Laurens van der Maaten and Robert Erdmann

Automatic Thread-Level Canvas Analysis

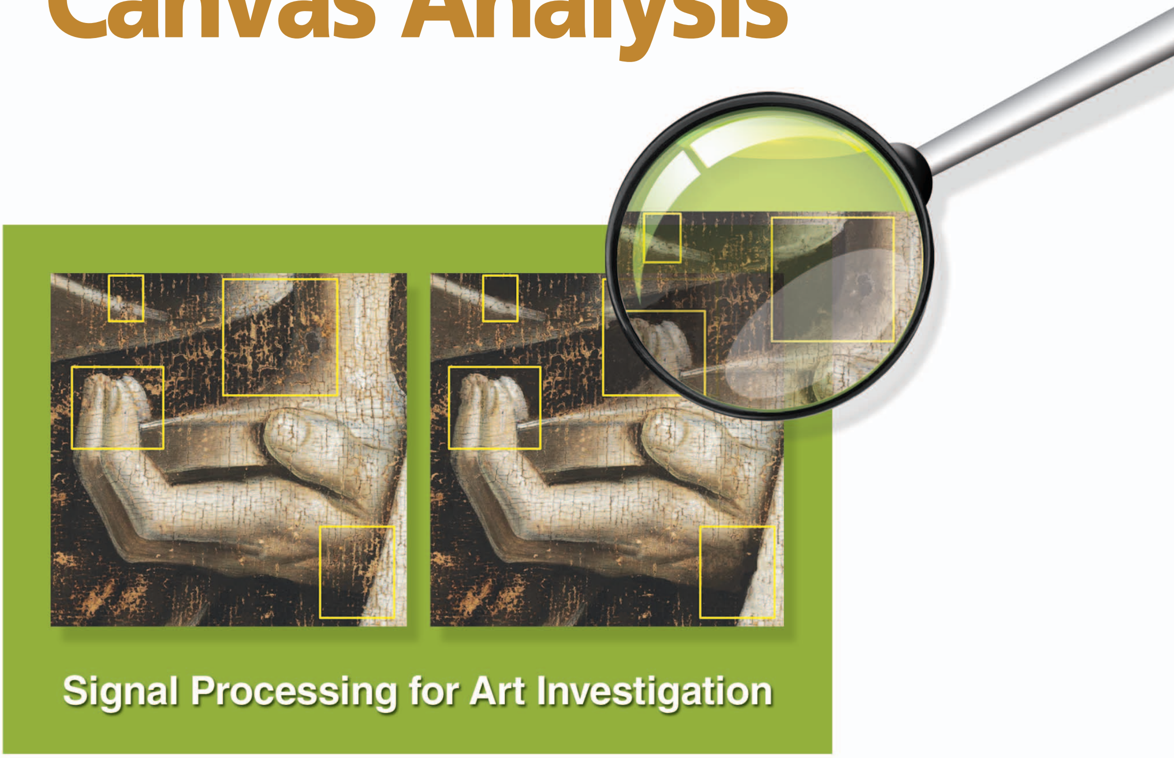

A machine-learning approach to analyzing the canvas of paintings

Canvas analysis is an important tool in art-historical studies, as it can provide information on whether two paintings were made on canvas that originated from the same bolt. Canvas analysis algorithms analyze radiographs of paintings to identify (ir)regularities in the spacings between the canvas threads. To reduce noise, current state-of-the-art algorithms do this by averaging the signal over a number of threads, which leads to information loss in the final measurements. This article presents an algorithm capable of performing thread-level canvas analysis: the algorithm identifies each of the individual threads in the canvas radiograph and directly measures between-distances and angles of the identified threads. We present two case studies to illustrate the potential merits of our thread-level canvas analysis algorithm, viz. on a small collection of paintings ostensibly by Nicholas Poussin and on a small collection of paintings by Vincent van Gogh.

Digital Object Identifier 10.1109/MSP.2015.2407091
Date of publication: 15 June 2015

INTRODUCTION

The analysis of paintings is increasingly aided by the availability of imaging and image-processing tools, including various types of imaging to reveal underpaintings and underdrawings [1], [2], techniques for automatic brushstroke segmentation and analysis [3]–[5], and automatic face analysis techniques [6]. These

 1053-5888/15©2015IEEE

tools can help art historians, conservators, and restoration artists understand the way in which different painters worked and may provide clues about the attribution of a painting to a particular artist. One of the most commonly performed analyses is the analysis of radiographs (X-rays) of paintings, as such radiographs can reveal visible and hidden paint layers according to the radio-opacity of the paint. Radiographs do not only provide essential information on the materials that need to be used in restorations, but they may also form the basis for valuable art-historical insights.

RADIOGRAPHS DO NOT ONLY PROVIDE ESSENTIAL INFORMATION ON THE MATERIALS THAT NEED TO BE USED IN RESTORATIONS, BUT THEY MAY ALSO FORM THE BASIS FOR VALUABLE ART-HISTORICAL INSIGHTS.

In addition to information on the radio-opacity of the paint, radiographs also reveal the individual threads in the canvas (see Figure 1) because the ground layer, which generally contains lead, varies in thickness when it is spread over the textured surface of the bare canvas. Until a few years ago, art experts generally considered the display of the threads as a disturbance because it was obfuscating what they were truly interested in: the composition of the different paint layers. More recently, however, scholars have realized that the canvas threads visible in radiographs may carry important art-historical information [7]. This information arises from the fact that the thicknesses of the threads are irregular because of the way a loom works. Some threads are thinner than others because of natural variations in the manufacturing process: a thread with higher tension on it tends to be narrower. Such irregularities persist throughout the entire bolt of canvas. As a result, paintings made on canvas that was cut from the same bolt will likely have the same irregularities in their thread thicknesses. Thread thickness measurements may be used to identify the bolt from which the canvas originates. Specifically, if we find two paintings that have the same canvas thread thicknesses, we have obtained a strong indication that these two paintings were made in the same workshop in the same period [8].

