Reuniting Poussin's Bacchanals Painted for Cardinal Richelieu through Quantitative Canvas Weave Analysis

Robert G. Erdmann*[1], C. Richard Johnson[2], Mary Schafer[3], John Twilley[3] and Travis Sawyer[4]

[1]Department of Materials Science and Engineering, University of Arizona
[2]School of Electrical and Computer Engineering, Cornell University
[3]Nelson-Atkins Museum of Art
[4]College of Optical Sciences, University of Arizona

Abstract

Art historical debate has percolated for more than half a century over the identity of the three Bacchanals (1635-6) commissioned from Nicolas Poussin by Cardinal Richelieu for his chateau in Poitou, France. The subjects have been identified as the Triumphs of Pan, Bacchus, and Silenus, respectively. The success of this commission led to demand for the production of copies close in date to the originals, some of which are represented in major collections. There has been general consensus that *Triumph of Pan* (National Gallery of London, NG6477) was an original part of that commission. However, the status of *Triumph of Bacchus* (Nelson-Atkins Museum of Art, 31-94) as the original commissioned version has at various times been doubted, in spite of sharing a common provenance with *Pan* until 1850. An incomplete early provenance and apparent stylistic discrepancies, particularly in relation to *Pan*, have caused serious doubts about the inclusion of *Silenus* (National Gallery of London, NG42) among Richelieu's original commission. It is today classified by the National Gallery as a painting "after Poussin".

In preparation for the Nelson-Atkins' catalog of its French painting collection we sought to employ recent innovations in computerized "thread counting" for automated analysis of canvas weave variations to resolve the status of *Bacchus*. This approach has proven particularly powerful in demonstrating commonalities and distinctions between canvases of nominally identical average thread counts with art historical implications, as among the paintings of Vincent van Gogh (*The Burlington Magazine*, February 2012), court portraits by Velazquez (*The Burlington Magazine*, September 2012), and paintings by Johannes Vermeer (*Metropolitan Museum Journal*, 2012).

Automated weave comparison held the potential to settle the questions surrounding *Bacchus* in the event that the artist had employed part of the same bolt of canvas for *Pan*. Existing National Gallery radiographs of *Pan* were digitized and made available for our comparison with *Bacchus*

*Corresponding author. erdmann@arizona.edu

through this process, resulting in a very close match of their warp thread-spacing variations (resembling a bar code when color-coded and mapped). When subsequently performed on radiographs provided for the National Gallery *Triumph of Silenus*, its warp spacing variations were also matched to a very high degree with those of *Pan*, providing compelling evidence that all three paintings were executed on the same bolt of canvas. This form of evidence, connecting the three works to a single canvas as it does, relates them more closely than other forms of analysis might, such as a demonstration of the shared use of a common set of pigments, and is unaffected by variations in condition of the paintings. This outcome should lead to a reassessment of the relationship between *Silenus* and the other paintings as well as the workshop practices of Poussin whose engagement of assistants is unclear. A comparison of painting materials and working methods employed for *Bacchus* and *Pan* is now proceeding, which will be underpinned by the certainty that they were contemporary products of the same studio.

1 INTRODUCTION AND BACKGROUND

In 1635, Nicolas Poussin received a prestigious commission from Armand-Jean du Plessis, Cardinal de Richelieu, to paint a series of bacchanal scenes devoted to Bacchus, Pan, and Silenus. The series is known from the *Triumph of Bacchus* (31-94, Nelson-Atkins Museum of Art), *Triumph of Pan* (6477, National Gallery of Art, London), and *Triumph of Silenus* (42, National Gallery of Art, London), although the attribution of the first and third of these as the original versions by the hand of Poussin has been the subject of debate.[1] The series was destined for Richelieu's chateau in Poitou, France, where it was installed in his Cabinet du Roi alongside earlier works by Perugino, Mantegna, and Costa.

As early as 1665, a number of high-quality copies were produced of the successful series. Although little is known about the copyists, Hugh Brigstocke (1994) points out that Poussin probably did not make the copies, since he was more likely to have improved upon or slightly modified a successful composition rather than faithfully replicate it, as was the case with his numerous variants of the popular holy family subject. By the mid-eighteenth century, Poussin's original paintings at the château de Richelieu were themselves replaced by a set of copies now located at the Musée des Beaux-Arts, Tours. For *Bacchus* alone, at least seven copies exist today. Over the years, the presence of so many seventeenth- and eighteenth-century copies gave rise to complicated attribution questions for a number of Poussin paintings, including the Richelieu Bacchanal series.

The National Gallery classifies *Silenus* as an excellent early copy, referencing the fact that certain passages lack the level of quality displayed by *Pan*, and that its provenance can be traced no earlier than 1824. In addition, there has been a general belief that Poussin, described as a perfectionist who "neglected nothing," worked independently and would not have risked engaging studio assistants for such an important commission and patron (Wine 2001). This belief has lead many art historians to rule out the possibility that, as an explanation for perceived differences in quality, other hands may have been involved in the creation of the painting. Opinion on the attribution of *Silenus* is somewhat divided among Poussin scholars in the recent literature; most art historians regard it as a copy, while others allow for the possibility that Poussin may have worked on the Richelieu commission with studio assistants.

Although the authenticity of the National Gallery's *Pan* is unchallenged today, it too was once deemed a copy. A painting in the Louvre, now firmly classified as a copy of *Pan*, was considered the original version at the time of an exhibition held at the Petit Palais in 1925. The National Gallery painting was exhibited as the autograph version in 1960 and became widely accepted following an exhibition at the National Gallery of Scotland in Edinburgh in 1981, which included a reunion of the Richelieu Bacchanal series (Wine 2001, 364n62).

Doubt was initially cast on the Nelson-Atkins painting during the 1925 exhibition mentioned above, and Anthony Blunt expanded on these concerns in his 1966 catalogue raisonné, publishing for the first time the attribution of *Bacchus* as a copy: "It is cold and mechanical in handling, and has nothing of the delicacy and sensitiveness of the [National Gallery] picture" (Blunt 1966, 98). In addition to overlooking the fact that *Pan* and *Bacchus* shared an identical provenance until 1850, Blunt implied that when the paintings left the Richelieu family in the mid-eighteenth century, a dealer replaced the original *Bacchus* with a copy in order to profit from two sets of bacchanals, each containing an original Poussin.

Pierre Rosenberg called for the reinstatement of *Bacchus* as the original in 1977, but it was not until the 1981 Edinburgh exhibition that the painting was reassessed and accepted by a majority of scholars as autograph. The entry for the painting in the exhibition catalog addresses the inconsistency of certain passages and the authenticity question: "On balance it seems likely that *The Triumph of Bacchus* is an original. It is difficult to imagine any copyist achieving the vigour of the centaur holding a torch and flowers, the warmth of the landscape background under a golden sky, or the subtle and lively handling of the ivy, vine leaves, and other foliage which surrounds and decorates the foreground figures. At the same time, condition alone cannot explain the extremely poor modeling of certain figures, most notably the putto in the left foreground and the dancing female, seen from behind, at the extreme right" (National Gallery of Scotland 1981, 52). Blunt (1982) eventually reversed his earlier opinion but proposed that, despite Poussin's apparent standard practice, studio assistants were involved in the making of both *Pan* and *Bacchus*. Since 1981, most Poussin scholars - Pierre Rosenberg, Hugh Brigstocke, and Doris Wild to name a few - accept the Nelson-Atkins painting as original, but questions persist in the literature. In 1994, Jacques Thuillier (18) described *Bacchus* as a copy that "has found defenders." Given that the three paintings in the Richelieu commission were executed at roughly the same time and intended for the same location, he strongly urged a comparative study of the canvas, preparatory layers, and brushwork in anticipation of new findings that might settle the ongoing debate.

