

Automatic registration and mosaicking of technical images of Old Master paintings

Damon M. Conover¹ · John K. Delaney^{1,2} · Murray H. Loew¹

Received: 3 November 2014 / Accepted: 26 March 2015 / Published online: 23 April 2015
© Springer-Verlag Berlin Heidelberg (outside the USA) 2015

Abstract The registration of technical art conservation images of Old Master paintings presents unique challenges. Specifically, X-radiographs and reflective infrared (1000–2500 nm) images reveal shifted, or new, compositional elements not visible on the surface of artworks. Here, we describe a new multimodal registration and mosaicking algorithm that is capable of providing accurate alignment of a variety of types of images, such as the registration of multispectral reflective infrared images, X-radiographs, hyperspectral image cubes, and X-ray fluorescence image cubes to reference color images taken at high spatial sampling (300–500 pixels per inch), even when content differences are present, and a validation study has been performed to quantify the algorithm’s accuracy. Key to the algorithm’s success is the use of subsets of wavelet images to select control points and a novel method for filtering candidate control-point pairs. The algorithm has been used to register more than 100 paintings at the National Gallery of Art, D.C. and The Art Institute of Chicago. Many of the

resulting registered datasets have been published in online catalogues, providing scholars additional information to further their understanding of the paintings and the working methods of the artists who painted them.

1 Introduction

Conservators regularly collect X-radiographs and infrared (IR) images (750–2500 nm) of paintings to look for changes in the painting’s composition and for original preparatory sketches [1, 2]. Such information provides insight into the working method of the artist and to some of the materials that were used. For example, a painting, found to have a preparatory sketch that has been modified, points to a composition that was being worked out during the painting process, whereas the same analysis of a “production painting” often shows that a template, or compositional study, was used to lay out the figures. Detailed comparison of the IR images with the corresponding color image often reveals further changes made by the artist, thus giving additional information.

In cases where the artist has made large changes in the composition, these may be visible only by identifying changes in contrast of the paints used, achieved by comparing the IR images and the X-radiographs to the color image. In IR images, highly scattering, low-IR-absorbing pigments are transparent, while umbers, some greens, and organic black pigments remain opaque. In contrast, in X-radiographs, high-density materials such as lead white (a pigment often used for painting faces and hands) remain opaque. An example of this can be seen by comparing the color image with the X-radiograph and IR image of Johannes Vermeer’s *Girl with the Red Hat* (Fig. 1). The IR

Electronic supplementary material The online version of this article (doi:10.1007/s00339-015-9140-1) contains supplementary material, which is available to authorized users.

✉ John K. Delaney
j-delaney@nga.gov

✉ Murray H. Loew
loew@gwu.edu

Damon M. Conover
dconover@gwu.edu

¹ Department of Electrical and Computer Engineering, George Washington University, 800 22nd Street NW, Ste 5000, Washington, DC 20052, USA

² Department of Scientific Research, National Gallery of Art, Washington, DC, USA