Thread densities or thread spacings are good surrogates for the thread thicknesses that we would like to measure. Various recent studies have attempted to measure thread densities and/or spacings across the canvas, in particular, in paintings by Nicolas Poussin [9], Vincent van Gogh [10], Johannes Vermeer [11], [12], and Diego Velázquez [13]. In particular, these studies estimate the thread density in a small patch of the painting using a two-dimensional (2-D) Fourier analysis [14] or an approach based on measuring autocorrelations in small canvas patches [9]. These analyses provide valuable information, but they average information across relatively large patches of canvas (over five threads or more), which leads to low-resolution thread density maps. The averaging may hide variations in the thickness of individual threads, which makes it harder to obtain conclusive evidence that two canvases originated from the same roll.

In contrast to most prior work (the work by [12] is a notable exception), this article proposes an algorithm for thread-level analysis of the canvas. Our approach involves training a machine-learning model to identify thread crossings in the canvas based on their visual appearance. The resulting model is used to automatically identify the millions of thread crossings inside a canvas, which, in turn, form the basis for measuring thread spacings. We show the merits of thread-level canvas analysis by using it to study a collection of three alleged Nicolas Poussin paintings as well as a small collection of paintings by Vincent van Gogh.

THREAD-LEVEL CANVAS ANALYSIS

Our approach to thread-level canvas analysis comprises four main parts: 1) we extract features from the radiographs that are sensitive to the signals produced by the threads, 2) we train and deploy a machine-learning model that automatically detects thread crossings based on these features, 3) we use the response of this detector to estimate the distance between neighboring threads, and 4) we automatically try to match the resulting thread-distance maps produced for different canvases to determine whether or not these canvases likely originate from the same roll. The details of these four parts of our approach are described separately below. A MATLAB implementation of our canvas analysis algorithm is publicly available from http://lvdmaaten.github.io/canvas.

FEATURE EXTRACTION

Thread-crossing detection can be performed with very high accuracy because thread crossing corresponds to visually salient locations in canvas radiographs. Our thread-crossing detector: 1) extracts histograms-of-oriented-gradient (HOG) features from the image region around the canvas location of which we want to determine whether or not it corresponds to a thread crossing and 2) uses a linear support vector machine to determine based on these features whether or not the location is a thread crossing or a "nonthread crossing."

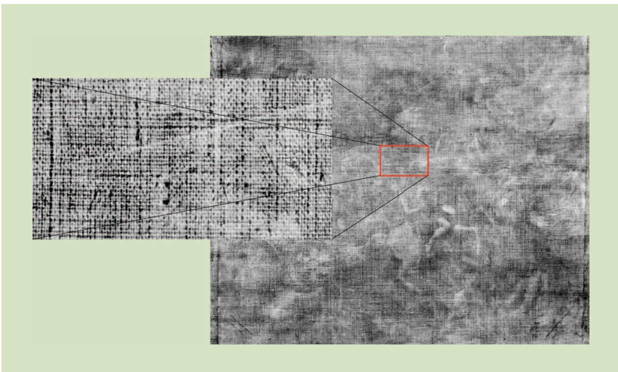

[FIG1] An example of a high-resolution radiograph of the Nicolas Poussin painting *Triumph of Bacchus*. The inset shows the individual threads in a small part of the canvas. (Radiography reproduced with permission from the Nelson-Atkins Museum of Arts in Kansas City.)

HOG features [15] describes an image location by a histogram of image gradient magnitudes for a number of quantized gradient orientations (we used eight orientations in our study). The histograms are constructed over small image patches; depending on the type of canvas and the resolution of the canvas radiographs, we used image patches of 4×4 or 8×8 pixels in this study [in radiographs scanned at 600 dots per inch (dpi)]. Subsequently, the histograms are normalized for contrast differences by normalizing the L2-norm of all the histograms in a square, spatially connected block of four image patches. The advantage of the use of image gradients and the subsequent contrast normalization is that it produces partial invariance to larger-scale signals in the radiograph images that stem from the paint layers (in particular, from layers of white paint that contain relatively large amounts of lead, which in turn lead to strong radiograph responses). To obtain additional invariance to small variations in the gradient magnitudes, the contrast-normalized histograms are clipped at 0.2 and then renormalized according the L2-norm to produce the final HOG features. The resulting features have a particular structure near thread crossings, which is illustrated in Figure 2.

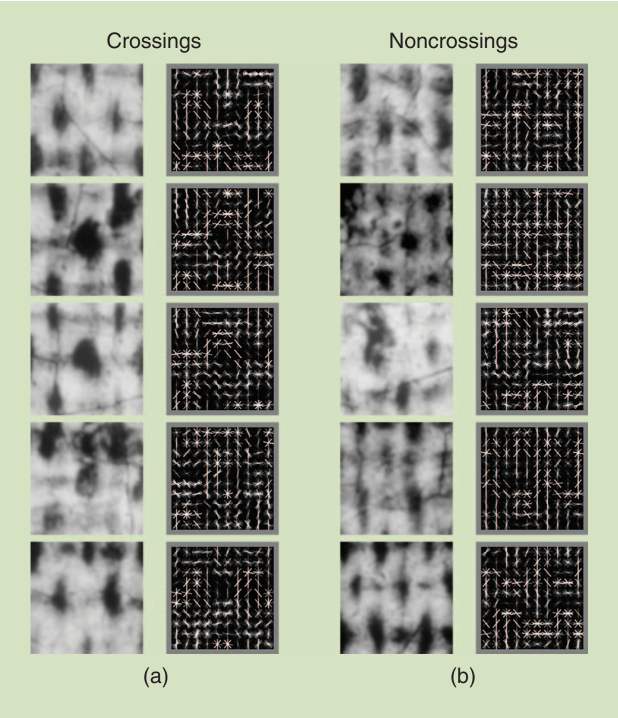

[FIG2] **Examples of five canvas patches around (a) a thread crossing and (b) five randomly selected canvas patches along with the corresponding HOG feature-representation of these patches.**