In preparation for a scholarly catalog of its French paintings collection, the Nelson-Atkins has undertaken a series of scientific investigations in support of curatorial and conservation questions about certain works. The Nelson-Atkins partnered with the Thread Count Automation Project (TCAP), initially to make a comparison of the canvas of *Bacchus* with that of *Pan*, for which the National Gallery had existing radiographs. In the event that a direct correlation could be demonstrated between the canvas of *Bacchus* and that of *Pan*, whose authenticity is no longer questioned, lingering concerns about the Nelson-Atkins painting could be assuaged in a way that mere similarity of painting materials and technique would be unlikely to do. Since the weave comparison is inherently independent of the technique of painting or differences in the condition of the works, the complications that might arise for other forms of technical study due to surface losses or overpainting are overcome in this approach. In the event of a direct match between the canvases, questions of connoisseurship related to the use of studio assistants or the role of a copyist would take on a fundamentally different character. The weave analysis posed the additional advantage of utilizing existing radiographs without requiring significant time from museum staff or a disruptive period with the paintings off view. In the event of a lack of correlation between the canvases, other means of comparison would have been required whose completion might not have been feasible in advance of the catalog deadline. Based upon the success of the method described below, and with further cooperation of the National Gallery, the canvas of the third painting in the series, the strongly doubted *Triumph of Silenus*, was added to the study.

Radiographs of the three Richelieu Bacchanal paintings are fairly similar. All are dominated by the canvas weave and reveal low contrast among the compositional forms. Although all three paintings are lined, and the original canvases are thus hidden from inspection, the radiographs confirm that the original canvas is a coarsely-woven, plain weave fabric with variations in thread thickness. The canvases are primed with a double ground, and based on the published technical notes for the National Gallery paintings, the ground colors on *Bacchus* (lower reddish-brown and upper beige) appear to be comparable to those of *Pan* (lower reddish-brown and upper buff-color) and *Silenus* (lower red-brown ochre and upper buff-color) (Wine 2001) [2]. The minimal contrast of the radiographs can be attributed to the added density of the double ground and the artist's thin paint applications. Whereas the low contrast and prominent canvas weave interfere with the identification of artist changes and indications of painting technique, these features provide ideal conditions for computer-automated canvas weave analysis.

2 AUTOMATED AND MANUAL CANVAS ANALYSIS

Consider the set of average thread counts taken in evaluation squares surrounding test points on a grid covering an entire painting. The result was first visualized in (Johnson, Hendricks, et al., 2009) using colors to represent thread counts above (redder) and below (bluer) the average for the painting. Stripes of color appeared due to the weaving process with bundles of adjacent threads retaining their relative spacing as they traverse the canvas. Matching these striped patterns reveals original canvas rollmate candidates (Johnson, Hendricks, and Johnson, 2013). In combination with other knowledge of the artist's studio practice and stylistic development, such weave matches provide evidence useful in addressing art historical issues in studies of the paintings of Van Gogh (van Tilborgh et al., 2012), Velazques (Perez d'Ors, Johnson, and Johnson, 2012), and Vermeer (Liedtke, Johnson, and Johnson, 2012).

Previous efforts at automating whole-painting canvas analysis have utilized a Fourier transform approach for the extraction of thread density from small square swatches of the radiograph (Johnson, Erdmann, and Johnson 2011) (Johnson, Johnson, and Erdmann, 2013), but this technique has two main weaknesses: it requires manually selecting a small window of wave vectors in Fourier space, and it encounters difficulty when analyzing irregular canvas weaves, such as those arising from erratic spacings or from twill, or swatches with non-thread features. The three canvases studied here were too irregular to be successfully analyzed with the previous technique. To overcome this weakness and to create an analysis technique with less need of human intervention, we developed a new, more robust, automated analysis technique based on autocorrelation analysis and pattern recognition algorithms. The new technique and its application to the canvases of *Pan*, *Bacchus*, and *Silenus* is described below, with full details of the algorithm provided in the Appendix.

In order to validate the results of the automated analysis and also to obtain thread-level correspondence among the three canvases, we also developed a quantitative procedure for manual comparison of thread spacings among canvases. In this method, the user manually marks the locations of all of the threads crossing a "guide thread" in each canvas, so that the local thread spacings are known for every thread. This method, and its application to the three canvases studied here, is also presented below.

2.1 AUTOMATED ANALYSIS PROCEDURE

The automated analysis of a canvas proceeds in three main phases. First, the individual overlapping radiographs tiling a painting are digitized and stitched into a single whole-painting radiograph using a multiscale image analysis technique. Second, the whole-painting radiograph is decomposed into a large number of strongly overlapping small square swatches, each of which is analyzed to extract its so-called "canvas basis vectors", from which the dominant spacing and orientations of both horizontal and vertical threads can be extracted. Third, these geometric parameters from each patch are assembled and visualized using whole-canvas thread spacing and angle maps, from which the similarity of the spacing patterns among a set of canvases can be assessed visually.

2.1.1 Step 1: Multiscale Radiograph Stitching

For each canvas, the analysis begins with a collection of digitized high-resolution radiographs. These must be stitched into a whole-painting radiograph free of artifacts such as visible seams or ghosts. We have described our procedure for stitching the radiographs elsewhere (Johnson, Erdmann, and Johnson 2011), but we summarize here for completeness. First, each scanned radiograph is analyzed to mask out the bright non-radiograph border region which occurs when the radiograph film is not perfectly aligned to the scanner bed or in the typical case where it has rounded corners. The PCA-SIFT algorithm (Ke and Sukthankar 2004) is then used to detect and catalog approximately 20,000 characteristic features in each source radiograph. The collection of characteristic features across all source radiographs is then analyzed to find features with similar appearance such as those resulting when the same feature is captured by more than one overlapping source radiograph. These apparent matches are then filtered to remove false positive matches using the RANSAC statistical algorithm (Fischler and Bolles 1981). Next, a global optimization is performed to determine the optimum placement of each source radiograph relative to the others such that the root-mean square (r.m.s.) distance over all pairs of matching feature points is minimized. Each source image is then resampled using high-order Lanczos interpolation, and the collection of images is blended into a single image using a multiscale Laplacian pyramid (Burt and Adelson 1983) to ensure that subtle exposure differences between the radiographs do not result in visible seams or other artifacts.

2.1.2 Step 2: Analysis of individual radiograph swatches to extract thread spacing and orientation

Once a whole-painting radiograph is assembled, the next step is to extract the characteristic canvas geometry from each of a large number of strongly-overlapping square radiograph swatches. The desired geometry is shown schematically in figure 1. The basic model is that for a small swatch of the radiograph, the canvas spacings and orientations should be approximately constant, in which cases there will be a characteristic horizontal offset vector, **h**, needed to move from one thread crossing to another along a horizontal thread, and a vertical offset vector, **v**, needed to move from one thread crossing to another along a vertical thread. These two vectors form the bases for a local grid of thread crossings, and are sufficient to extract the local trajectory angle for the horizontal and vertical threads, θ_h and θ_v, the spacing between the horizontal and vertical threads, d_h and d_v and the area A associated with each thread crossing.

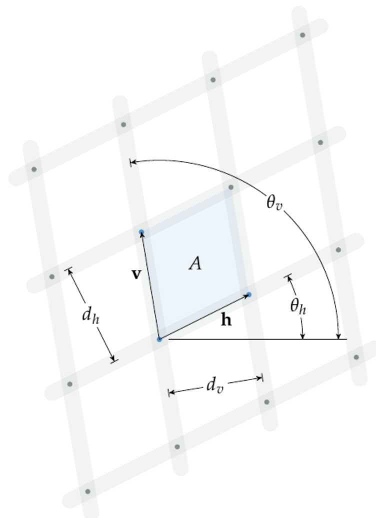

Figure 1: Schematic illustration of a local patch of canvas with threads indicated by light gray bars and thread-crossings indicated by dots. The canvas basis vectors, **v** and **h**, are extracted from the local radiograph patch. From these, the angles of the horizontal and vertical threads, θ_h and θ_v, respectively, and the spacing between the horizontal and vertical threads, d_h and d_v, respectively, and the area per thread crossing, A, can be computed.

The grid of overlapping patches is chosen with a large degree of overlap in order to be robust to small anomalies in the canvas since a median filter is applied in step 3 after the results are assembled. We utilize a uniform grid of 2 cm square swatches with horizontal and vertical center-to-center spacing of 8 pixels (339 µm at 600 dpi). Thus, any given pixel in the interior of the radiograph is actually contained in $59^2 = 3481$ different swatches. Each swatch can be analyzed independently of the others, so that the analysis of the entire canvas can be easily decomposed into many smaller tasks and performed simultaneously across a large array of processors in order to keep the overall time for the analysis manageable.

