

A THREAD COUNTING ALGORITHM FOR ART FORENSICS*

Don H. Johnson

Elec. & Comp. Engineering
Rice University

William A. Sethares

Elec. & Comp. Engineering
University of Wisconsin

C. Richard Johnson, Jr.

Elec. & Comp. Engineering
Cornell University

H. Lee

Elec. & Comp. Engineering
University of Wisconsin

Andrew G. Klein

Elec. & Comp. Engineering
Worcester Polytechnic Institute

Ella Hendriks

Conservation Department
Van Gogh Museum

ABSTRACT

Thread counting algorithms seek to determine the vertical and horizontal weave densities of the canvas underlying a painting. X-ray images of a painting serve as the raw data. Canvas texture can be modeled as a sum of two sinusoids having nearly orthogonal spatial frequencies. The actual painting serves as an additive signal that only distracts from the thread-counting process. The algorithm employs a variant of short-time Fourier analysis to the image domain. In simple cases, two-dimensional Fourier transforms of small areas reveal isolated peaks at the proper vertical and horizontal frequencies. In more challenging cases, even within the same painting, additional peaks appear in the expected frequency range, requiring heuristics to extract from the spectral measurements representative thread-count estimates for a painting. The resulting algorithm, while still under development, can analyze weave density for entire paintings with an accuracy comparable to human measurements.

1. THE CANVAS THREAD COUNTING PROBLEM

To investigate the history and authenticity of paintings by the great masters, signal processing algorithms can provide new insights [1,2]. Our focus here is on x-ray images that can reveal much about what's below the visible surface [3,4]. The greater the radiographic-absorbing paint thickness along the beam, the greater the opacity, meaning that x-ray image intensity variations correspond to paint thickness variations. Figure 1 shows an x-ray taken of a painting by Vincent van Gogh. At the proper scale, when the canvas weave is made visible by the thicker ground layer of lead-white paint in the grooves between canvas threads as in Figures 2a and 2c, the thread density can be determined. Thread count data is commonly used as evidence for dating, linking pictures from the same canvas roll, and attribution [3–5]. Considering how a loom works reveals what detail can be discerned from thread count measurements. The vertical threads mounted in 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 may correspond to either the vertical or horizontal threads in the painting. Thread count spread across a painting 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 [4, p. 100]. In addition, paintings made from the same canvas rolls may not have been made on pieces cut the same way: the warp may be horizontal in one painting and vertical in another. Thread counts along with related forensic data allow the art historian to pose strong hypotheses about how the canvas roll was used for paintings contemporary with each other.

Figure 2 shows typical test swatches taken from the x-ray image of the painting of study here. The vertical and horizontal weave patterns can be clearly seen. Thread counting algorithms seek the weave density, measured in threads/cm, in both the horizontal and vertical directions within a swatch and to study how these counts vary throughout the painting. The current standard 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, 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. By using signal processing techniques, the hope is to develop algorithms that are “better” than manual measurements, but also allow whole painting analysis, which would allow quantification of an entire painting's canvas weave pattern and thus permit a more detailed analysis of temporal painting sequences and painting orientation on the canvas roll.

The Thread Count Automation Project seeks to develop signal processing algorithms that can detail the variations in

***Acknowledgement.** This work represents a collaboration between the Thread Count Automation Project directed by Professor Rick Johnson at Cornell University and the Van Gogh Museum, Amsterdam.



Fig. 1. X-ray image of van Gogh’s *Portrait of an Old Man with a Beard* (F205/JH971). The painting, the wood stretcher (the lighter border) and canvas-mounting nails can be clearly seen. Magnifying the x-ray reveals the canvas weave as well (see Figure 2). Image sampled at 600 dpi and provided by the Van Gogh Museum.

the canvas thread density across a painting [6]. 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 wooden stretcher with nails. The hypothesis is that paintings from the same canvas roll should possess common thread count (density) variations horizontally and vertically. In this way, works associated with the same canvas roll can be presumed to have been painted at about the same time. The thread used to make a canvas is transparent to x-rays. Fortunately, artists typically prepared their canvases with a white undercoat (also called a ground layer) to smooth the surface. The small variations in undercoat thickness filling the valleys of the canvas weave lead to variations in x-ray opacity that can be measured. Painstaking counting of the number of threads in short vertical and horizontal stripes at several locations in the x-ray *by humans* is the current standard. The Project seeks to produce reliable, accurate techniques that will allow entire paintings to be analyzed with minimal human intervention. In an internal competition among Project members, the algorithm described here achieved the best performance [7].