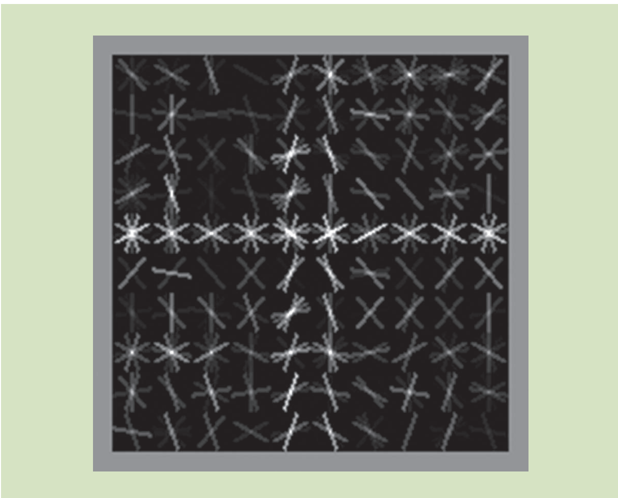

[FIG3] **A visualization of our thread-crossing detector. The figure shows that the detector identifies crossings as a location at which prolonged horizontal and vertical edges (caused by the boundaries of the threads) cross.**

THREAD-CROSSING DETECTION

To obtain a model that can automatically distinguish thread crossings from other structures in canvas radiographs, we train a logistic regression model to discriminate a set of image patches that contain manually annotated thread crossings (positive examples) from a set of image patches that are randomly sampled from the canvas (negative examples). Ideally, the set of positive examples describes the variation in the visual appearance of thread crossings, while the set of negative examples captures the visual variation in nonthread-crossings. Denoting an image patch in the training data by $\mathbf{I}$, the corresponding label by $z \in \{-1, +1\}$, and the HOG feature function by ϕ, the logistic regressor builds the following probabilistic model:

$$p(z \mid \mathbf{I}; \theta) = \frac{\exp[z(\mathbf{w}^\top \phi(\mathbf{I}) + b)]}{\exp[-(\mathbf{w}^\top \phi(\mathbf{I}) + b)] + \exp[\mathbf{w}^\top \phi(\mathbf{I}) + b]}.$$

Herein, the parameters $\theta = \{\mathbf{w}, b\}$ comprise a vector $\mathbf{w}$ and a scalar bias b. After incorporating an isotropic Gaussian prior over $\mathbf{w}$, $p(\mathbf{w}) = \mathcal{N}(\mathbf{w} \mid 0, \sigma^2 \mathbf{I})$, the parameters θ are estimated via maximum a posteriori (MAP) estimation on the aforementioned data set $\mathcal{D} = \{(\mathbf{I}_1, y_1), (\mathbf{I}_2, y_2), \ldots, (\mathbf{I}_N, y_N)\}$:

$$\theta^* = \underset{\theta}{\operatorname{argmax}} \sum_{n=1}^{N} \log p(z_n \mid \mathbf{I}_n; \theta) - \frac{1}{2}\sigma^2 \|\mathbf{w}\|^2.$$

Herein, the hyperparameter σ^2 is set via cross-validation. The resulting weights $\mathbf{w}^*$ are visualized as a HOG feature in Figure 3. The figure shows that they are a template for the visual appearance of a typical thread crossing. This training procedure need to be performed only once for a particular type of canvas, assuming the imaging conditions are similar across the collection of canvas radiographs.

After training, the trained model (i.e., the template) is applied to all image patches in a canvas radiograph to predict the likelihood $p(z \mid \mathbf{I}; \theta)$ that a location in the canvas contains a thread crossing. An example of the resulting likelihood map is shown in Figure 4(c); brighter colors indicate a higher likelihood of the location containing a thread crossing according to the logistic-regression model. The quality of the likelihood map can be substantially improved by exploiting that the likelihood map ought to be quite regular: the likelihood of the thread-crossing presence at location (x, y) should be high when there is a high likelihood of