2.1.3 Step 3: Assembly of canvas geometry and visualization

The final step of the automated analysis procedure is the assembly of the local canvas basis vector **h** and **v** for each of the many square swatches, followed by geometry extraction, filtering, and visualization. First, the results from the analysis of each swatch are collected from the parallel analysis procedure and reassembled into their original grid. Next, simple trigonometry is used to obtain θ_h, θ_v, d_h, d_v, and A at each grid point in a two-dimensional array covering the original radiograph. Each array is then filtered using a 5×5 moving window median filter to remove noise and artifacts arising from irregularities in the original radiograph. Finally, histograms and summary statistics are generated for each of the fields along with whole-painting "canvas maps" showing the local spacings and angles for the horizontal and vertical threads, along with thread intersection angles and local areal thread intersection densities.

A comparison of the summary statistics for the mean and standard deviations of the horizontal and vertical thread spacings serves as a first crude indicator of whether a pair of canvases may have matching spacing patterns and whether an alignment between them would require rotating one of the canvases

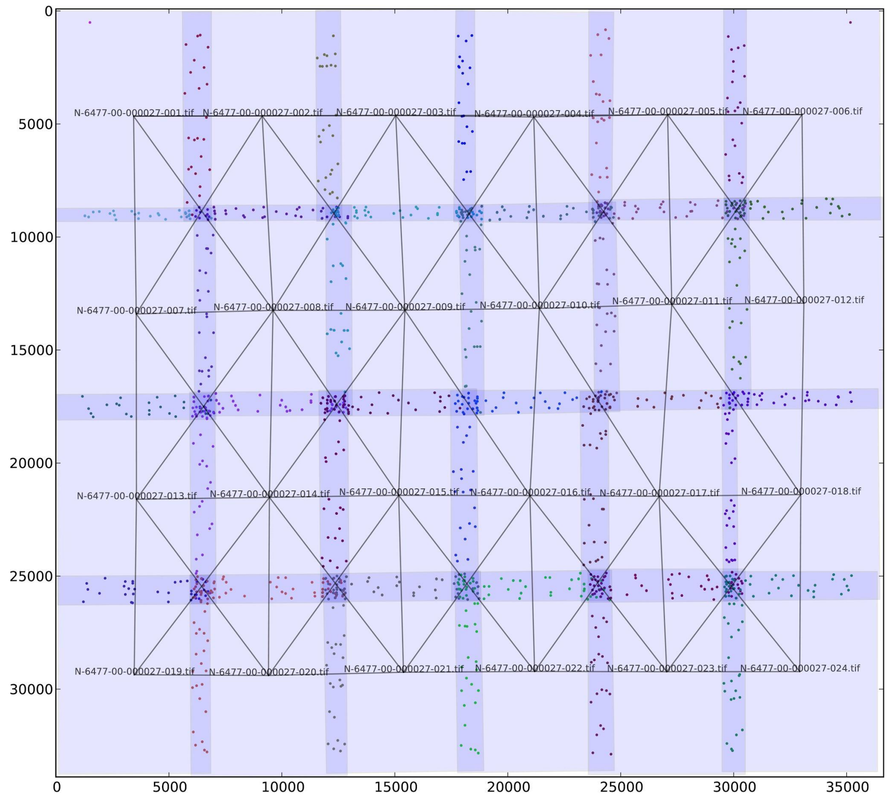

Figure 2: Radiograph overlap structure of *Triumph of Pan*. Units are in pixels. Small circles show the locations of shared keypoints. When a pair of images is known to share keypoints, their centers are connected with a solid black line. Each radiograph is rendered using a translucent blue to indicate the locations and extent of the overlap.

by $\pm 90^{\circ}$. Next, a comparison of the histograms can further reveal differences in the textures between a pair of canvases, and can also provide evidence within a single painting about which of the two thread directions is the warp and which is the weft since the weft typically shows greater variation (van der Wetering, 2000) [3]. Finally, a visual inspection of the thread spacing maps can quickly reveal qualitatively whether a pair of canvases shares a characteristic spacing pattern, and individual canvas angle maps can reveal the pattern of secondary and primary cusping within an individual painting.

2.2 RESULTS OF AUTOMATED ANALYSIS FOR *PAN*, *BACCHUS*, AND *SILENUS*

The radiographs for *Pan* and *Silenus* were provided by the National Gallery in London at 600 dpi (236 pixel/cm) and 1200 dpi (472 pixel/cm), respectively. The radiographs for *Bacchus* were provided by the Nelson-Atkins Museum of Art at a resolution of 500 dpi (197 pixel/cm). Each collection was stitched into a whole-painting radiograph using the procedure described above. *Bacchus* was then upsampled to 600 dpi and *Silenus* was downsampled to 600 dpi to ensure consistency of analysis across all three canvases. A sample of the source radiograph overlap structure for *Pan* is shown in figure 2.

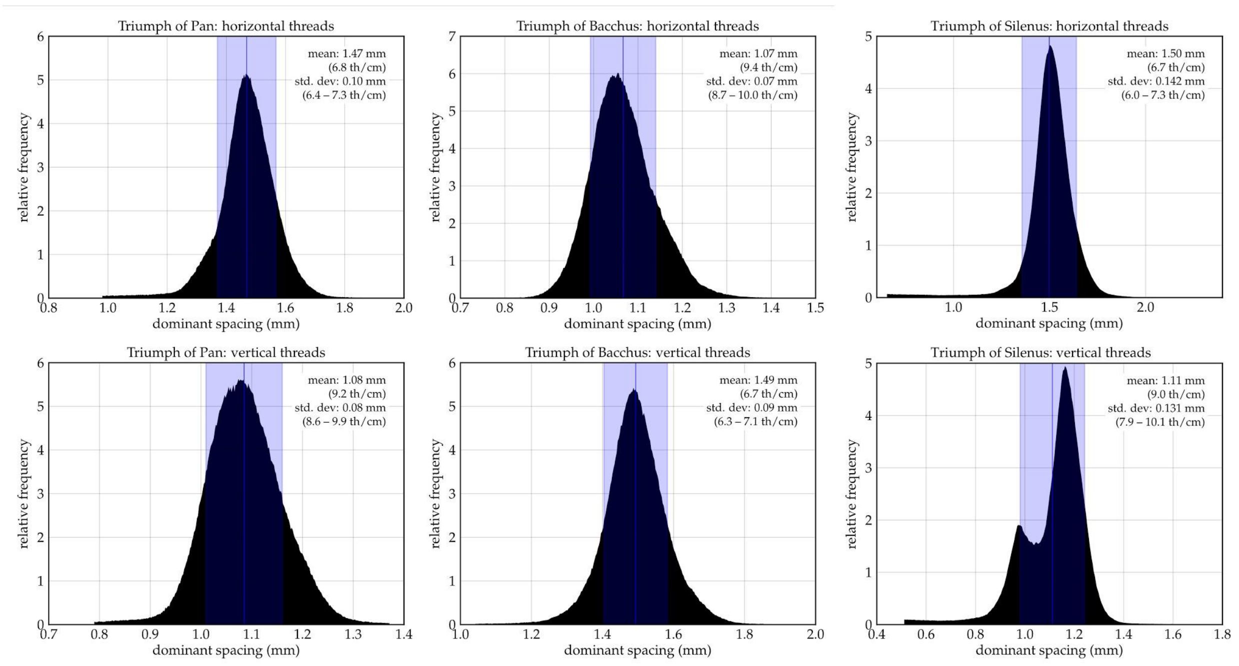

Figure 3: Histograms and summary statistics from the analyses of *Pan*, *Bacchus*, and *Silenus*. Blue shaded ranges indication mean plus and minus one standard deviation. Values within 1 thread/cm of each other are considered to indicate a possible canvas match. The bimodal distribution of the vertical threads of *Silenus* corresponds to the fact that there are two distinctive regions of vertical thread spacing as shown in figure 4 (bottom right panel), approximately divided by the line $x = 85$ cm, each of which has a unimodal distribution.

2.2.1 Thread Spacing Statistics

The autocorrelation-based analysis procedure was performed across each whole-painting radiograph, resulting in the histograms and summary statistics shown in figure 3 and summarized in table 1. The mean spacings all match closely for the horizontal, vertical, and horizontal thread spacings for *Pan*, *Bacchus*, and *Silenus*, respectively, with mean thread spacings close to 1.48 mm in all three. The values for vertical, horizontal, and vertical mean thread spacings for *Pan*, *Bacchus*, and *Silenus*, respectively, are also very similar, with values of around 1.08 mm. These are well within the 1 thread/cm difference under which canvases can be considered to possibly match (van der Wetering, 2000). The statistics also indicate that the *Bacchus* canvas should be rotated $\pm 90°$ relative to its normal orientation to search for an alignment of the thread spacing patterns among the three canvases.