2. MODEL OF THE CANVAS WEAVE PATTERN

A simple model for the canvas weave leads to a simple processing algorithm under ideal conditions. Examining the

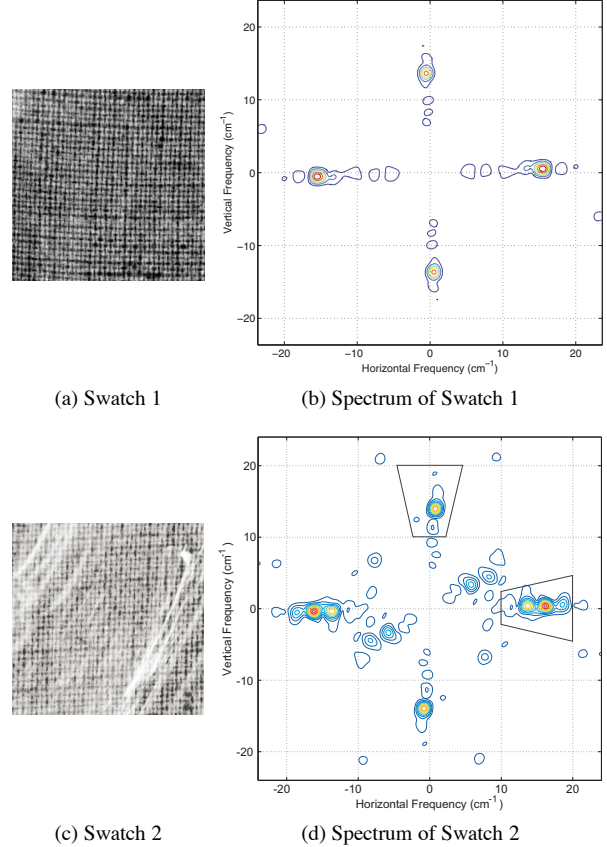


Fig. 2. Panels a and c show 1" × 1" swatches taken from the x-ray of F205 shown in Figure 1. Panel (a) is taken from the area in front of the face; panel (c) is taken from the area below the ear. Panels b and d show detailed spectra computed from a smaller square (1/2" × 1/2") located in the upper left corner of each swatch. The wedges indicate areas where weave-related spectral peaks are found.

canvas-related patterns in x-ray images showed that representing the canvas weave as a sum of sinusoids in the x and y directions can serve as a baseline model.*

$$c(\vec{x}) = p(\vec{x}) + \frac{A}{2} \cdot (2 + a_h \sin(2\pi \vec{f}_h \cdot \vec{x} + \phi_h) + a_v \sin(2\pi \vec{f}_v \cdot \vec{x} + \phi_v)) \quad (1)$$

The quantity $c(\vec{x})$ is the canvas x-ray image that depends on the 2D spatial variable \vec{x} ; $p(\vec{x})$ represents the artist’s contribution (the painting) to the x-ray. The remainder of the model describes the canvas weave. The constants A , a_h and a_v determine the average intensity and the amplitudes of the horizontal and vertical weave. Because the image must be non-negative, the amplitudes a_h and a_v must be less than one. They and the phase terms are ignored, the focus being the weave densities, which are conveyed by \vec{f}_h and \vec{f}_v , the vector frequencies corresponding to the vertical and horizontal thread counts, respectively. In our model, the *vertical* threads

*The product-of-sinusoids leads to very similar spectra. The simpler sum model captures all the features needed to determine thread counts.