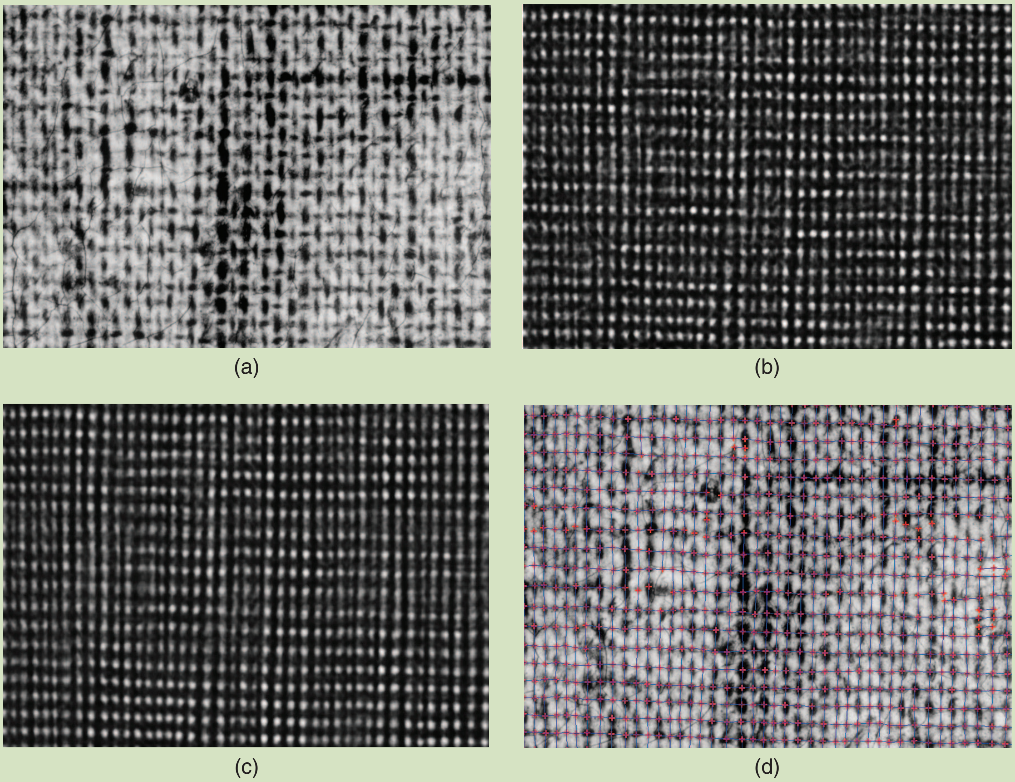

[FIG4] An illustration of our canvas analysis algorithm: (a) a small patch of canvas taken from Poussin's *Triumph of Bacchus*, (b) the response of our thread crossing classifier on the patch of canvas, (c) the response of our model after incorporating the pictorial-structures model that exploit canvas regularity, and (d) the final thread-crossing detections and the identified neighbor relations between these detections. In the response images, a brighter color corresponds to a higher likelihood of a thread crossing being present in the canvas (according to our model). In (d), detected thread crossings are indicated by red crossings. The blue lines indicate the detected neighbor relations: to construct distance maps, distances that were measured between all neighboring thread crossings (i.e., over all blue lines). At locations where blue lines are absent, the distances are interpolated from neighboring thread crossings. (Figure is best viewed in color.)

thread-crossing presence near the locations $(x-d_x, y)$, $(x+d_x, y)$, $(x, y-d_y)$, and $(x, y+d_y)$, where d_x and d_y are the average distances between threads in the warp and weft directions, respectively. We employ a pictorial-structures model [16] that can exploit this information to also detect thread crossings for which little visual evidence is present (e.g., because the thread crossing is hardly visible due to the presence of lead white paint). Our pictorial-structures model computes the score s for thread-crossing presence based on the image patch $\mathbf{I}(x, y)$ extracted at location (x, y) as follows:

$$\begin{aligned} s(x, y; \theta) &= p(z=1 \mid \mathbf{I}(x, y); \theta) \\ &+ \sum_{t_x} \max_{\delta_x, \delta_y} [p(z=1 \mid \mathbf{I}(x+t_x d_x+\delta_x, y+\delta_y); \theta) - \alpha(\delta_x^2+\delta_y^2)] \\ &+ \sum_{t_y} \max_{\delta_x, \delta_y} [p(z=1 \mid \mathbf{I}(x+\delta_x, y+t_y d_y+\delta_y); \theta) - \alpha(\delta_x^2+\delta_y^2)], \end{aligned}$$

where $\alpha \geq 0$ is a manually set discount factor and where $t_x \in \{-1, +1\}$ and $t_y \in \{-1, +1\}$. Intuitively, the score for a location is thus given by the sum of the likelihood for that location and the likelihood of the highest-scoring locations in a four-lattice surrounding that location, where the score of those locations is discounted by the distance to their expected location. While the resulting score is not technically a likelihood, it may be employed in the same way. An example of a pictorial-structures score map is shown in Figure 4(c). The figure illustrates that incorporating prior knowledge on the typical structure of canvas greatly improves the performance of the thread-crossing model. The final thread-crossing detections are obtained by applying nonmaxima suppression on the score map. Nonmaxima suppression finds local maxima in the score map that are above a predefined threshold τ. An example of the resulting thread-crossing detections is shown in Figure 4(d). While some detection errors are present, the result in Figure 4(d) illustrates that the majority of thread-crossings and neighborhood relations is correctly identified.

ESTIMATING THREAD DISTANCES

After detecting the thread crossings, we need to identify which crossings are the warp and weft "neighbors" to be able to measure the distance between two weft threads or two warp threads at that

