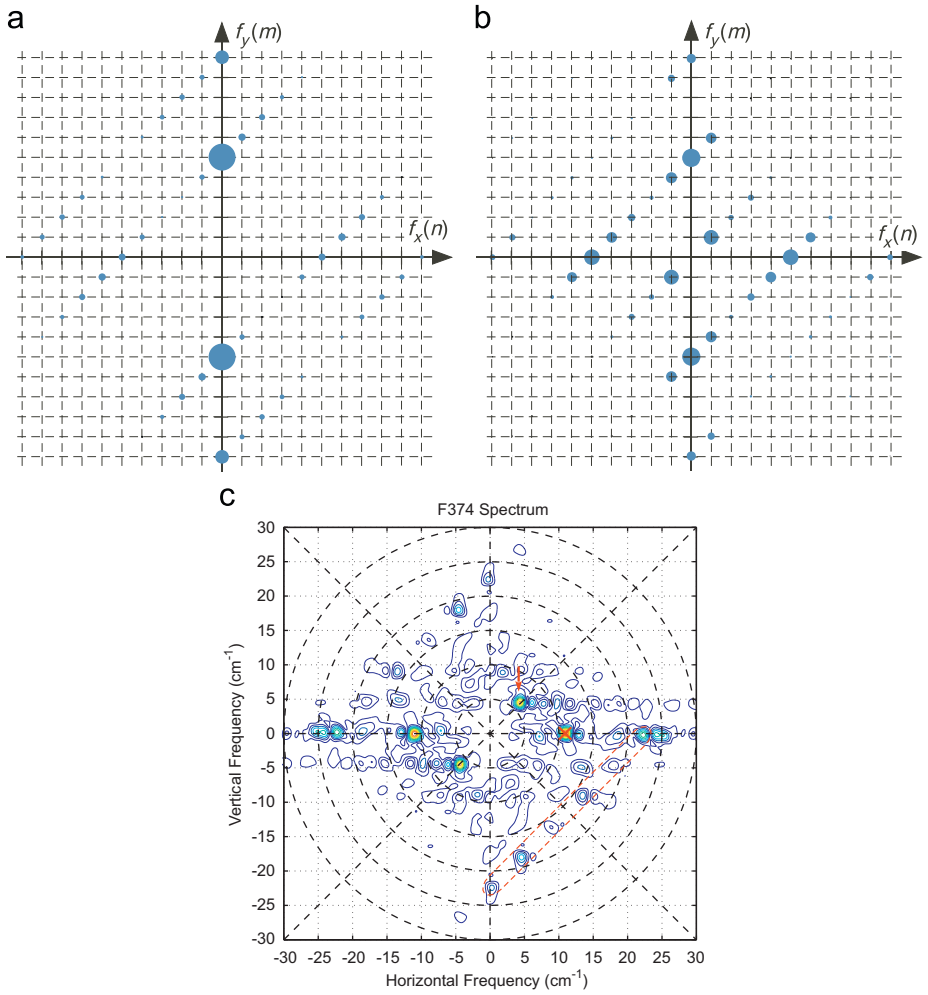

Fig. 5. The spectra predicted by Eqs. (7) and (8) are shown in panels (a) and (b) for the 4:1 and 4:2 twills, respectively. Panel (c) shows the spectrum of the twill swatch shown in Fig. 2b. The encircled sequence of peaks correspond to the line of peaks in the ideal spectrum shown in the fourth quadrant of (b). The peak indicated by the arrow corresponds to the strong peak in the ideal spectrum located at (1,1). The strong peak on the horizontal axis marked by an '$\times$' occurs at half the frequency of the actual vertical-thread density peak we seek. A color version of this figure can be found online.

to the horizontal-thread density. Simplistically, we would search for spectral peaks on the horizontal and vertical axes, the locations of which corresponds to the thread densities we seek. The theory predicts that this approach would work for both the weave patterns considered here.

But in more detail, weave-related spectral peaks cannot be presumed to be precisely located on the frequency axes. As Fig. 3b shows, tilting of the horizontal and vertical threads from alignment to the sampling grid leads to a corresponding rotation of the spectrum, a property of the two-dimensional Fourier transform [6]. Thread angles are also distorted by the phenomenon known to art conservators as *cusping*: the displacement of the threads from a rectilinear pattern caused by attaching raw canvas to the stretcher or priming frame prior to surface preparation. To envision cusping, consider vertical threads that lie next to the edge of a stretcher's vertical member. When the canvas is pulled laterally with a point-applied force induced by the attachment, the vertical threads closest to the edge and the attachment point are displaced. At horizontal and vertical locations away from the point force application, little thread displacement occurs. In this way, the originally linear paths of threads now show scalloping, which means locally that the thread count may decrease slightly while the thread angle swings positively then negatively (or vice versa) at the attachment points. Possible angular rotation and the expected range for the thread counts imply that the search for weave-pattern-related spectral peak locations can focus on annular wedges centered on the frequency axes, indicated by the red annular regions in Fig. 3b. The parameters of these wedges – the two radii and the angular spread – must be determined on a painting-by-painting basis. We use a simple graphical user interface to set the parameters for the nominally vertical and horizontal spectral search regions. In effect, focusing the search amounts to bandpass filtering the x-ray image to remove non-weave related "interference," such as stretcher boundaries, tacks and stretcher bars, and the painting itself. This

interference has no typical spectral signature but in most cases, but certainly not all, occurs "out of band." Once the rotations and expected ranges are used to focus the spectral analysis, the vertical and horizontal thread counts are taken to be radii of the nearly horizontal and nearly vertical spectral peak locations, respectively, that lie in the annular regions. If multiple peaks occur within a given region, the algorithm chooses the location closest to spectral peak locations found in single-peak regions. The thread angles correspond to the angles of the chosen spectral peak locations. About 90% of the image excerpts surrounding the locations of spot hand counts on the examined paintings by van Gogh have peak locations within 1 thread/cm of the manual thread count, as do approximately 80% of examined spots in paintings by Matisse, approximately 75% of examined spots in paintings by Vermeer, and about 65% of examined spots in paintings by Rembrandt [7].