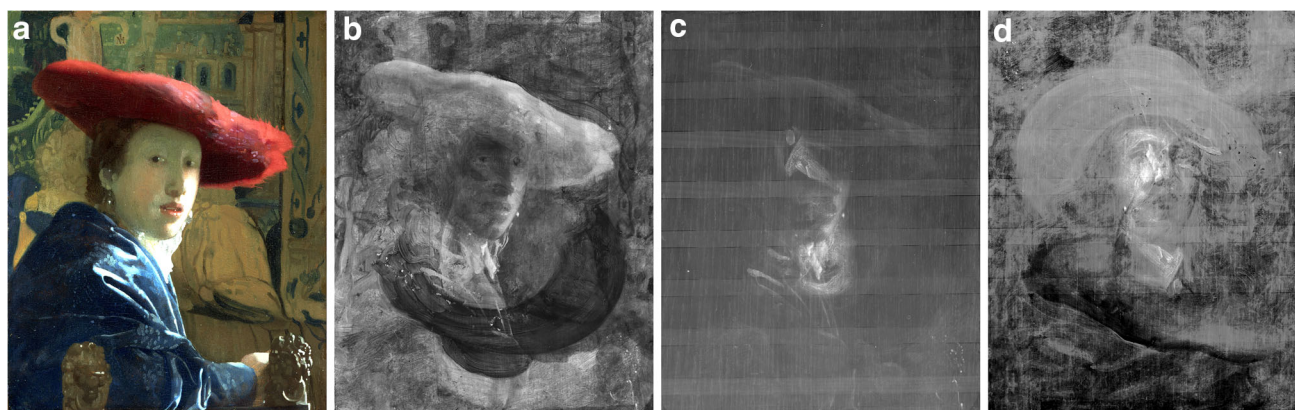


Fig. 1 **a** Color image of Johannes Vermeer's *Girl with the Red Hat* (1665/1666). Andrew W. Mellon Collection, 1937.1.53, National Gallery of Art (NGA), Washington, D.C., **b** infrared reflectance (2100

2400 nm), **c** X-radiograph, and **d** summation of the rotated X-radiograph and the intensity-inverted and rotated infrared reflectance image [3]

image (Fig. 1b) reveals an upside-down hat painted with an IR-absorbing pigment, and the X-radiograph (Fig. 1c) reveals an upside-down face of a man (in addition to the girl) painted, in part, with lead white. Accurate registration of these images allows the construction of a composite (Fig. 1d) that shows both the man's hat and face of the prior composition. Additionally, using the IR image and the X-radiograph, we can examine the painting style of subsurface paint. For this painting, it was determined that the painting style of the reversed portrait is significantly different from that of Vermeer and that it was likely the work of a different artist [3].

2 Background

Automatic registration of images acquired using different modalities ("modalities"), or multimodal image registration, of Old Master paintings presents several challenges. First, these modalities often contain both similar and unique information. While identification of the unique information is the *raison d'être* for multimodal imaging, the lack of equivalent information in the other modalities complicates the registration process. Additionally, some objects that are found in two modalities can shift with respect to one another. For example, an artist might add or reposition a person when moving from the sketch to the final painted composition. More often, one observes small changes in the positions of hands, fingers, eyes, and features of clothing [4]. These small changes complicate the task of accurately aligning images using image content, especially when compared to other image registration applications. In remote sensing of the Earth, for example, these effects do not occur, and in medical imaging, the variation between modalities is due mostly to differences in contrast between objects, rather than to small alterations and shifts.

The second challenge comes from the way the images are collected. In remote sensing, ephemeral data are collected to allow orthorectification of the images, which simplifies the alignment of different imaging modalities. In medical imaging, fiducial markers are often placed in the field of view as common references visible in each modality. Unfortunately, in art conservation, the technical images often are collected by different groups without the required metadata or a sufficient number of fiducial markers to allow geometrical model-based image registration. Moreover, the imaging sensors used differ greatly in their collection modes and associated geometrical distortions. For example, the color images, which have the highest spatial sampling, often are collected with a large-format color array (approximately 20–50 megapixels) and consist of one or a few frames stitched together. While such cameras are of high quality, with lenses having near zero distortion, they produce images that are rectilinear and thus are affected by a scale factor that varies slightly as a function of the distance to the optical axis. For example, a camera (focal length of 100 mm, 8000 horizontal pixels, 6- μ m pixel pitch) 83.8 cm from a painting will produce a maximum shift of 24 pixels at the edges of the image, relative to a regular grid, at 565 pixels per inch. X-radiographs and IR images are not affected to the same degree by this varying scale term. The X-radiographs used are often created by scanning X-ray films, which have been illuminated by an X-ray source. And the IR images typically are composited by mosaicking hundreds of small image frames, with small fields-of-view, captured using IR cameras with arrays typically 0.25 megapixels in size.

Currently such IR and X-radiographs are registered and/or mosaicked using overlap regions between frames either manually, using tools like Photoshop, or semi-automatically using Photoshop Photomerge, Vasari [5], or custom software [6]. A more advanced approach has been

taken to produce well-mosaicked X-radiographs using an algorithm running on a large cluster of networked computers¹ [7]. The use of the overlap regions alone, however, does not guarantee that the resulting mosaic will be able to be registered accurately to the color image using a simple affine transform.