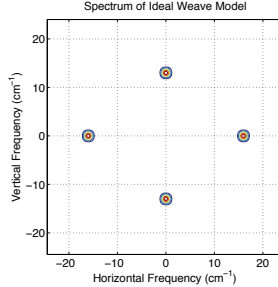


Fig. 3. The panel shows the spectrum corresponding to the weave model expressed by equation (1) when computed when $\vec{f}_h = 16\vec{i}_x \text{ cm}^{-1}$ and $\vec{f}_v = 13\vec{i}_y \text{ cm}^{-1}$. The average was subtracted before computing the spectrum, thereby removing the peak at the origin.

create oscillations of x-ray intensity in the x direction, which leads to a *horizontal* frequency component; a similar idea applies to the horizontal threads and the vertical frequency.

Shown in Figure 3 is the spectrum resulting from the model. The predicted spectrum consists of impulses at the origin (removed) and at the horizontal and vertical frequencies $\pm\vec{f}_h$ and $\pm\vec{f}_v$, respectively. This spectrum is somewhat simpler than that predicted for modern fabric weave [8,9], indicating that the model could possibly be improved. Be that as it may, clearly identifiable peaks at the thread frequencies dominate the spectrum, and help interpret the spectra shown in Figure 2. Swatch 1’s spectrum shows similar peaks along with other variations that don’t disturb obtaining thread counts from the dominant spectral peaks. Swatch 2’s spectrum is shows weave-related peaks along with spectral variations off-axis and on-axis. Substantially off-axis portions are due to the painting itself and can be ignored. On-axis peaks (visible to the right and left of the weave peak on the positive horizontal frequency axis) mean that heuristics need to be employed as a post-processing stage. Assuming the heuristics work well, a 2D spectral approach will clearly produce a clearer frequency estimate than measuring the number of peaks in the horizontal and vertical directions, which is the current benchmark approach. Furthermore, there is no guarantee that the weave will align with the x and y axes. The sampling grid used to digitize the x-ray will not align with the canvas weave and canvas mounting artifacts such as canvas stretch, which can slant the weave near the painting’s edges, rotate the frequency vectors from the axes. These effects make space-domain peak counting algorithms hard to implement. Spectrally, the weave peaks will be rotated by the same angles as the warp and weft. Consequently, the frequencies \vec{f}_h and \vec{f}_v must be vector valued, pointing in the direction of local weave variations. We seek measurements of the frequency *magnitudes* that are unaffected by these distortions: $|\vec{f}_h|$ corresponds to the vertical thread count and $|\vec{f}_v|$ to the horizontal thread count.

3. SHORT-SPACE SPECTRAL ANALYSIS

“Short-space” Fourier analysis was chosen since the artist’s renderings creates in the x-ray what amounts to a non-

stationary *additive* pattern to the canvas weave. The additivity arises because it is the *integral* along the x-ray beam of an object’s opacity that determines the x-ray image. Since paints of different colors have nearly the same opacity, the thickness of the paint determines x-ray image intensity. Thus, a well-imaged x-ray of a painting should show approximately sinusoidal canvas-related variations throughout a painting, whether the artist paints thinly or uses impasto.[†] Weave stretching varies across a painting, further reinforcing local rather than global spectral estimates. The most important deviation from the ideal model is the painting itself, a phenomenon that is difficult to describe mathematically. When the painting has “out of band” components, they are simply ignored. In-band components pose a more serious challenge. In our experience, the spectral approach is very robust and reliable thread-count estimates could be produced for nearly all of the painting surface.

The “short-space” spectral analysis algorithm mimics the familiar short-time Fourier analysis used in analyzing non-stationary time series (e.g., speech) [10]. A square section is extracted from the x-ray, the average subtracted and a window applied. To obtain spectral detail, the 2D Fourier transform size is much larger than the windowed section. Spectral resolution is determined by the window size and shape, but accuracy in locating spectral peaks is determined by the transform length (about every 0.1 thread/cm for the data shown here). The swatch shown in Figure 2a represents a simple case: the associated spectrum in Figure 2b evidences four clear peaks evidenced by the sinusoidal canvas weave pattern of equation (1). Note that the peaks do not lie precisely on the axes; they are rotated slightly, indicating that the sampling grid was not precisely aligned with the weaving. Since such precise alignment is impossible, the algorithm searches for spectral peaks within a radius range known for canvas thread counts (between 10 and 20 threads/cm) within small wedges spanning $\pm 20^\circ$ centered on the axes as shown in Figure 2d. In this easy case, only one peak in the expected spectral range occurs, yielding an unambiguous estimates of vector frequency magnitude in both the horizontal and vertical directions.