2.2.2 Canvas Maps

Whole-canvas thread spacing maps created for all three canvases are shown in figure 4. Several observations can be made immediately. First, comparing the horizontal and vertical spacing maps of *Pan*, it is seen that the vertical spacing presents a more splotchy appearance, with features persisting for less than the entire height of the canvas. By comparison, the horizontal thread spacings are smoother in ap-

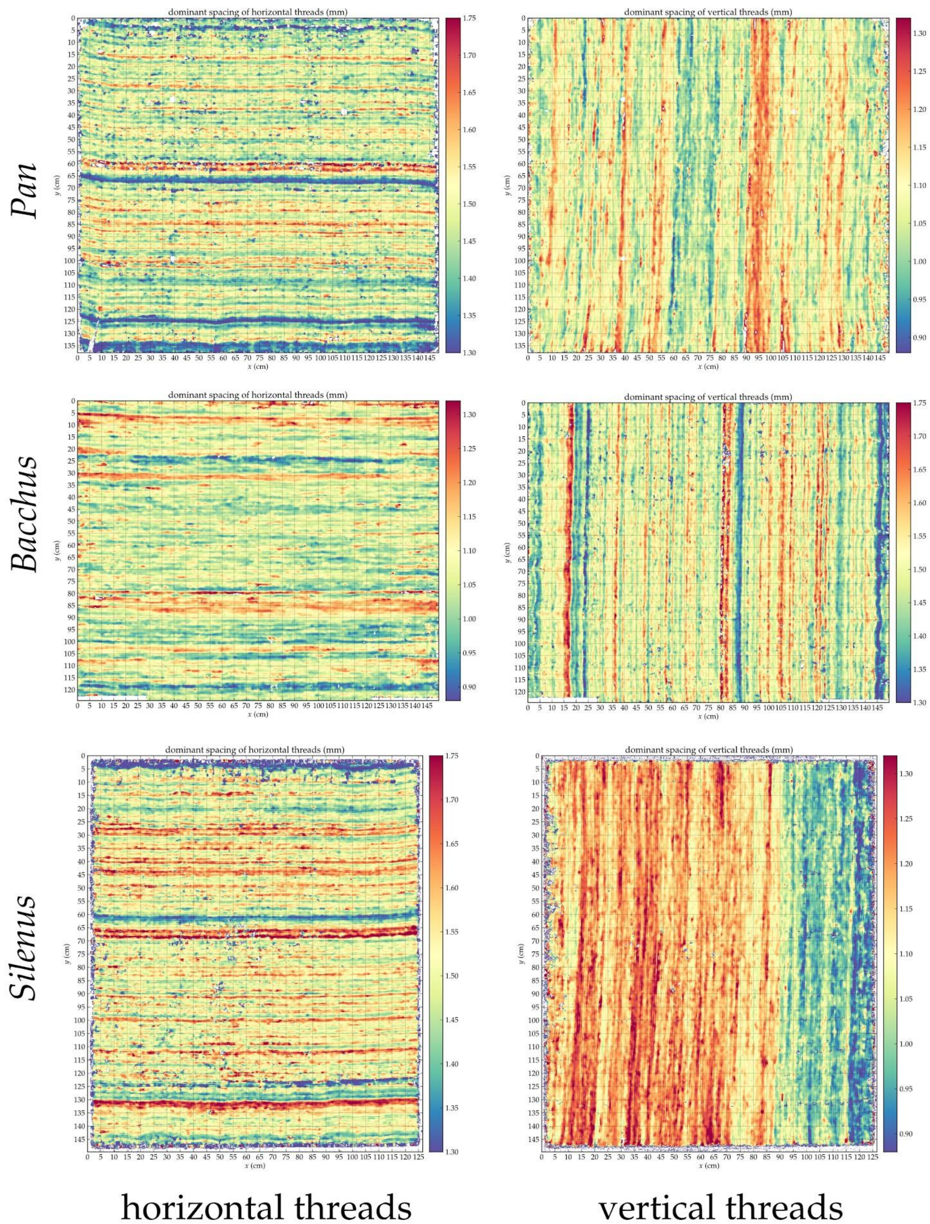

Figure 4: Horizontal and vertical thread spacings for all three canvases shown in their usual orientations. The color scales for the horizontal, vertical, and horizontal threads for *Pan*, *Bacchus*, and *Silenus*, respectively, are identical. The color scales for the vertical, horizontal, and vertical dominant thread spacings for *Pan*, *Bacchus*, and *Silenus*, respectively, are also identical.

Table 1: Summary statistics for automated canvas analysis for all three canvases.

	Horizontal threads			Vertical threads		
	Spacing (cm)	Density (th/cm)		Spacing (cm)	Density (th/cm)	
Painting	mean ± 1 std. dev.	mean	range	mean ± 1 std. dev.	mean	range
Pan	1.47 ± 0.10	6.8	6.4 – 7.3	1.08 ± 0.08	9.2	8.6 – 9.9
Bacchus	1.07 ± 0.07	9.4	8.7 – 10.0	1.49 ± 0.09	6.7	6.3 – 7.1
Silenus	1.50 ± 0.14	6.7	6.0 – 7.3	1.11 ± 0.13	9.0	7.9 – 10.1

pearance, with features persisting with only slight variation across the entire width of the canvas. This difference is consistent with the theory that the horizontal threads in *Pan* correspond to the warp threads and that the vertical threads belong to the weft. This is consistent with the observation of higher standard deviation of the vertical thread spacings compared to the horizontal ones. This pattern of distinct difference in appearance between horizontal and vertical threads is also observed in *Bacchus*, where it appears that the vertical threads belong to the warp, and in *Silenus*, where the horizontal threads apparently belong to the warp.

It is also immediately observed that the horizontal, vertical, and horizontal thread spacing patterns for *Pan*, *Bacchus*, and *Silenus*, respectively, bear a strong similarity to each other. As a first qualitative comparison, the three spacing maps were resized to scale and then rotated and aligned in order to obtain the best visual match, resulting in figure 5. To obtain this alignment, *Bacchus* was rotated 90° clockwise and *Silenus* was rotated 180°. The visible images for the three paintings are also shown, scaled, rotated, and translated in the same fashion as the spacing maps. All of the features of the spacing pattern, from the large-scale trends down to the individual small details, appear to match precisely, giving us very strong evidence that the canvases were cut from the same bolt of cloth. This observation is consistent with the previous identification of these aligned threads as the warp threads: canvases were not generally manufactured with widths exceeding approximately 2 m, eliminating the possibility that all three canvases could share weft threads due to their cumulative size. It should be noted that the sequence of images shown in figure 5 is arbitrary; a re-sequencing of the images would apparently result in a match of equal visual quality since the matching features persist across the entirety of each painting. A horizontal mirroring of any of the images would leave the quality of the visible match similarly unchanged. It is also interesting to note that the *Bacchus* and *Silenus* canvases have matching dimensions in this orientation, and, further, that the best match between the spacing maps for these two canvases occurs when the edges of the canvases are aligned.