Hardware-specific solutions to remove the mosaicking requirement for IR images have been explored, such as a sensor that uses a single scanning detector or small-area or line arrays that scan the image plane of a large-format lens. These single-pixel raster-scanning systems are slow, taking a few hours to scan a square-meter area [8]. Multi-spectral versions of these raster-scanning systems were also produced that were capable of capturing tens of channels with scan times of seven to fourteen hours per square meter of coverage [9]. Those systems provide moderate resolution (100 pixels per inch) but are not portable.

Truly portable sensors place the raster-scanning detector elements(s) in the focal plane of a lens having a large field of view (FOV). Such systems provide 16-megapixel images (4k-by-4k) by scanning, using either a small-area array [10] or a line array [11]. Those systems have limitations given their low collection efficiency (area covered per unit time), which diminish their utility when used with spectral band filters. Lastly, many conservators use a step/stare collection method in which a painting is moved by a 2-D movable easel in front of a small IR area array with interference filters designed to optimize the contrast of the preparatory drawing in the 750–2500 nm spectral range with sampling of 300–500 pixels per inch [1, 2].

Given the different collection-related distortions inherent to the various modalities and the limited utility of hardware-specific solutions, we sought to design a robust algorithm to register and mosaic the various types of image datasets that are used commonly in conservation science. Here we will discuss a point-based approach for the registration and mosaicking of IR images and X-radiographs to color images. The algorithm, moreover, is capable of registering other imaging modalities, including those obtained from raster-scanning X-ray fluorescence (XRF) systems [12] and line-scanning cameras, such as hyper-spectral reflectance cameras [13–16]. The resulting algorithm was implemented in MATLAB and has been compiled to run on various operating systems. It can run on multi-core desktop computers that are commercially available and found in the photography departments of most museums. Additionally, the user interface of the compiled software was intended to be suitable for use directly by conservators.

3 Methods

Point-based registration algorithms generally consist of four major steps: (1) candidate control-point identification, (2) candidate control-point matching, (3) calculation of a spatial transformation function, and (4) image resampling and transformation [17, 18]. Below we describe the means by which we automatically identify and match control-point pairs to address the challenges of registering mismatched image content that arises between the different imaging modalities, as well as account for and correct the various distortions contributed by the imaging modalities.

3.1 Identify an initial set of control-point pairs

Control points are the corresponding points in two images that will be used to define the transformation function between the images. Selecting a good set of control points is the first step to registering two images, particularly when they are acquired using different modalities. By carefully selecting which control points are initially identified, we can limit the matching step, as described below, to only regions (surrounding the control points) that are likely to match between the two modalities. In the case of a painting, the regions likely to match are the painting's texture features, such as cracks, brushstrokes, bubbles, and blisters, so we select regions that contain objects the same size as those features.

Before control-point identification, we first manually rescale and rotate the images (to within 5–10 %) so that they are at the same approximate size and orientation. We then use the modulus maxima of the wavelet transform to identify regions that contain objects that are the same approximate size as the painting's texture features. The modulus of the wavelet transform is computed by filtering images using low- and high-pass, direction-sensitive filters of various sizes (scales). As shown in Fig. 2, an image with vertical features emphasized, $L_x H_y(n)$, at scale n is created by low-pass filtering (LP) an image with a horizontally-

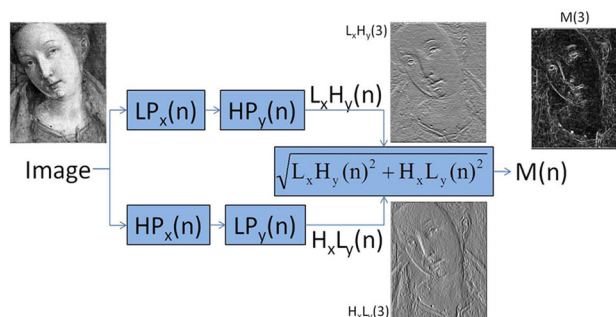


Fig. 2 Steps for computing the modulus of the wavelet transform (*LP* low-pass filter, *HP* high-pass filter)

¹ Personal communication with R. Erdmann, 2013.