Fig. 6 shows the results of spectral-based thread counting analysis for the example plain-weave painting. As described in [8], the count maps show variations of thread densities characteristic of the original canvas section from which this painting's support was cut. Consequently, by searching for other paintings having a similar pattern, we could possibly reconstruct the positions of matching paintings on the canvas roll, which could provide dating information [4,5]. The horizontal-thread angle map shows strong cusping along the top but not along the bottom. This result indicates that a larger canvas was prepared with a common ground layer and then this painting's canvas was cut from that prepared canvas.

Obtaining thread count results for the twill-weave painting in Fig. 2b is not as straightforward. If only the (approximately) horizontal- and vertical-axis peaks are used, inconsistent results are obtained because of weaving variability and a relatively poor-quality x-ray. Obtaining thread counts required exploitation of the twill weave's spectral properties. Our spectral theory predicts that the off-axis trail of spectral peaks in the fourth quadrant of Fig. 5a and b is completely determined by the horizontal and vertical thread patterns. Using vector notation for the spectral peak locations, the *n*th peak is

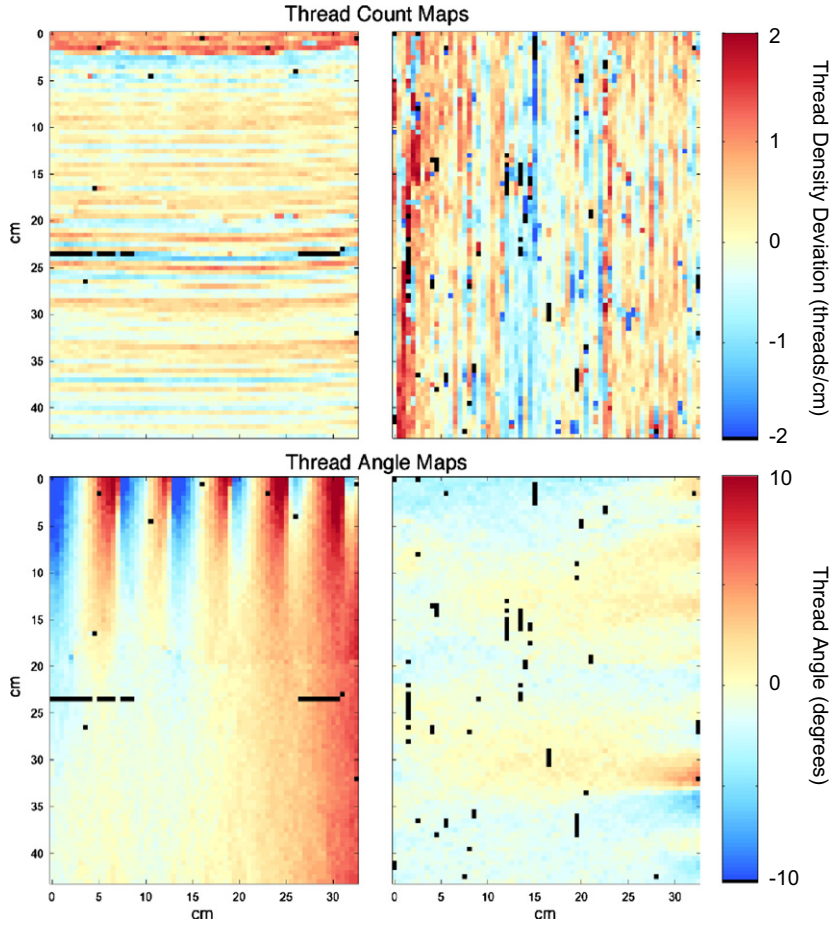

Fig. 6. Count maps – displays of departure of thread densities from painting-wide average values – are shown in the top row for painting van Gogh's painting F205. The average thread counts are 13.3 threads/cm (horizontal) and 16.0 threads/cm. The color map (from deep blue to deep red) spans ± 2 threads/cm. Angle maps – displays of the departure of the thread from horizontal/vertical are shown in the bottom row. The color map spans $\pm 10^{\circ}$. Black areas in the count and angle maps correspond to locations at which the current algorithm could not determine a value. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this article.)

located at

$$\vec{f}_n = \frac{(T+1-n)\vec{f}_h - n\vec{f}_v}{T+1}, \quad n = 0, \ldots, T+1$$

where $\vec{f}_h$ and $\vec{f}_v$ are the sought spectral peak locations on the horizontal and vertical frequency axes, with $n=0$ corresponding to the horizontal frequency location and $n=T+1$ the vertical frequency location. Here, T is the twill factor, equal to 4 in this case. This system of equations can be used to determine $\vec{f}_h$ and $\vec{f}_v$ by measuring the locations of the second-quadrant peaks, each of which was obtained by establishing annular off-axis regions for peak location searches. In many cases, not all of the predicted peaks can be located, which is an effect due to inaccuracies in our model of threads as rectangles. Because the set of equations relating measured values of $\vec{f}_n$ to $\vec{f}_h$ and $\vec{f}_v$ is not invertible and supports multiple solutions, our algorithm uses a consensus algorithm that finds the L_1 centroid of the various solutions for $\vec{f}_h$ and $\vec{f}_v$ to estimate the thread densities we seek. Fig. 5c shows that not all of these fourth-quadrant peaks have a significantly large amplitude. In our example, $\vec{f}_3$ was not measured; the remaining five peak locations were used to estimate the thread counts and these results are shown in Fig. 7. Again, cusping is present only along the right side, indicating the same canvas preparation as in the plain-weave example.