A second swatch (Figure 2c), taken from the same painting as the first swatch, shows clear signs of van Gogh’s artistry. For the thread-counting problem, these brush strokes amount to *additive* interference. Consequently, interference from the artist’s brush strokes can theoretically be removed by *linear* filtering. The spectrum (Figure 2d) of the second swatch shows a much more complicated structure than that of the first. Energy slanting through the frequency origin corresponds to the slanted brush strokes but it doesn’t interfere with the spectral peaks attributed to the canvas weave. Be-

[†]Several x-rays do not satisfy the criterion of being “well-imaged.” In the digitization process, clipping sometimes occurs in heavily painted areas, which means the small intensity variations due to the canvas weave are absent. Automatic thread counting algorithms must allow for the possibility that *no* thread counts can be measured in some areas of the painting.

cause we search only within the expected frequency range, energy in other frequency regions does not affect our estimates. Interference only becomes an issue when it lies in the same spectral region as the weave. Because of the swatch's small size, spectral spread from non-canvas related portions of the spectrum can interfere with the spectral peaks we seek. Consequently, the x-rays are high-pass filtered with a zero-phase FIR filter to mitigate such effects before the Fourier spectrum is calculated.

However, in the region of interest, more than one spectral peak sometimes occurs. Figure 2d shows an example. We are unsure why these interfering peaks are present. They might indicate a low-frequency modulation of the weave pattern. When multiple peaks in the frequency ranges of interest occur, we resorted to heuristics that consider what the spectra are in other areas of the painting, which are discussed subsequently in the context of whole painting analysis. Also note that the peaks again do not lie on the frequency axis, but they aren't rotated as they are in the first example. Instead, the spectral peak locations indicate a stretching of the canvas that warps the weave slightly, angling the threads in both directions. This stretching can be seen clearly in Figure 2c. We found that rotation and/or stretching to be a common occurrence in the paintings we analyzed. Consequently, the thread count is determined by the *radius* of the spectral peak and the angle of the peak providing an estimate of the stretching direction. This stretching justifies searching in angular wedges straddling the frequency axes shown in Figure 2d.

4. WHOLE PAINTING ANALYSIS

Short-space Fourier analysis can be applied to an entire painting, not only to obtain thread count measurements across the entire area, but also to facilitate spectral peak extraction when the individual spectra yield complicated results. As described earlier, relatively small square swatches ($1/2" \times 1/2"$) are extracted, windowed with a separable raised-cosine window and a spectrum for each computed. Swatches overlap each other by half in each direction, meaning we are sampling the short-space spectrum every $1/4"$ in both directions. When spectra have several peaks close to each other in the same two-dimensional frequency area, we rely on those nearby spectra having one dominant peak to serve as consistency benchmarks. Consistency is judged by computing the median (to lessen the effect of outliers) and only considering those spectra whose peak locations lie within a radius determined by the average of absolute errors. We throw out frequency estimates that occur too far from the median frequency, declaring that no estimate could be obtained for the corresponding swatch. We then return to the multiple-peak cases, accepting a single peak from each spectrum (in both the horizontal and vertical frequency axes) that is the closest to the single-peak-median and close enough. Calculations were made in Matlab and took about three hours to analyze F205 on a laptop computer.

Viewing the results across the entire painting (Figure 4)

provides the information we seek about the canvas. First of all, the largest spread of measurements occurs in the vertical direction, implying it is the weft direction. In more detail, patterns of thread count variations emerge in both deviation plots. The horizontal thread counts are much smaller than the vertical thread count. The horizontal thread counts show two areas of largest deviation from the average: across the top and in the middle. The vertical thread counts show the largest deviations down the left side (positive deviation) and just to the left of the center (a negative deviation). A less prominent positive deviation lies to the right of center. These characteristics can be used as a kind of fingerprint for the canvas underlying the painting.