2.3 MANUAL ANALYSIS PROCEDURE

The striking similarity among the three spacing maps shown in figure 5 provides strong evidence that the canvases were cut from the same bolt of cloth, but three possible criticisms of the whole-canvas automated analysis method immediately arise: (1) the method has not been validated against canvases with known spacing; (2) both the alignment among the canvases and the apparently high similarity of their spacing patterns are only qualitative; and (3) the analysis procedure operates on swatches of canvas

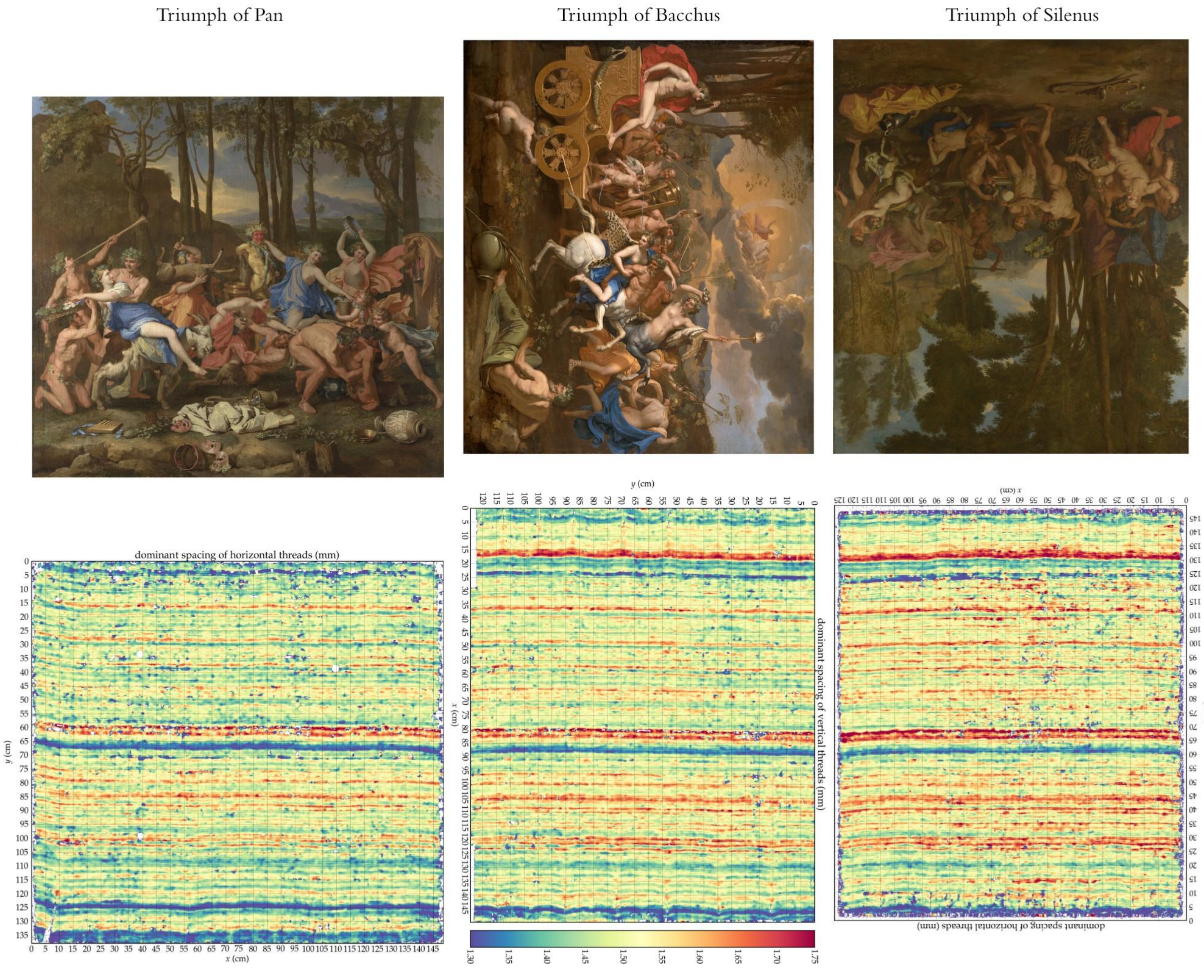

Figure 5: Manually rotated and aligned thread spacing maps for all three canvases. To be put into alignment, *Bacchus* was rotated 90° clockwise and *Silenus* was rotated 180°, so the horizontal, vertical, and horizontal dominant thread spacings for *Pan*, *Bacchus*, and *Silenus*, respectively, from figure 4 are juxtaposed here. Note that the ordering shown here is arbitrary. Also note the matching canvas dimensions and the matching alignment of thread spacing patterns relative to the canvas edges between *Bacchus* and *Silenus*. *Pan* and *Silenus* images courtesy of The National Gallery, London, and *Bacchus* image courtesy of the Nelson-Atkins Museum of Art.

containing many threads, rendering the method unable to detect patterns of thread spacing at the scale of single threads.

To address these criticisms, we developed a simple and straightforward method for direct thread-level manual comparison between a pair of canvases. It is seen from figure 5 that the pattern of spacing of the (apparent) warp threads persists nearly unchanged over the entirety of each painting. Thus, if the center-to-center spacings of the apparent warp threads are manually extracted along a cross-painting cut which intersects all of the warp threads (a vertical cut across each of the canvases in their aligned orientations as shown in figure 5), the particular horizontal location of the cut should be relatively unimportant because we would expect to observe nearly the same pattern of spacings from any vertical cut across the same painting. Thus, the sequence of center-to-center warp thread spacings for the warp threads in a canvas should provide a good characteristic sample of the "canvas fingerprint" for the warp threads in that canvas. Extraction and quantitative analysis thus addresses each of the above criticisms: (1) the thread spacings are directly measured, so their statistics can serve as the ground truth for validating the automatic method; (2) the similarity between two spacing sequences is easy to quantify and interpret; and (3) the method operates at the level of individual threads so even single-thread anomalies can be detected.

The analysis of these so-called "canvas fingerprints" proceeds in two stages. In the first stage, we extract individual thread spacings across each of the canvases to be analyzed, and in the second phase we compare these spacing sequences to obtain a quantitative metric of the similarity of two canvases in a given orientation.

The manual extraction of the thread spacings from a canvas proceeds as follows.

1. The whole-painting radiograph is loaded into the *nip2* image analysis program or a program with similar capabilities. (This program is chosen because it is free, it allows for easy navigation and viewing of very large radiograph images, and it enables the user to manually record the locations of arbitrarily many positions within the image conveniently.)

2. A warp thread is chosen after visual inspection of the whole-painting radiograph according to the criterion that all of the intersections of the given weft thread with the warp threads in the canvas can be visually identified. This weft thread will serve as a "guide thread" along which thread intersections will be marked.

3. Starting at the edge of the canvas, the user manually marks the best estimate for the exact center of the first identifiable thread crossing of the guide thread with a warp thread. The software displays the intersection and records the coordinate, (x_0, y_0).

4. The user then proceeds in a similar fashion, marking every intersection of the guide thread with successive crossing threads until the opposite edge of the canvas is reached. A screenshot from this phase of the analysis is shown in figure 6 . The sequence of intersection coordinates (x_0, y_0) is saved to disk for subsequent analysis.

5. The same procedure can be repeated using any other guide thread on any canvas in any orientation for which quantitative comparison is desired.

The comparison phase of the analysis then proceeds as follows:

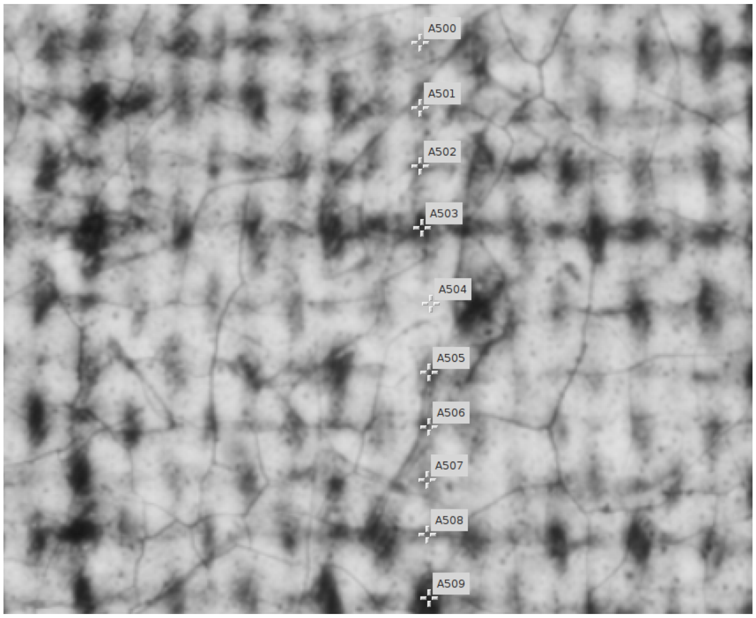

Figure 6: Screenshot from the extraction phase of the manual analysis showing manually marked estimates for the locations of warp threads crossing the vertical weft guide-thread in the whole-painting radiograph of *Bacchus*.