Several circumstances contribute to the algorithm's inability to obtain a reliable thread density estimate at certain locations. The bandpass-filtering implicit in the algorithm cannot remove in-band interference from the brushstrokes, the wood grain of some stretchers and tacks. Such interference has no single model, certainly not over the 1 cm^2 squares we typically use in processing. When multiple peaks occur in-band that cannot be related to single-peak results, the algorithm does not produce an estimate for one or both thread counts. Furthermore, even for plain weave patterns, "on-axis" spectral peaks can be weak or non-existent. One fallback is to use an approach similar to that employed for twill: use predicted off-axis spectral peaks to infer the frequency of the obscured or missing "on-axis" peaks. As shown in Fig. 3, our theory predicts a spectral peak will be present at the vector-average of the desired thread densities $\vec{f}_v \pm \vec{f}_h$. Just as with twill, all the off- and on-axis spectral peaks judged to be present are used to infer values for $\vec{f}_v$ and $\vec{f}_h$.

4. Thread density map matching

Once the horizontal- and vertical-thread count maps have been measured for two paintings, we want to determine if their maps match (agree). Before searching for a match, we found it convenient to relate horizontal and vertical to warp and weft. In essence, we want to convert from painting coordinates to canvas coordinates. Using the convention that each painting's count and angle maps are oriented so that the warp-thread direction is vertical, we mimick the appearance of a loom's thread layout from the weaver's viewpoint (Fig. 1). Paintings made from canvas cut to the left or right of an analyzed painting should share the same variation pattern in weft (always running horizontally after being re-oriented) while ones cut above or below should share the same warp-thread density variations. Because of thread count consistency along the thread direction, shown by the

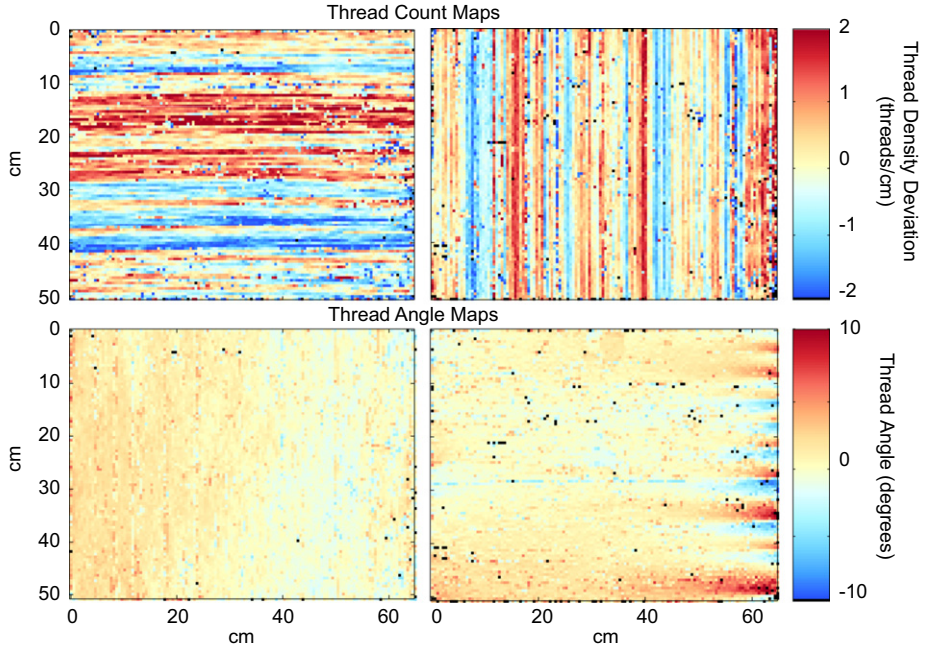

Fig. 7. Using the same format as in Fig. 6, the count and angle maps for van Gogh's painting F374 and executed on a twill canvas are shown. The average thread counts are 22.2 threads/cm horizontal and 22.7 threads/cm vertical. A color version of these count maps can be found online.

striping in both the vertical and horizontal count maps [7], we averaged the vertical and horizontal counts along their respective thread directions to create what we call thread count *profiles* for the vertical and horizontal threads. We take into account an imperfect alignment of the canvas with the stretcher by rotating the measured thread counts by the negative of the average thread angle before calculating profiles. Thus, for the weft thread counts, count map values in each row were averaged; for the warp thread counts, columns were averaged. For painting locations where no count was made, no value contributed to the average. We demanded a minimum number (ten) 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 eight combinations of profiles and their reversed versions [8].