To illustrate how whole-painting thread counts can be used forensically, we compared the thread count analysis for F205/JH971 (Figure 1 and for the painting *Backyards of Old Houses in Antwerp in the Snow* (F260/JH970) known to have been painted the same month. Figure 5 shows their thread count deviation maps. The results indicate that these paintings were likely cut from the same canvas roll, sharing the more variable weft direction. We are working to develop a quantitative measure of similarity.

5. CONCLUSIONS

Spectral techniques for measuring thread counts in artist canvases offers clear advantages over manual, space-domain methods. Typical problems affecting human thread counting are artist's marks on a larger scale than the thread width, contrast-compressing ground thickness/opacity, tilt and cusping of the canvas. Interference from the artist's brush strokes that lie outside a narrow range of frequencies have no effect on the measurements made by locating peaks. Tilt and cusping are compensated by considering the magnitudes of the horizontal and vertical frequency vectors. Even so, artifacts do occur that demand a more detailed and extensive analysis than has been performed to date. Heuristics will probably always be needed to reconcile difficult cases.

Whole-painting analysis from two paintings known to be cut from the same roll should show nearly identical spatial variations in thread count over the spatial extent of the paintings. Consequently, we believe our "short-space" spectral analysis approach can provide more detailed and extensive measurements to strengthen the forensic evidence for sequencing paintings.

6. 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] A. Pelagotti, A. Del Mastio, A. De Rosa, and A. Piva, "Multispectral imaging of paintings," *Signal Processing Magazine*, vol. 25, pp. 25–36, July 2008.