1. The crossing coordinates of the intersections, (x_0, y_0), are extracted and projected onto the axis along which the guide thread runs. For example, in the case of a primarily vertical guide thread, the sequence of y-coordinates y_i of the intersections of the guide thread with the crossing threads is extracted. The same procedure is followed for the second guide thread to which the first is to be compared. We designate the projected coordinates for the second guide thread as Y_i.

2. A sequence of thread-to-thread spacings is determined by computing the differences of the y-coordinates of successive intersections, $s_i = y_{i+1} - y_i$. This sequence comprises the so-called "canvas fingerprint" with which the canvas can be compared with others. The spacing sequence is extracted for both canvases to be compared. The spacing sequence for the second thread is designated as S_i

3. The spacing sequences for the two canvases are exhaustively compared for all possible relative offsets, including those where there is only partial overlap among the sequences but for which at least 300 threads overlap. For each relative offset, the root-mean square difference δ between the two shifted sequences is determined and recorded: $\delta = \sqrt{\sum_{i=1}^{n} (s_i - S_i)^2}$.

4. After all relative offsets are compared, the offset with the lowest r.m.s. difference among the two sequences is taken as the best relative offset between the two sequences. An example comparison is shown in figure 7. Each sequence is truncated at its ends as necessary to retain only those threads occurring within the overlapping portion of the sequence.

5. The end-to-end distance between the first and last thread in each sequence is found. For the

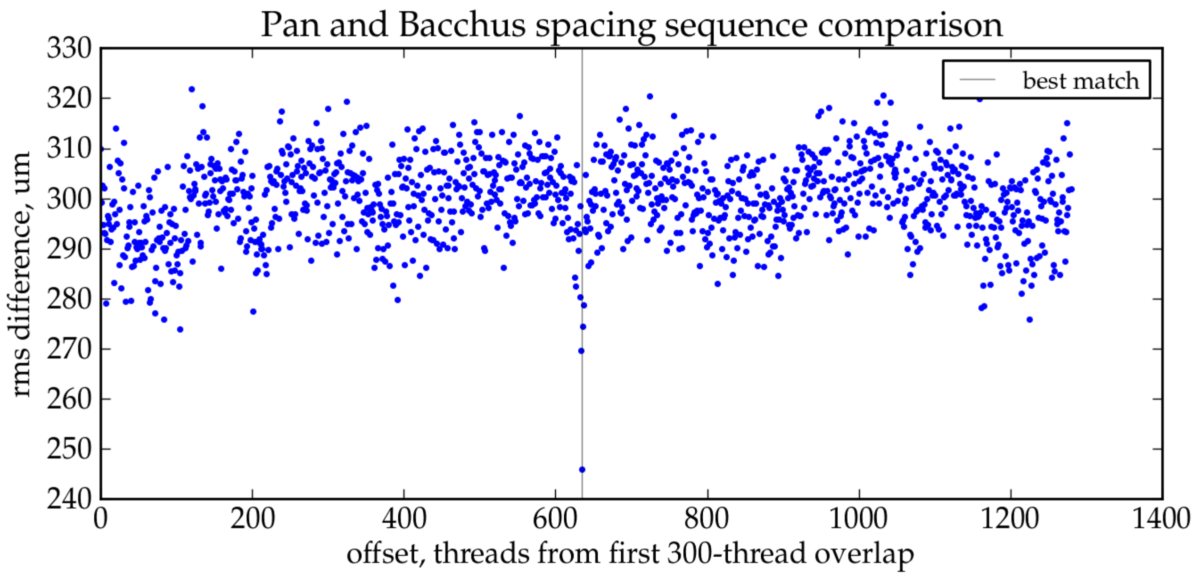

Figure 7: Result of step 3 of the comparison of the spacing sequences along the center guide threads of *Pan* and *Silenus* showing the r.m.s. difference between the overlapping portions of the sequences at different relative offsets. The best offset is shown with a vertical line.

sequence with the lower end-to-end distance, the sequence is uniformly scaled up to match the span of the longer one. This is motivated by two considerations. First, either a slight misalignment of the angles of the two radiographs with respect to each other or a non-orthogonal trajectory of the guide thread relative to the edge of the painting could both result in compression when the sequence of intersections is projected onto the y-axis. Second, it is anticipated that differences in handling of two canvases after they are cut will result in differences in their accumulated after they are cut. Direct comparison between two canvases to determine whether they were cut from the same bolt should then compare their spacing sequences in a same-span normalized form. We scale up the shorter of the two sequences rather than scaling down the longer one in order to obtain a pessimistic estimate of their difference.

6. The r.m.s. difference between the two sequences (one in its original form, and one stretched to give it the same length) is computed and recorded for direct quantitative comparison of the two sequences.

2.4 RESULTS OF MANUAL ANALYSIS FOR *PAN*, *BACCHUS*, AND *SILENUS*

The above procedure for manual comparison of a pair of thread sequences was performed for five guide threads on each of the Pan, Bacchus, and Silenus canvases in the orientations shown in figure 5, corresponding to approximately 15,000 manually marked thread intersections. For each canvas, the five vertical guide threads were chosen at locations near the left edge, at approximately $\frac{1}{4}$, $\frac{1}{2}$, and $\frac{3}{4}$ of the distance from the left to the right edge, and near the right edge of the canvas. We avoid the extreme

Table 2: Summary statistics for the spacings manually extracted along the center guide thread for all three canvases.

Painting	Spacing mean ± 1 std. dev. (cm)	Density mean ± 1 std. dev. (thread/cm)
Pan	1.52 ± 0.20	6.59 ± 0.87
Bacchus	1.54 ± 0.20	6.51 ± 0.86
Silenus	1.54 ± 0.20	6.50 ± 0.94

edges due to the severity of cusping there since it is not known to what extent cusping affects the sequence of spacings for threads approaching the edge of the canvas. Summary statistics for the manually extracted thread spacings for the center thread on each canvas are shown in table 2. The mean spacings are very similar, as are the standard deviations for all three. The mean values are slightly higher than those obtained through the automated procedure, but the differences are well within one standard deviation. It is also noted that the standard deviations for the manually extracted thread spacings are higher than those from the automated procedure. We hypothesize that this is because the automated procedure extracts the dominant spacing over an area of canvas rather than the spacings of the individual threads, resulting in a reduced variance due to the aggregation of several spacings in the estimate for one swatch. figure 8 shows an exhaustive comparison among all possible pairs of threads from each of the five guide threads for each of the three canvases. It shows that the choice of guide thread has little effect on the quality of the match, with all spacing sequences matching quite well. The mean r.m.s. difference between thread-to-thread spacings among all pairs of threads is 0.243 mm (standard deviation 0.019 mm; self-self comparisons are omitted from these statistics). As can be seen by comparison with figure 7, this figure is considerably lower than the typical r.m.s. mismatch between a pair of misaligned sequences such as would be found between a pair of thread sequences having similar statistics but mismatched patterns. figure 8 also shows that the quality of match between pairs of spacing sequences between two paintings is similar to the match quality found between spacing sequences within one painting.

The selection of multiple guide threads at different locations within the same painting was motivated by a desire to further investigate the relationship between spacing sequences and distance along the roll. If the spacing sequences along two closely-spaced guide threads within the same canvas tended to exhibit greater similarity than those along distant guide threads, we might infer something about the roll spacing of two canvases from the similarity of their spacing sequences. At this time, however, we do not observe an apparent decay in the quality of the match between spacing sequences in the same painting as their separation increases, so the match results obtained here are unable to resolve questions about the relative sequence of the paintings on the bolt or of their proximity to each other on the bolt.