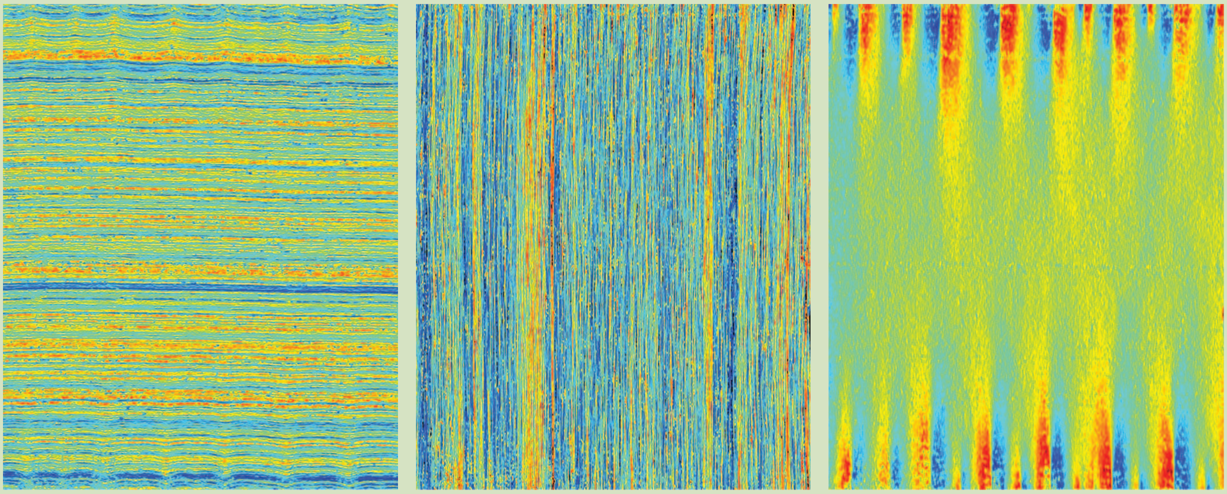

[FIG5] An illustration of the results of our canvas analysis algorithm on Nicolas Poussin's *Triumph of Bacchus*: (a) spacing between warp threads, (b) spacing between the weft threads, and (c) orientation of the warp threads. In the spacing maps, a blue color corresponds to a small distance between threads, while a red color corresponds to a large distance between threads. For the warp threads, thread spacings range between 1.2 and 1.9 mm, while for weft threads, thread spacings range between 0.8 and 1.45 mm. The warp orientation map shows strong cusping on the top and bottom, caused by the deformation of the canvas as it is placed on the stretcer. (Figure is best viewed in color.)

location. To this end, we center an anisotropic Gaussian distribution at each thread crossing location that has much more variance in either the warp or the weft direction (depending on whether we want to measure interweft or interwarp distances). The resulting density map tends to follow the direction of the threads in the canvas. The density maps may be further improved by rotating the Gaussians according to the estimated orientation of the threads at each location to become more robust to cusping, but for simplicity, we omit that in this study. To determine the neighbor of a thread crossing, we perform a large number of random walks that emanate from the thread crossing under investigation and terminate at the next thread crossing we encounter. The random walks are forced to go in a "forward" or "backward" direction, while using the density map to determine whether or not to move in the direction perpendicular to the thread. We construct a histogram over neighbor candidates that counts how often a random walk terminated in each thread crossing, and we select the crossing that has the highest count as the final neighbor candidate. The process is performed both in the "forward" and in the "backward" direction, and a neighborhood relation is only accepted if both thread crossings pick each other as neighbor to eliminate any inconsistencies (i.e., when the neighborhood relation is reciprocal). The detected thread-crossing relations are indicated by blue lines in Figure 4(d). Finally, we estimate interthread distances on all locations where the thread identification procedure has a high confidence, while interpolating in low-confidence regions of the canvas and removing small outliers using a median filter. In a similar manner, we can measure the orientation of the thread connections to produce a thread orientation map. An example of the resulting distance and orientation maps (for both warp and weft threads) is shown in Figure 5.

WE SHOW THE MERITS OF THREAD-LEVEL CANVAS ANALYSIS BY USING IT TO STUDY A COLLECTION OF THREE ALLEGED NICOLAS POUSSIN PAINTINGS AS WELL AS A SMALL COLLECTION OF PAINTINGS BY VINCENT VAN GOGH.

MATCHING THREAD DISTANCE MAPS

To identify potential matches between different canvases based on the thread distance maps, we adopt an approach similar to that described by [17]. Specifically, we extract a small band of the distance map and take the median along this band (in the direction of the threads) to obtain an estimate of the thread distance signal. The thread distance signal is convolved with a Gaussian kernel to remove very fine-grained structure: empirically, we found matching is more accurate when performed based on features in the thread distance signal that live on a coarser scale. We match the thread-distance signals of two canvases by sliding one signal over the other (enforcing a minimum overlap), while measuring the mean absolute distance between the signals in the overlapping region. We use mean average distance as it is less sensitive to outliers than squared errors. The match is repeated for a flip of one of the thread signals as one of the canvases may have been "upside down" compared to the other. A match is only accepted if one of the two minimum mean average distances is a below a certain threshold.

EXPERIMENT 1: NICOLAS POUSSIN

We used our canvas-analysis algorithm to study a collection of three Nicolas Poussin paintings that were studied before by [9]: 1) *Triumph of Pan*, 2) *Triumph of Bacchus*, and 3) *Triumph of Silenus*. This set of paintings is of particular art-historical interest because the three paintings were part of single commissioning in 1636 to Cardinal de Richelieu for the Cabinet du Roi in his castle in Poitou, France. Their authenticity has been subject to strong debate: some have considered *Bacchus* to be a copy [18], [19], but most Poussin scholars now believe that *Bacchus* and *Pan* are authentic Poussin paintings. *Silenus*, however, is considered to be an early copy by its owners, the National Gallery London (Poussin's Bacchanals quickly became very popular, with the first copies being produced as early as 1665. For instance, at least seven known copies of *Bacchus* exist today). Recent canvas analysis results have challenged this belief by finding a canvas match between all three *Triumph*-paintings but were inconclusive because they were unable to perform thread-level canvas analysis. We obtained digital versions of radiographs (scanned at 600 dpi, 500 dpi, and 1,200 dpi for *Triumph of Pan*, *Triumph of Bacchus*, and *Triumph of Silenus*, respectively) and stitched them into whole-painting radiographs using algorithms described in [20]. Thereafter we manually annotated a total of 11,954 thread crossings in these radiographs and trained our thread-crossing detector on these manually annotated positive examples (negative examples were sampled randomly from the same radiographs).

THREAD-CROSSING DETECTION CAN BE PERFORMED WITH VERY HIGH ACCURACY BECAUSE THREAD CROSSING CORRESPONDS TO VISUALLY SALIENT LOCATIONS IN CANVAS RADIOGRAPHS.