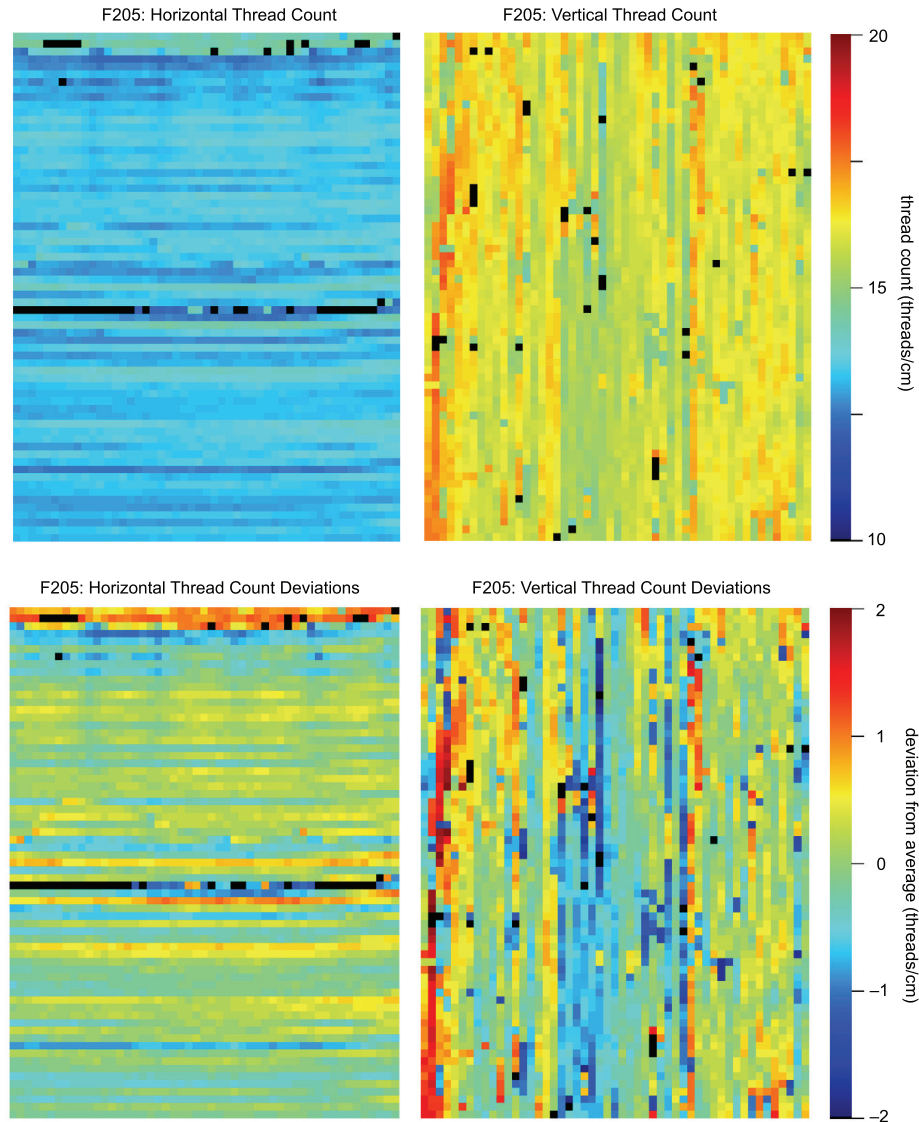


Fig. 4. Thread counts in threads/cm for the van Gogh painting shown in Figure 1. Each square corresponds to a spectrum sampled every 1/4" across the surface in both x and y . Black pixels indicate where the algorithm made no thread-count estimate. The top row shows as a heat map the measurements in threads/cm for the horizontal and vertical thread counts. The bottom row showed how these measurements deviated from their respective averages: 13.33 threads/cm horizontally and 16.01 threads/cm vertically.

- [3] 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.
- [4] E. van de Wetering, *Rembrandt: The Painter at Work*, Amsterdam University Press, Amsterdam, 1997.
- [5] A. Kirsh and R.S. Levenson, *Seeing through Paintings: Physical Examination in Art Historical Studies*, Yale University Press, 2000.
- [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] J.S. Ng, "A common platform for thread count automation," Design Project Report, Cornell University, 2008.
- [8] E.J. Wood, "Applying Fourier and associated transforms to characterization in textiles," *Textile Research Journal*, vol. 60, pp. 212–220, 1990.
- [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.
- [10] L.R. Rabiner and R.W. Schafer, *Digital Signal Processing of Speech Signals*, Prentice Hall, Englewood Cliffs, NJ, 1978.