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 F205 shown in Fig. 6 span ± 1 thread/cm about an average of 13.3 threads/cm. Because the thread density variations are much smaller than the average, the raw cross-correlation function will be insensitive to thread density variations. Second, 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 two-stage method was developed. The first ensured that the histograms of vertical and horizontal thread densities agreed well as detailed in the next section. Once agreement was found, the second stage, described in Section 4.2, subtracted an average count from the profile and computed a cross-correlation function with a new method that demands similar count variations in both space and amplitude.

4.1. Comparing large-scale thread density variations

The problem of determining if the average thread densities agree must take into account the measurement's variability. Our approach was to determine whether the empirical thread density distributions agree. The hypothesis testing problem for determining if two measurements $\mathbf{x}_1$ and $\mathbf{x}_2$ came from the same probability distribution or from different ones can be formally stated as [9]

$$H_0 : p_{\mathbf{X}_1}(\mathbf{x}) = p_{\mathbf{X}_2}(\mathbf{x}) \equiv p_{\mathbf{X}}(\mathbf{x})$$

$$H_1 : p_{\mathbf{X}_1}(\mathbf{x}) \neq p_{\mathbf{X}_2}(\mathbf{x})$$

The likelihood ratio solution to this hypothesis test is, assuming the two measurements are statistically independent,

$$\frac{p_{\mathbf{X}_1}(\mathbf{x}_1)\cdot p_{\mathbf{X}_2}(\mathbf{x}_2)}{p_{\mathbf{X}}(\mathbf{x}_1)\cdot p_{\mathbf{X}}(\mathbf{x}_2)} \gtrless \eta \tag{9}$$

where $p_{\mathbf{X}}(\cdot)$ denotes the shared distribution of the two measurements under hypothesis H_0. If the likelihood ratio evaluated at the measurements is greater than the threshold η, the data are judged to have come from different distributions, which would mean in the context of comparing thread counts that the paintings could not share a common canvas.

The main issue is that the probability distributions are not known. In the following parametric framework, the form is known (Gaussian) but not the values of crucial parameters (mean and variance). In a non-parametric framework, no assumptions are made about the distributions. The well-known solution to both scenarios is to use maximum likelihood estimates of the parameters or of the distributions calculated separately for the numerator and denominator since the two hypotheses (models) differ.

Parametric solution: Here we assume that $X \sim \mathcal{N}(m,\sigma^2)$ under all models, with the distributional "fit" determined by the maximum likelihood parameter estimates. The logarithm of the multivariate Gaussian density with statistically independent, identically distributed components becomes, upon substitution of the maximum likelihood estimates, $-(N/2)\ln 2\pi e\widehat{\sigma^2}$, where $\widehat{\sigma^2}$ denotes the maximum likelihood estimate of the variance.

For the different-distribution hypothesis H_1, we simply employ this result for both datasets *separately*. For the common-distribution hypothesis, the two datasets are merged to produce single estimates of the mean and variance. After evaluating the logarithm of the likelihood ratio (9), which becomes $-(N_1+N_2)/2\ln 2\pi e\widehat{\sigma^2_{1,2}}$ after substituting the maximum likelihood estimate, we obtain a simple test based *entirely* on the variance estimates:

$$-\frac{N_1}{2}\ln\widehat{\sigma_1^2} - \frac{N_2}{2}\ln\widehat{\sigma_2^2} + \frac{N_1+N_2}{2}\ln\widehat{\sigma_{1,2}^2} \gtrless \gamma$$

The variance estimate $\widehat{\sigma_{1,2}^2}$ is computed from the merged datasets. The resulting mean and variance estimates become

$$\widehat{\sigma_{1,2}^2} = \frac{1}{N_1+N_2}\left(\sum_i (X_{1,i}-\widehat{m_{1,2}})^2 + \sum_i (X_{2,i}-\widehat{m_{1,2}})^2\right)$$

$$\widehat{m_{1,2}} = \frac{1}{N_1+N_2}\left(\sum_i X_{1,i} + \sum_i X_{2,i}\right) \tag{10}$$

If they were drawn from different distributions, even if it is only in the means, the common-distribution term will differ greatly from the different-distribution terms.

When using maximum likelihood estimates in the likelihood ratio, the choice of the threshold becomes very important. In this case, the sufficient statistic can be shown to be non-negative, equaling zero when the empirical estimates agree. Because the sufficient statistic is always positive, the threshold γ must be a positive quantity. Furthermore, to obtain a test that is insensitive to the amount of data, we need to divide the sufficient statistic by N_1+N_2 because of Stein's Lemma [10]. In this way, the threshold will be a constant chosen to reflect a given performance requirement (like false alarm probability).

The test becomes

$$\ln \widehat{\sigma_{1,2}^2} - \frac{N_1}{N_1+N_2}\ln \widehat{\sigma_1^2} - \frac{N_2}{N_1+N_2}\ln \widehat{\sigma_2^2} \gtrless \gamma \tag{11}$$

Another consideration needs to be made before applying this statistical test to thread count data. Each dataset $\mathbf{x}_1$ and $\mathbf{x}_2$ actually consists of two components, the horizontal and vertical thread counts. These need to be handled separately in the variance estimates. Moreover, in the common-distribution estimate, we need to determine which counts from one measurement are best associated with the second set of counts. In other words, weft measurements need to be associated with weft measurements, warp with warp. This requirement amounts to a hidden similarity test, but one that does not declare which measurements are warp or weft. Quite simply, the proper choice amounts to yielding best-case results in terms of the statistical hypothesis test, which minimizes the variance of the combined datasets.