We set the value of the L2-regularization parameter in the logistic regression, σ^2, by performing a grid search guided by the classification error on a small held-out validation set. The error of our model on the validation set was approximately 9%; most of these errors were likely due to the set of negative examples containing some actual thread crossings by chance. The average thread distance parameters d_x and d_y were estimated by running the entire canvas-analysis procedure without the pictorial structures and taking the median interthread distance in the warp and weft direction. The nonmaxima suppression step used a window size of 5×5 pixels, and a threshold $\tau = 0.4$ (on a scale from zero to one). The final thread distance maps were cleansed with a 7×3 or 3×7 median filter (depending on the orientations of the threads being analyzed). For weft maps, we removed distance values below 0.85 mm and above 1.45 mm from the map, while for warp maps, we removed distance values below 1.2 mm and above 1.9 mm.

Figure 5 presents the results of our analysis of *Triumph of Bacchus*. Figure 6 presents the results produced by our thread-level canvas analysis algorithm after matching the three Poussin paintings. Different colors correspond to different spacings between individual threads. The results presented in the figure provide very strong evidence that all three canvases originated from the same roll. The results are in line with earlier results from automatic and manual canvas analyses of these three

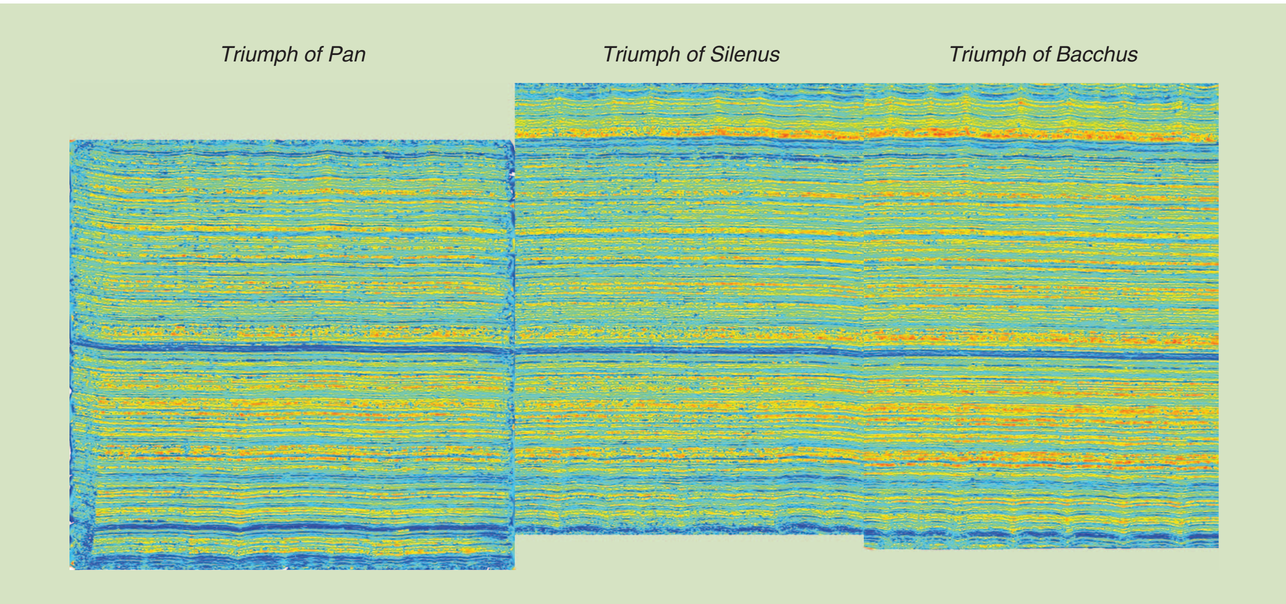

[FIG6] The results of our thread-level canvas analysis of the three *Triumph* paintings. Different colors indicate the distance between (detected) neighboring thread crossings. (The figure is best viewed in color.)

paintings [9] but provide stronger evidence because the evidence is on the level of individual threads and not on the level of multiple-thread averages. Indeed, the presented results make it highly unlikely that *Silenus* was copied 30 years later in a different location (Poussin was working in Rome, Italy, whereas a copyist likely would have worked in France), which strongly suggests that the current art-historical description of the three paintings needs to be revised. We leave such art-historical interpretations to other scholars; they are outside the scope of this work.

EXPERIMENT 2: VINCENT VAN GOGH

We also performed analyses of a small collection of paintings by Vincent van Gogh. Unlike the canvases of 17th-century Poussin, 19th-century van Gogh used canvas produced by the textile industry that has much finer threads and smaller irregularities in thread spacing [10]. Moreover, because van Gogh applies very thick paint layers, the thread structure is much harder to see in the X-rays. As a result, thread-level canvas analysis of van Gogh paintings is substantially harder than the analysis of Poussin paintings.

A large collection of roughly 180 van Gogh paintings has been studied intensively in the context of the Thread Count Automation Project [10], the goal of which is to assign all van Gogh paintings to a particular roll, as this may provide information on the order in which van Gogh made his paintings.

For this study, we had access to a small collection of ten radiographs van Gogh paintings that were scanned at 600 dpi:

- F402 *Two White Butterflies*
- F482 *Bedroom in Arles*
- F490 *Mother Roulin with Her Baby*
- F511 *Orchard in Blossom*
- F633 *The Good Samaritan*
- F692 *The Thresher*
- F699 *The Shepherdess*
- F720 *Enclosed Wheat Field with Rising Sun*
- F734 *The Garden of Saint-Paul Hospital*
- F822 *The Cows*.

(The F-numbers are the catalogue numbers used by the van Gogh Museum in Amsterdam, The Netherlands). Some of the ten canvases are surmised to originate from the same roll but the results of current analyses are inconclusive. Our analysis results in a group of at least four matching canvases, as illustrated in Figure 7. An extensive study on all 180 van Gogh paintings [10] is planned for a future work.