3 DISCUSSION

Computer-automated weave analysis has established compelling evidence that the canvases of *Pan*, *Bacchus,* and *Silenus* are contemporaneous and were cut from the same bolt of cloth. For the purposes of the

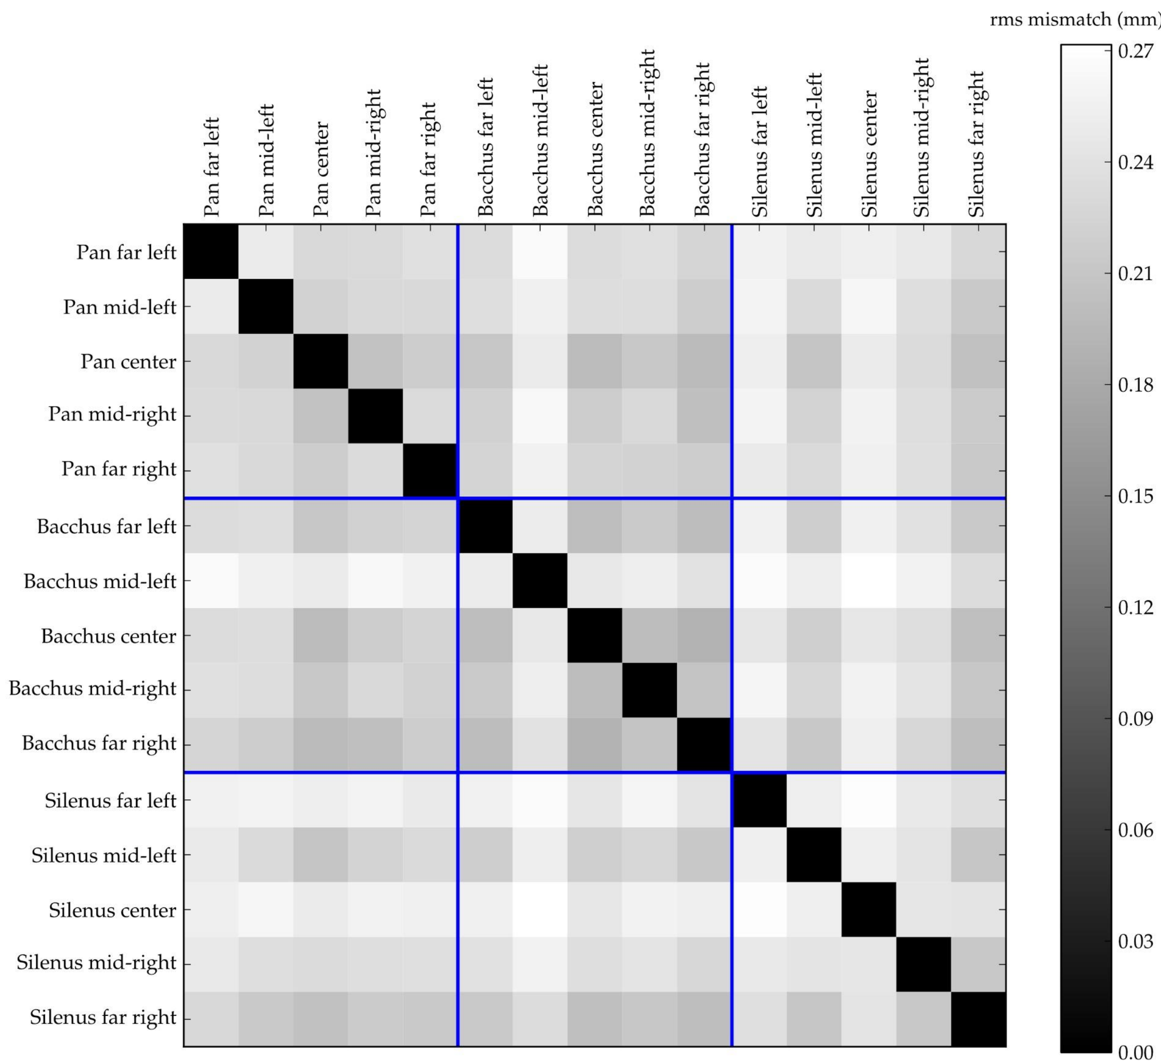

Figure 8: Exhaustive comparison of the best match between spacing sequences along each of five guide threads in each of the three canvases. The matches among all threads are very close, and the inter-canvas matches are of similar quality to the intra-canvas ones.

Nelson-Atkins study, the close warp thread match between *Bacchus* and *Pan*, a painting that is unchallenged today, strengthens the widely held view that *Bacchus* is indeed an autograph work.

The surprising weave match between *Pan* and *Bacchus* and the debated *Silenus* confirms a much closer connection among the canvas materials of the Richelieu Bacchanal series than was previously known. Considering that all three canvases originate from the same bolt, the National Gallery *Silenus*, whether it is the original version or not, must have been painted around the same time period as *Pan* and *Bacchus*. While the results of the weave analysis yield important new information about the Richelieu Bacchanal series, further research on Poussin's studio practice is necessary to resolve the complex attribution issues of *Silenus*.

In the National Gallery catalog, Wine (2001) provides a measured summary of the *Silenus* debate and emphasizes that the primary objection to the painting is its lack of quality in localized areas: "Given the nature of the double ground and the composition of the pigments, similar to those in [*Pan*] (see Technical Notes), the quality of the paintwork in parts and the fact that assessment is made difficult by the picture's overall wear, the view that [the National Gallery *Silenus*] is autograph is not unreasonable. However, the technical evidence is not conclusive – it may be that [*Silenus*] was executed soon after the original by someone with access to Poussin's studio..." (Wine 2001, 380). This possibility, however, remains ambiguous given the lack of information available on Poussin's studio. The fact that connoisseurs have identified both inferior and acceptable passages on *Silenus* raises the question of whether Poussin and assistant(s) may have cooperated in the execution of the third bacchanal.

Poussin's slow and systematic method of painting has lead supporters of the National Gallery painting to speculate that Poussin may have struggled to complete the Richelieu Bacchanal series in a timely manner (Brigstock 1994). Letters document the arrival of two bacchanals, assumed to be *Pan* and *Bacchus*, to Richelieu's chateau shortly after their completion in May of 1636. However, the original *Triumph of Silenus* may have been completed and delivered at a slightly later date (Wine 2001). One argument in favor of the authenticity of the National Gallery painting is that time constraints may have impacted *Silenus* in terms of quality or level of finish, perhaps even prompting Poussin to engage studio assistants to expedite the completion of the commission (Brigstock 1994). On the other hand, most art historians agree that it is unlikely that Poussin would have allowed a substandard painting to enter Richelieu's collection.

If Poussin diverged from the studio tradition of his seventeenth century contemporaries and did in fact work alone—leaving the production of replicas to independent copyists—under what circumstances could he and a copyist have painted from the same bolt of cloth? One explanation for the *Silenus* weave match is that a copyist purchased canvas from Poussin's supplier around the time of the Richelieu commission. That a copyist would begin a reproduction so close in time to the original is feasible given the fact that Poussin completed *The Plague at Ashdod* (7276, Musée de Louvre) in 1631 and the existence of a copy by Angelo Caroselli was documented in a court trial only six months later (Costello 1950) [4]. At the same time, the probability of purchasing a length of canvas for the reproduction that corresponds to the same bolt as the original Richelieu Bacchanal series seems quite low when one considers the active art market in Rome and the numerous suppliers that must have met this demand. Finally, it has been proposed that a copyist may have acquired a scrap of canvas lying about Poussin's studio, but this is also problematic given the size of the canvas, which is roughly equal to the dimensions of *Pan* and *Bacchus* and well suited to any number of projects.

4 CURRENT TECHNICAL STUDIES UNDERWAY AND FUTURE WORK TO BE UNDERTAKEN

Scientific analyses of the palette of *Bacchus* and idiosyncrasies that can be found among its pigments have been undertaken and will be the subject of a future publication. These results will be reviewed in the context of similar studies undertaken on Poussin's works by the scientific staffs of the National Gallery London, and the Louvre. Further investigations into automating the thread-level spacing extraction using machine learning techniques are also underway, as are investigations into the spatial evolution of the canvas fingerprint as a function of separation distance.

ACKNOWLEDGMENTS

Portions of this research were completed with the financial support of the Andrew W. Mellon Foundation endowment for conservation science at the Nelson-Atkins Museum of Art. RGE also gratefully acknowledges the generous support provided by a resident fellowship at the Netherlands Institute for Advanced Study (NIAS). The authors thank the staff at the National Gallery of Art, London, for their cooperation in the digitization and supplying of radiographs for their paintings. The authors acknowledge Nicole Myers, associate curator, Nelson-Atkins Museum, for research and editorial support; Joe Rogers, conservation associate, Nelson-Atkins Museum, for radiography of *Bacchus*; and Lise Koenig for translation of French publications.

ENDNOTES

1. There is some debate as to whether a fourth painting, *The Birth of Venus* (E1932-1-1, Philadelphia Museum of Art), often referred to as *Triumph of Neptune*, was part of the series. Although the painting was also commissioned by Richelieu and may date to the mid-1630s, its format differs and it was not installed in the Cabinet du Roi. Its connection to the Richelieu Bacchanal series remains unclear.