Non-parametric solution: While the parametric solution to testing for thread count similarity between two measurements only requires computation of sample variances, it does rely on the distribution of counts being Gaussian. Most measured distributions, as estimated by thread count histograms, do indeed appear to be unimodal and approximately symmetric. But the Gaussian model is not always correct and detailed examination may well reveal that histograms appearing to resemble a Gaussian are not. Without a model for the count distribution, we turn to non-parametric methods that do not require a model. In particular, a non-parametric hypothesis testing rule derived from an information theoretic standpoint has proven optimality properties and we employ it here.

The main assumption is the horizontal- and vertical-thread counts have each been each summarized with a histogram, thereby portraying which and how often thread count values occur for each measurement over the entire painting. With $\mathbf{X}$ representing a set of measurements, $h_{\mathbf{X}}(k)$ denotes the number of counts from dataset $\mathbf{X}$ having values within the range of the kth bin Δ_k. We used fixed-width bins with $\Delta = 0.1$ threads/cm. The following theory, which shows that the maximum likelihood estimate of a discrete probability distribution is the empirical histogram, does not demand fixed-width bins. Assuming the dataset contains statistically independent, identically distributed components:

$$\Pr[\mathbf{X}] = \prod_{i=1}^{N} \Pr[X_i] \quad \text{or} \quad \log \Pr[\mathbf{X}] = \sum_{i=1}^{N} \log \Pr[X_i]$$

The logarithm of the joint probability can be re-written as a sum over the histogram index rather than the data index to yield [10]

$$\log \Pr[\mathbf{X}] = -N\left[\mathcal{D}\left(\frac{h_{\mathbf{X}}(k)}{N}\middle\| P_k\right) + \mathcal{H}\left(\frac{h_{\mathbf{X}}(k)}{N}\right)\right] \tag{12}$$

$\mathcal{D}(h_{\mathbf{X}}(k)/N\|P_k)$ is the Kullback–Leibler measure of distance [11,12] between the empirical and theoretical probability distributions and $\mathcal{H}(h_{\mathbf{X}}(k)/N)$ is the entropy of the empirical histogram.[6] Note that this expression applies in general; it describes how the probability of a particular dataset depends on the empirical distribution – the histogram (known in information theoretic parlance as the type) – and the true distribution. We can consider this result as a log-likelihood function, with the P_k, the probabilities of X having a value within bin k, taking the role of unknown parameters. To find the maximum likelihood estimate, we want to minimize the quantity in the bracketed quantity in Eq. (12). Because the entropy term is fixed, we need only minimize the Kullback–Leibler term. As the Kullback–Leibler distance term is non-negative, the smallest it can be is zero, which occurs *only* when $P_k = h_{\mathbf{X}}(k)/N$, showing that the histogram is the maximum likelihood estimate of the underlying probability function.

This key result allows us to determine the optimal non-parametric decision rule for deciding if two datasets $\mathbf{X}_1$ and $\mathbf{X}_2$ have the same probability distribution or different ones. We assume the datasets are statistically independent of each other and have statistically independent, identically distributed components. As before, we use the likelihood ratio test.

$$\frac{\Pr_1[\mathbf{X}_1]\Pr_2[\mathbf{X}_2]}{\Pr[\mathbf{X}_1]\Pr[\mathbf{X}_2]} \gtrless \eta$$

Because we do not know the distributions, we use the maximum likelihood estimates of the probabilities in the likelihood ratio. Using these estimates, the probability $\Pr[\mathbf{X}]$ equals $2^{-N\mathcal{H}(h_{\mathbf{X}}(k)/N)}$, which makes the log likelihood ratio

$$\log \frac{\Pr_1[\mathbf{X}_1]\Pr_2[\mathbf{X}_2]}{\Pr[\mathbf{X}_1]\Pr[\mathbf{X}_2]} = -N_1\mathcal{H}(h_{\mathbf{X}_1}(k)/N_1) - N_2\mathcal{H}(h_{\mathbf{X}_2}(k)/N_2) + (N_1+N_2)\mathcal{H}\left(\frac{h_{\mathbf{X}_1}(k)+h_{\mathbf{X}_2}(k)}{N_1+N_2}\right) \tag{13}$$

Because of the convexity of the entropy function, we know that this quantity is positive. Furthermore, the left side of (11) can be seen as a special case of this result for the Gaussian case. The same quantity emerges if we divide by N_1+N_2. The more general result of (13) does not rely on the Gaussian assumption. As described previously, each painting's dataset has two components, the thread counts for the horizontal and vertical threads. We determine which of the two associations of thread count results yields the smallest value of the test statistic, then use that value in the similarity test. We do not attempt to match pairs of paintings that do not pass this histogram similarity test.

4.2. Correlating thread density variations

For those pairs of paintings that pass the thread-count histogram similarity test, the algorithm proceeds by subtracting the average value from the profile and then correlating the resulting thread density deviations-from-average (examples of which are shown in Figs. 6 and 7). The correlation coefficient is rooted in the Cauchy–Schwarz inequality: $|\langle x,y\rangle| \le \|x\| \cdot \|y\|$. The problem we face is that equality, equivalent to the maximal correlation coefficient value of one, occurs when $x \propto y$. We need to demand that maximal correlation occurs when the two quantities – deviations of the thread counts from their

[6] Note that although entirely an empirical measurement, $h_{\mathbf{X}}(k)/N$ is a probability mass function in that it is non-negative and sums to one.

average values – are equal, not just proportional. A simple extension of the bound leads to what might be called the maximal linear correlation coefficient[7]:

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

Now, dividing the inner product by the maximum squared norm yields a value of one only when $x=y$. Note that if a constant is subtracted from each the same result applies: $|\langle x-m,y-m\rangle| \le \max\{\|x-m\|^2,\|y-m\|^2\}$. Removing the average thread count in this way leads to a similarity measure that focuses on the same waveform and amplitude of count deviations. In the following expression, $\overline{w}$ is the average of the two profiles $w_1(m)$, $w_2(m)$ combined. Thus, if the two profile's averages differ, the maximal linear correlation coefficient will be reduced. The resulting maximal linear cross-correlation function is

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

Fig. 8 shows the maximal linear cross-correlation between two paintings by van Gogh. The count maps for these paintings match in weft, the less consistent direction, implying that they were cut side-by-side from the same canvas roll. However, because of the consistency of count maps [7], these two paintings could be separated by an unknown amount of canvas, but must lie within the roll's width.[8] The maximal linear cross-correlation value for this case was 0.61. 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).

To determine a reasonable criterion for the threshold value, the details must be considered. For a majority of paintings by van Gogh, the histogram of thread counts has a decidedly Gaussian appearance. Consequently, we can describe a thread-count profile (defined to be the average of thread counts along the thread direction), as a discrete-time, one-dimensional Gaussian process having some correlation function $R(\ell)$. Thread count profiles for two x-rays are cross-correlated to search for weave matches. The preliminary criterion is that the cross-correlation is greater than a threshold at some lag. Because of analytic difficulties presented by the maximal linear correlation function, we instead consider the more easily analyzed correlation function to establish thresholds.

Let $X_1(n)$, $X_2(n)$ be two zero-mean, stationary, discrete-time Gaussian processes having correlation functions $R_1(\ell)$, $R_2(\ell)$, respectively. The measured cross-correlation coefficient over N samples at zero lag is defined to be

$$\widehat{R} = \frac{\frac{1}{N}\sum_{n=0}^{N-1} X_1(n)X_2(n)}{\sqrt{R_1(0)R_2(0)}}$$

This formula assumes that the variances of the two processes are known. Furthermore, we assume that these variances are equal since demanding histogram matches as a precondition for considering the cross-correlation means nearly equal variances: $R_1(0)=R_2(0)\equiv R(0)$.

Assuming the two processes are statistically independent, the expected value of this empirical cross-correlation is zero. To establish a threshold for correlated thread-density patterns, we seek the probability that this cross-correlation exceeds a threshold despite the two profiles being unrelated. We shall assume the quantity $\sum_n X_1(n)X_2(n)$ is approximately Gaussian, which means we only need to evaluate its variance to estimate the probability the cross-correlation exceeds a threshold:

$$\begin{aligned}\mathsf{E}[\widehat{R}^2] &= \frac{1}{N^2R^2(0)}\mathsf{E}\left[\sum_{m,n=0}^{N-1} X_1(n)X_2(n)X_1(m)X_2(m)\right]\\ &= \frac{1}{N^2R^2(0)}\left(NR^2(0)+2\sum_{\ell=1}^{N-1}(N-\ell)R_1(\ell)R_2(\ell)\right)\end{aligned}$$

The most interesting case has equal correlation functions. That is, the two profiles have the same statistical structure but are statistically independent of each other. After a little simplification that incorporates this assumption, the expression for the variance of the estimated correlation coefficient becomes

$$\mathsf{E}[\widehat{R}^2] = \frac{1}{N}\left[1+2\sum_{\ell=1}^{N-1}\left(1-\frac{\ell}{N}\right)\rho^2(\ell)\right] \tag{14}$$

where $\rho(\ell)$ is the correlation-coefficient function $R(\ell)/R(0)$.

To estimate the variance, we computed the term in brackets for every x-ray in the van Gogh database. In this database, each thread count is marked as representing the warp or weft direction. Profiles for warp- and weft-thread counts were formed and the normalized correlation functions of each computed. These were substituted into Eq. (14) and histograms formed of these values (see Fig. 9). The histogram of the weft thread pattern's term tends to be broader than that for the warp. The variance expression (14) will increase according to the width of the normalized correlation function term $\rho(\ell)$. The wider the correlation function (a greater low-frequency content in the power spectrum), the larger the variance. Larger weft-thread count variances result from the lower-frequency nature of weft-thread count variations in comparison to warp-thread count variations. Because of histogram spread, no typical value portrays the behavior for either warp or weft threads.

With a Gaussian model, the probability an empirical cross-correlation between two statistically independent profiles exceeding two standard deviations is very low: $\Pr[\widehat{R} > 2\sigma] = 0.023$. This rule-of-thumb can be used to determine a threshold the *maximal linear* cross-correlation must exceed to declare a thread-count pattern match even though we used the usual cross-correlation function in the analysis. We found that the threshold established this way worked well. From Eq. (14), this threshold depends on two quantities, the value of the bracketed term (plotted in Fig. 9) and N, the amount of

[7] The maximal correlation coefficient between two random variables X and Y is defined as the maximal value of the covariance $K(\phi(X),\eta(Y))$ with respect to all possible functions $\phi(\cdot)$, $\eta(\cdot)$ [13].

[8] Indeed, a painting has been found that has a weft weave match with both of these and fits between them on the putative canvas roll [14, pp. 164–165].