CONCLUSIONS AND OUTLOOK

We have presented a novel canvas-analysis approach that is able to perform thread-level analyses of canvas. We believe the method has two main advantages over prior work: 1) it provides more conclusive evidence on whether or not two patches of canvas have the same thread patterns and 2) it is easier for art experts to understand exactly what is being measured. We believe the second advantage is essential to get canvas-analysis

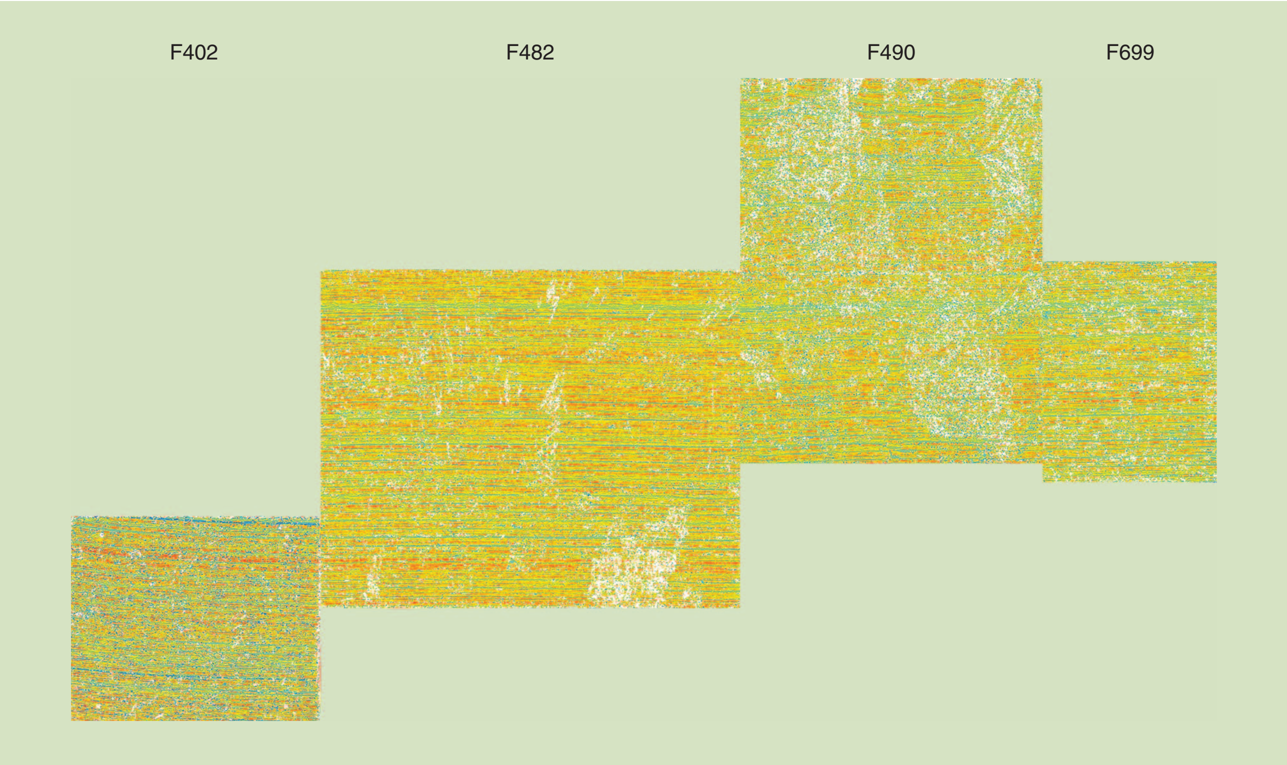

[FIG7] An illustration of canvas weave matches between four van Gogh paintings: 1) F402 *Two White Butterflies*, 2) F482 *Bedroom in Arles*, 3) F490 *Mother Roulin with Her Baby*, and 4) F699 *Shepherdess (after Millet)*. Different colors indicate the distance between (detected) neighboring thread crossings. In white regions, hardly any thread crossings were detected because the crossing signal was obfuscated by thick paint layers; these regions were ignored in the thread spacing measurements. (The figure is best viewed in color.)

technology widely used: showing art experts visualizations such as those in Figure 4 allows art experts to understand the analysis process, to identify potential errors in the measurements, and to manually correct such errors when desired, and to assess whether thread spacings are a good surrogate for thread thicknesses for the canvas a hand.

A substantial drawback of the proposed approach is that a trained thread-crossing detection model is likely only applicable to canvas of a similar type that was imaged under similar conditions: for instance, models trained on the Poussin paintings do not work well on the van Gogh paintings because van Gogh's canvases have much finer threads, which results in a different visual appearance of thread crossings. This implies that to apply our approach to a new type of canvas, it may be necessary to manually annotate a few hundred thread crossings for that canvas type. To resolve this problem, it would be very useful to establish a database with a large collection of canvas radiographs along with a crowdsourcing annotation tool. Such a database would not only facilitate systematic comparisons between canvas-analysis algorithms, but it would also allow for training thread-crossing detectors that can be applied to a wide variety of canvas types. Similar data-gathering and annotation efforts have proven instrumental in improving the state of the art in other computer-vision problems, such as object recognition [21].

ACKNOWLEDGMENTS

We thank the National Gallery London, the Nelson-Atkins Museum of Art in Kansas City, Missouri, and the van Gogh Museum in Amsterdam for generously providing the X-rays that were used in this study. In particular, we are indebted to Mary Schafer, Nicole Meyers, Teio Meedendorp, and Louis van Tilborgh. The manual annotation of thread crossings for the Poussin canvases was done by Travis Sawyer of the University of Arizona. The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007–2013) under grant agreement 600849 (INSIDDE). We also acknowledge support from The Netherlands Institute for Advanced Study.