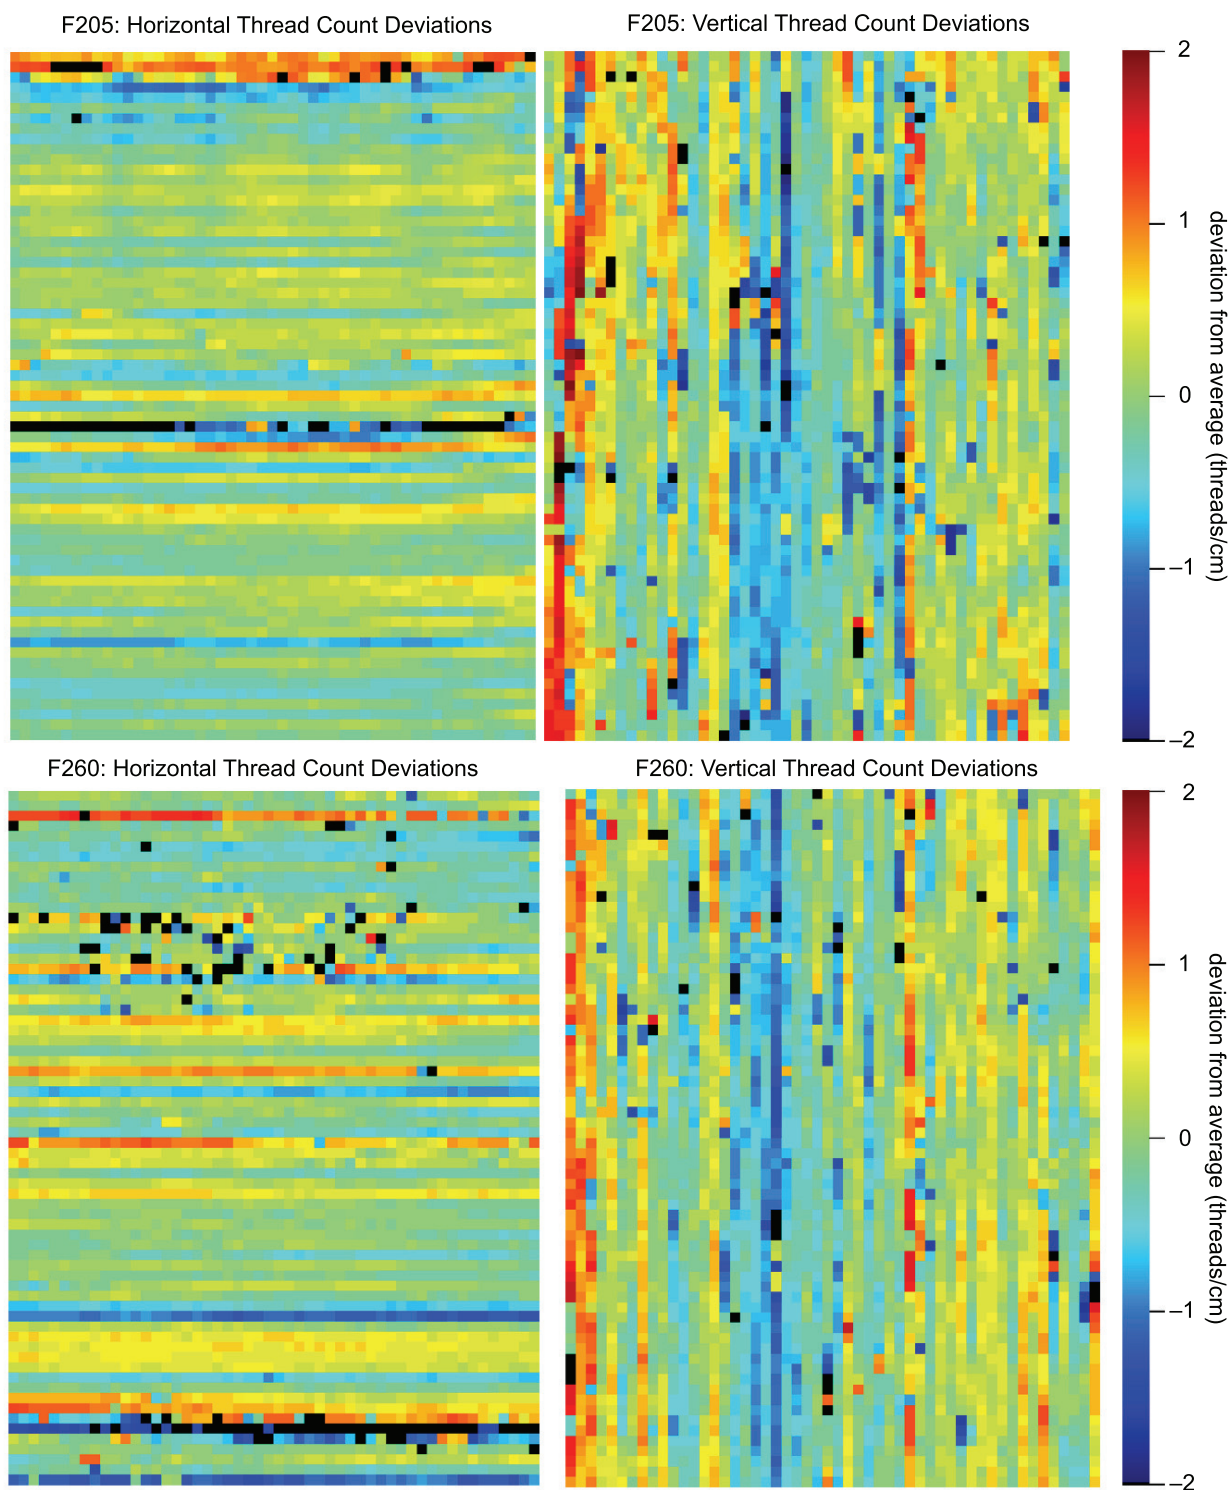


Fig. 5. Comparison for thread count deviations from average for two paintings (F205 shown in Figure 1 and F260) known to have been painted from the same canvas roll. Same display format as in Figure 4. Average thread counts for F260: 13.18 threads/cm horizontally and 15.73 threads/cm vertically. The vertical thread count maps have been shifted horizontally to maximize the alignment of the maps for the two paintings.