2. At this time, the double ground layers of the Richelieu Bacchanal series have not been compared in terms of their pigmentation and binders.

3. "[I]f the average number of warp threads per centimetre differs by more than one thread, one can practically discount the possibility of the canvases in question coming from the same bolt of cloth. The number of weft threads, with their often more varying thicknesses, can differ a great deal more within a single bolt." on p. 96 from chapter 5 of (van der Wetering, 2000).

4. Caroselli's copy of *Plague at Ashdod* has been attributed to a painting at the National Gallery of Art, London (NG 165). It has been suggested that Caroselli may have copied directly from the original painting and, in the process, recorded an earlier version of Poussin's background architecture (Wine 2001). While this scenario could supply an example of a copyist working in Poussin's studio, there is no connection between the pentimenti on the Louvre painting and the architecture

of the National Gallery copy. Poussin's *Plague at Ashdod* was completed and delivered to Fabrizio Valguarnera in February or March of 1631, and an inventory of Valguarnera's possessions at the time of his court trial in July of 1631 included the Caroselli copy (Costello 1950). Rather than painting in Poussin's studio, Caroselli could have simply executed the copy in the five or six months between Valguarnera's acquisition of the original painting and the date of the trial.

REFERENCES

Blunt, A. 1966. *The Paintings of Nicolas Poussin: A Critical Catalogue*. London: Phaidon. 137: 95–98.

Blunt, A. 1982. "French Seventeenth Century Painting: The Literature of the Last Ten Years." *Burlington Magazine* 124(956): 707.

Brigstock, H. 1994. "Variantes, copies et imitations. Quelques réflexions sur les méthodes de travail de Poussin."*Nicolas Poussin (1594-1665): Actes du colloque organisé au musée du Louvre par le Service culturel. Paris: Louvre*. vol. 1 208.

Burt, P. and E. Adelson. 1983. "The Laplacian pyramid as a compact image code," *IEEE Trans. Commun.*, vol. 31(4), 532–540.

Costello, J. 1950. The Twelve Pictures "Ordered by Velasquez" and the Trial of Valguarnera. *Journal of the Warburg and Courtauld Institutes* 13: 237–284.

Ester, M., H.P. Kriegel, J. Sander, and X. Xu. 1996. "A density-based algorithm for discovering clusters in large spatial databases with noise." *Proc. Intl. Conf. on Knowledge Discovery in Databases and Data Mining (KDD-96)*, 226–231.

Fischler, M. A. and R.C. Bolles. 1981. "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography." *Communications of the ACM*, 24(6), 381–395.

Johnson, C. R., Jr., E. Hendriks, P. Noble, and M. Franken, 2009. "Advances in Computer-Assisted Canvas Examination: Thread Counting Algorithms," *37th Annual Meeting of American Institute for Conservation of Historic and Artistic Works*, Los Angeles, CA, May 2009.

Johnson, D. H., E. Hendricks, and C. R. Johnson, Jr., 2013. "Interpreting Canvas Weave Matches," *Art Matters*, vol. 5.

Johnson, D. H., C. R. Johnson, Jr., and R.G. Erdmann, 2013. "Weave Analysis of Paintings on Canvas from Radiographs," *Signal Processing* (Special Issue on Image Processing for Digital Art Work), vol. 93, 527–540.

Johnson D.H., R.G. Erdmann, and C.R. Johnson, Jr., 2011. "Whole-painting canvas analysis using high- and low-level features.", *Proc. 36th Int. Conf. on Acoustics, Speech and Signal Processing*, 969–972.

Ke, Y. and R. Sukthankar. 2004. "PCA-SIFT: A more distinctive representation for local image descriptors." *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, vol. 2, 506–513.

Liedtke , W., C.R. Johnson, Jr., and D. H. Johnson, 2012. "Canvas Matches in Vermeer: A Case Study in the Computer Analysis of Fabric Supports," *Metropolitan Museum Journal*, vol. 47, 99–106.

Lin, H., L. Wand and S. Yang. 1997. "Extracting periodicity of a regular texture based on autocorrelation functions" *Pattern Recognition Letters*, vol. 18(5), 433–443.

National Gallery of Scotland, Edinburgh. 1981. *Poussin Sacraments and Bacchanals: Paintings and Drawings on Sacred and Profane Themes by Nicolas Poussin, 1594-1665*. Exhibition catalog. Edinburgh: Trustees of the National Galleries of Scotland.

Perez D'Ors , P., C. R. Johnson, Jr., and D. H. Johnson, 2012. "Velazquez in Fraga: a New Hypothesis about the Portraits of El Primo and Philip IV," *The Burlington Magazine*, vol. CLIV, 620–625.

Thuillier, J. 1994. "Poussin et la laboratoire." *Techne* 1: 12–20.

Van Tilborgh, L., T. Meedendorp, E. Hendriks, D. H. Johnson, C. R. Johnson, Jr., and R. G. Erdmann. 2012. "Weave Matching and Dating of Van Gogh's Paintings: An Interdisciplinary Approach," *The Burlington Magazine*, vol. CLIV, 112–122.

Van der Wetering, E. 2000. *Rembrandt: The Painter at Work*. University of California Press.

Wine, H. 2001. *The Seventeenth Century French Paintings*. London: National Gallery.

Witkin, A. 1984. "Scale-space filtering: A new approach to multi-scale description", *IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP'84*, vol. 9, 150–153.

APPENDIX 1. DETAILED DESCRIPTION OF AUTOMATED CANVAS ANALYSIS PROCEDURE

For completeness, we provide the details of the algorithm for extraction of the canvas geometry from a single radiograph swatch here.

1. A square patch $2\,\text{cm} \times 2\,\text{cm}$ is extracted from the stitched radiograph centered on the point of interest. Its mean brightness is subtracted from all pixels for subsequent computation of the autocorrelation.

2. The patch is Gaussian blurred with standard deviations of 4 and 8 pixels, respectively, and the difference of these blurred patches is computed. Gaussian blurring with a standard deviation of pixels results in an image with features smaller than pixels effectively removed (Witkin 1984). Thus, the difference image contains primarily only those features from the original patch having length scales between approximately 4 and 8 pixels. This removes small features such as those arising from paint particles and small cracks, and also removes large features such as macroscopic brightness variations arising from the paint layer, stretcher bars, or other large non-canvas components of the radiograph.

3. The autocorrelation (Lin, Wand, and Yang 1997) of the difference image is computed after padding with zeros and center weighting with a circular Gaussian of standard deviation $0.7\,\text{cm}$ so that the computation procedure does not include any influence from the wrap-around portion of the image and so that areas near the point of interest are more heavily weighted.

4. The peaks of the autocorrelation correspond to offsets of a copy of the image relative to itself which would yield strong agreement between them. Quadratic refinement is used locally to refine the

estimate of the peak location by fitting each autocorrelation function with a local quadratic form and extracting the peak location therefrom.

5. The collection of peaks is analyzed to extract the vector offsets from each peak to its four nearest neighbors. While some anomalies may occur, these should contain a repetition of the dominant canvas basis vectors **h** and **v**. Only vectors with positive projections on the x- and y-axes are retained to avoid double-counting offset vectors.

6. The DBSCAN algorithm (Ester et al. 1996) is used to perform cluster analysis on the collection of offset vectors. The parametric input to the algorithm include the minimum distance within which points are considered to be "connected", which we take as 3 pixels, and the minimum number of connected points needed to constitute a cluster, which we take as 6 points.

7. The cluster centroids are computed, and the cluster centroid lying in the range $\pm 45^\circ$ which is closest to the origin is taken as the horizontal basis vector **h**, while the cluster centroid lying in the range 45°–135° closest to the origin is taken as the vertical canvas basis vector **v**.

AUTHOR INFORMATION

Robert G. Erdmann is Assistant Professor, University of Arizona, Department of Materials Science and Engineering and Program in Applied Mathematics.

C. Richard Johnson, Jr. is Geoffrey S. M. Hedrick Senior Professor of Engineering, Cornell University, School of Electrical and Computer Engineering

Mary Schafer is Associate Conservator of Paintings, Nelson-Atkins Museum of Art

John Twilley is Andrew W. Mellon Science Advisor, Nelson-Atkins Museum of Art

Travis Sawyer is an Optical Engineering major at the University of Arizona