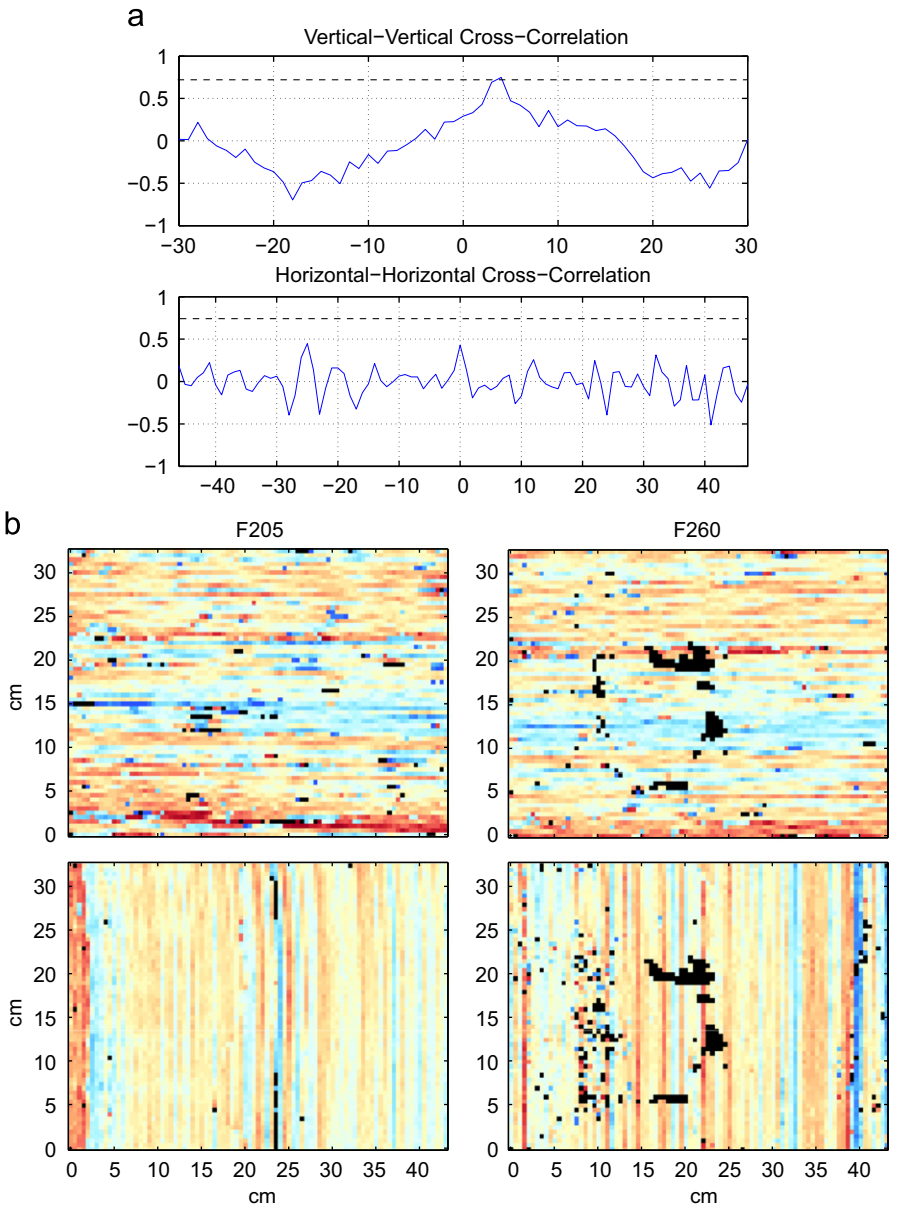

Fig. 8. Panel (a) shows the cross-correlation functions between the horizontal and vertical thread-count profiles for van Gogh's paintings F205 and F260. A clear correlation peak is evident for the vertical (weft) threads. The threshold for peak correlation for warp matches is indicated by the horizontal dashed line. Panel (b) shows both painting's count maps, rotated counter-clockwise to canvas coordinates with the weft maps aligned according to the correlation function peak. The same color scale is the same as that shown in Fig. 6. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this article.)

canvas overlap expressed in multiples of the spatial sampling grid (here 0.5 cm). Because of the larger value of the bracketed term for weft in (14), declaring a weft-thread count pattern match is more demanding on the data. Either the correlation function must exceed a larger threshold for the same amount of canvas overlap or more canvas overlap must be required for the same threshold value. This analysis indicates why weft matches tend to be more difficult to discern (many false-positives occur).

5. Summary

Signal processing has proven to play an important role in helping to date paintings, providing a better understanding of the sequence of artists' production and, to a lesser degree, determining authenticity. Whereas the signal processing tools described in [15] consider colors and brushstrokes evident at the paint surface, the algorithms outlined here help to fingerprint the different types of canvas picture supports used. As the size of the database increases to include