AUTHORS

Laurens van der Maaten (lvdmaaten@gmail.com) is an assistant professor in computer vision and machine learning at Delft University of Technology, The Netherlands. Previously, he worked as a postdoctoral researcher at the University of California, San Diego; as a Ph.D. student at Tilburg University, The Netherlands; and as a visiting Ph.D. student at the University of Toronto. His research interests include painting analysis, deformable template models, dimensionality reduction, classifier regularization, and tracking.

Robert Erdmann (erdmann@arizona.edu) is a senior scientist at the Rijksmuseum, Amsterdam, and is the Rijksmuseum Professor of Conservation Science at the University of Amsterdam, The Netherlands. He is also special professor for the visualization of art history at Radboud University in Nijmegen, The Netherlands. He was formerly an associate professor at the University of Arizona, Department of Materials Science and Engineering and the Program in Applied Mathematics. He is a member of the Bosch Research and Conservation Project and is the director of the Thread Counting Automation Project. His interests include computational materials science, microstructural analysis, and computational art history and art conservation.

REFERENCES

[1] A. Adam, P. Planken, S. Meloni, and J. Dik, "Terahertz imaging of hidden paint layers on canvas," *Opt. Express*, vol. 17, no. 5, pp. 3407–3416, 2009.

[2] J. Dik, K. Janssens, G. V. D. Snickt, L. van der Loeff, K. Rickers, and M. Cotte, "Visualization of a lost painting by Vincent van Gogh using synchrotron radiation based X-ray fluorescence elemental mapping," *Anal. Chem*., vol. 80, no. 16, pp. 6436–6442, 2008.

[3] C.R. Johnson, Jr., E. Hendriks, I. Berezhnoy, E. Brevdo, S. Hughes, I. Daubechies, J. Li, E. Postma, and J. Wang, "Image processing for artist identification," *IEEE Signal Process. Mag*., vol. 25, no. 4, pp. 37–48, 2008.

[4] J. Li, L. Yao, E. Hendriks, and J. Wang, "Rhythmic brushstrokes distinguish van Gogh from his contemporaries: findings via automated brushstroke extraction," *IEEE Trans. Pattern Anal. Mach. Intell*., vol. 34, no. 6, pp. 1159–1176, 2012.

[5] S. Lyu, D. Rockmore, and H. Farid, "A digital technique for art authentication," *Proc. Nat. Acad. Sci. U.S.A*., vol. 101, no. 49, pp. 17006–17010, 2004.

[6] W. Smith, C. Tyler, and D. Stork, "In search of Leonardo: Computer based facial image analysis of renaissance artworks for identifying Leonardo as a portrait subject," *Hum. Vis. Electron. Imaging*, vol. 8291, no. 82911D, pp. 1–7, 2012.

[7] E. V. der Wetering, *Rembrandt: The Painter at Work*. Oakland, CA: Univ. of California Press, 2000.

[8] E. Hendriks, C. R. Johnson, Jr., and D. Johnson, "Interpreting canvas weave matches," *Art Matters: Int. J. Techn. Art History*, vol. 5, pp. 53–61, 2013.

[9] R. Erdmann, C. R. Johnson, Jr.., M. Schafer, J. Twilley, and T. Sawyer, "Reuniting Poussin's Bacchanals painted for Cardinal Richelieu through quantitative canvas weave analysis," *J. Amer. Inst. Conserv*., 2013. [Online]. Available: http://erg.mse.arizona.edu/Erdmann_Reuniting_Poussins_Bacchanals.pdf

[10] L. van Tilborgh, T. Meedendorp, E. Hendriks, D. Johnson, and C. R. Johnson, Jr., "Weave matching and dating of van Gogh's paintings: An interdisciplinary approach," *Burlington Mag*., vol. 154, no. 1307, pp. 112–122, Feb. 2012.

[11] W. Liedtke, C. R. Johnson, Jr., and D. Johnson, "Canvas matches in Vermeer: A case study in the computer analysis of fabric supports," *Metropolitan Museum J*., vol. 47, pp. 99–106, 2012.

[12] B. Tobin, "Automated canvas spacing analysis," Master's thesis, Dept. Materials Science Eng,. Univ. Arizona, Tucson, AZ, 2014.

[13] P. P. d'Ors, C. R. Johnson, Jr., and D. Johnson, "Velázquez in Fraga: A new hypothesis about the portraits of el Primo and Philip IV," *Burlington Mag*., vol. 154, no. 1314, pp. 620–625, 2012.

[14] D. Johnson, C. Johnson, and R. Erdmann, "Weave analysis of paintings on canvas from radiographs," *Signal Process*., vol. 93, no. 3, pp. 527–540, 2013.

[15] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2005, pp. 886–893.

[16] P. Felzenszwalb, R. Girshick, D. McAllester, and D. Ramanan, "Object detection with discriminatively trained part based models," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 9, pp. 1627–1645, 2010.

[17] D. Johnson, C. R. Johnson, Jr., and E. Hendriks, "Matching canvas weave patterns from processing x-ray images of master paintings," in *Proc. Int. Conf. Acoustics, Speech, Signal Processing*, 2010, pp. 958–961.

[18] A. Blunt, *The Paintings of Nicolas Poussin: A Critical Catalogue*. London, U.K.: Phaidon, 1966.

[19] A. Blunt, "French seventeenth century painting: The literature of the last ten years," *Burlington Mag*., vol. 124, no. 956, p. 707, 1982.

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

[21] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. Berg, and L. Fei-Fei, "Imagenet large scale visual recognition challenge." [Online]. Available: http://arxiv.org/abs/1409.0575

[SP]