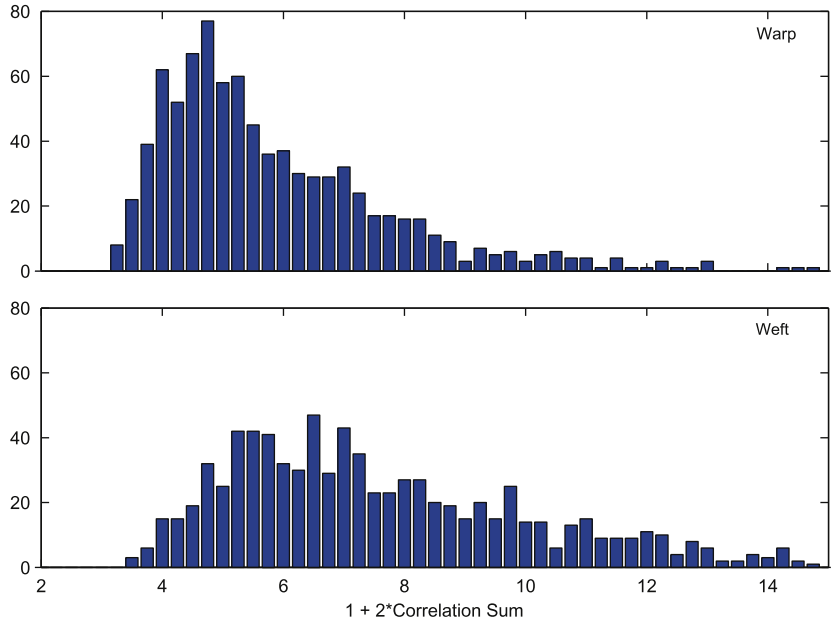

Fig. 9. The empirical values of the bracketed term in Eq. (14) for warp and weft thread density variations are shown for the entire database of over 400 paintings by van Gogh. Implicit in this calculation is the spatial sampling grid, here 0.5 cm. While these histograms have significant overlap, they reveal that the weft-thread density patterns tend to be more variable than those of the weft-direction threads.

artists' works spanning four centuries, we are characterizing the variety of weave and thread density patterns found in artist canvas. To date, over 600 paintings consisting of about 4000 radiographs have been examined. These patterns have different spectra [16], most of which are dominated by horizontal and vertical peaks as demonstrated by our theory. We can extend our spectral theory to describe rarer canvas weave patterns employed by master artists (such as herringbone and diamond).

The spectrum-based weave estimation and thread-count-map 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 understand the artist's process. Considered together, chemical and physical measurements coupled with signal processing provide valuable new insights into the artist's technical and creative processes, complementing the information gained from the study of historical sources [4].

Acknowledgments

The authors thank the Van Gogh Museum (Amsterdam) for partial support of this research and for allowing reproduction of the drawing and the radiographs. Travel grants from the National Science Foundation (CCF-1048344 to DHJ, CCF-1048352 to CRJ) also supported the research.

References

[1] K. Lister, C. Peres, I. Fiedler, Appendix: tracing an interaction: supporting evidence, experimental grounds, in: D. Druick, P. Zegers (Eds.), Van Gogh and Gauguin: The Studio of the South, Thames & Hudson, 2001, pp. 354–369.
[2] E. van de Wetering, Rembrandt: The Painter at Work, Amsterdam University Press, Amsterdam, 1997.
[3] A. Kirsh, R. Levenson, Seeing through Paintings: Physical Examination in Art Historical Studies, Yale University Press, 2000.
[4] L. van Tilborgh, T. Meedendorp, E. Hendriks, D.H. Johnson, C.R. Johnson Jr., R. Erdmann, Weave matching and dating of van Gogh's paintings: an interdisciplinary approach, Burlington Magazine 154 (2012) 112–122.
[5] E. Hendriks, D.H. Johnson, C.R. Johnson, Jr., Interpreting canvas weave matches, Art Matters, to appear.
[6] R. Mersereau, D. Dudgeon, Multidimensional Digital Signal Processing, Prentice-Hall, 1984.
[7] C.R. Johnson Jr., D.H. Johnson, N. Hamashima, H. Yang, E. Hendriks, On the utility of spectral-maximum-based automated thread counting from x-rays of paintings on canvas, Studies in Conservation 156 (2012) 104–114.
[8] D.H. Johnson, L. Sun, C.R. Johnson, Jr., E. Hendriks, Matching canvas weave patterns from processing x-ray images of master paintings, in: Proceedings of the ICASSP, 2010, pp. 958–961.
[9] M. Gutman, Asymptotically optimal classification for multiple tests with empirically observed statistics, IEEE Transactions on Information Technology 35 (1989) 401–408.
[10] T.M. Cover, J.A. Thomas, Elements of Information Theory, 2nd ed. John Wiley & Sons, Inc., New York, 2006.
[11] S. Kullback, R.A. Leibler, On information and sufficiency, Annals of Mathematical Statistics 22 (1951) 79–86.
[12] S. Sinanović, D.H. Johnson, Toward a theory of information processing, Signal Processing 87 (2007) 1326–1344.
[13] H. Gebelein, Das statistische problem der korrelation als variations- und eigenwertproblem, Zeitschrift für Angewandte Mathematik und Mechanik 21 (1941) 364–379.
[14] E. Hendriks, L. van Tilborgh, V. van Gogh, Paintings: Antwerp and Paris (1885–1888), van Gogh Museum, Amsterdam and Zwolle, 2011.
[15] C.R. Johnson Jr., E. Hendriks, I. Berezhnoy, E. Brevdo, S. Hughes, I. Daubechies, J. Li, E. Postma, J. Wang, Image processing for artist identification, Signal Processing Magazine 25 (2008) 37–48.
[16] J. Escofet, M. Millán, M. Ralló, Modeling of woven fabric structures based on Fourier image analysis, Applied Optics 40 (2001) 6170–6176.
[17] J.-B. de la Faille, The Works of Vincent van Gogh: His Paintings and Drawings, Meulenhoff, Amsterdam, 1970.
[18] J. Hulsker, The New Complete Van Gogh: Paintings, Drawings, Sketches: Revised and Enlarged Edition of the Catalogue Raisonné of the Works of Vincent van Gogh, Meulenhoff, 1996.