Fig. 3 **a** Detail IR image and **b** modulus image with control-point regions (red) identified (scale $n = 3$)

oriented filter and then high-pass filtering (HP) the result with a vertically-oriented filter. The opposite filter order is used to create an image with horizontal features emphasized, $H_x L_y(n)$. The widths of the filters are equal to $2n$. By changing the scale n , we can emphasize features of different sizes. The results then are merged by computing the pixel-by-pixel magnitude of the two images [4, 19]. This produces the modulus, $M(n)$, of the original image. The images, $L_x H_y(3)$, $H_x L_y(3)$, and $M(3)$, in Fig. 2, show the results for scale $n = 3$. To prevent spatial clumping, thus risking a poorly matched group from biasing the result toward an incorrect solution, we use the local maxima from the modulus of the wavelet transform as the initial set of points. A square neighborhood is defined around each maximum, and all maxima within it, except for the largest, are removed. Fig. 3 shows a detail of an IR image frame (Fig. 3a) and the corresponding modulus image at a single scale (Fig. 3b). The initial set of control-point regions that has been identified for the IR image is superimposed, in red, on the modulus image.

After identifying control-point regions in a *template* image (IR images or X-radiographs, in the examples described below), we find corresponding regions in a *reference* image (typically the color image) using a variant of phase correlation. This is performed by first computing the phase image² for each image and then finding the maximum of the normalized cross-correlation between them [4, 20]. Initially, this step is performed on downsampled (using bicubic interpolation) copies of the full IR image frame and full color image. The result is an estimate of the alignment of the IR image frame and the color image. Using the estimated alignment, the algorithm identifies small search

regions in the reference phase image for each of the corresponding control-point regions in the template phase image. The search regions and the control-point regions are then upsampled by a scale factor of three using bicubic interpolation (to minimize artifacts), to achieve subpixel control-point pairings. They are then matched by computing the maximum of the normalized cross-correlation between the full-resolution phase images of each control-point region in the template image and the corresponding (sub)regions in the reference phase image. The resulting cross-correlation matrix for one of those computations is shown in Fig. 4. In addition to this matching step, the candidate maximum correlation peak is tested to distinguish it from the noise floor by comparing it to the standard deviation (σ) of the cross-correlation window after the peak correlation value has been removed. A candidate is accepted if it is (a) greater than or equal to 4σ , and (b) unique (no other value in that (sub-) region is within σ of it). The result is a set of matched control-point pairs³, as shown in Fig. 5a. Those matches, however, are not perfect, which leads to a further filtering requirement, as described in Sect. 3.2. The red points and lines in Fig. 5a define the control-point pairs that are determined to be accurate in the next step, while the blue points and lines show the pairs that will be removed.

3.2 Filtering the initial set of control-point pairs

The sets of control-point regions from each image (reference and template) may represent pairs, i.e., the same point in space. The sets may also contain non-homologous pairs and inaccurate matches. This can be seen in Fig. 5a, where the good pairs (red lines) have the same approximate orientation, while some bad pairs (blue lines) have different orientations. Another way to visualize this is the horizontal pixel disparity plot⁴ in Fig. 5b. The blue and red points represent the *difference* between the corresponding horizontal pixel coordinates of a pair of “matching” control points identified in the previous step. The set of points is plotted in the coordinate space of the first image. If the sets of regions were accurately matched, one would expect the disparity points to be near the best-fit function (shown in black). Even in the situation where the template image is shifted, rotated, and/or magnified relative to the reference image, it would be expected that the disparity will vary slowly across space. Therefore, when a set of points conforms to a consistent disparity relationship, there is reason

² A phase image is produced by computing the 2-D Fourier transform of an image then computing the inverse Fourier transform of the result using only the phase information (i.e., using unit magnitude) [20].

³ While the matching is performed using phase images, Fig. 5a uses the color image and IR image to show more clearly the control points over their corresponding regions.

⁴ The vertical pixel disparity is also computed, and analyzed, in the same manner as the horizontal pixel disparity.

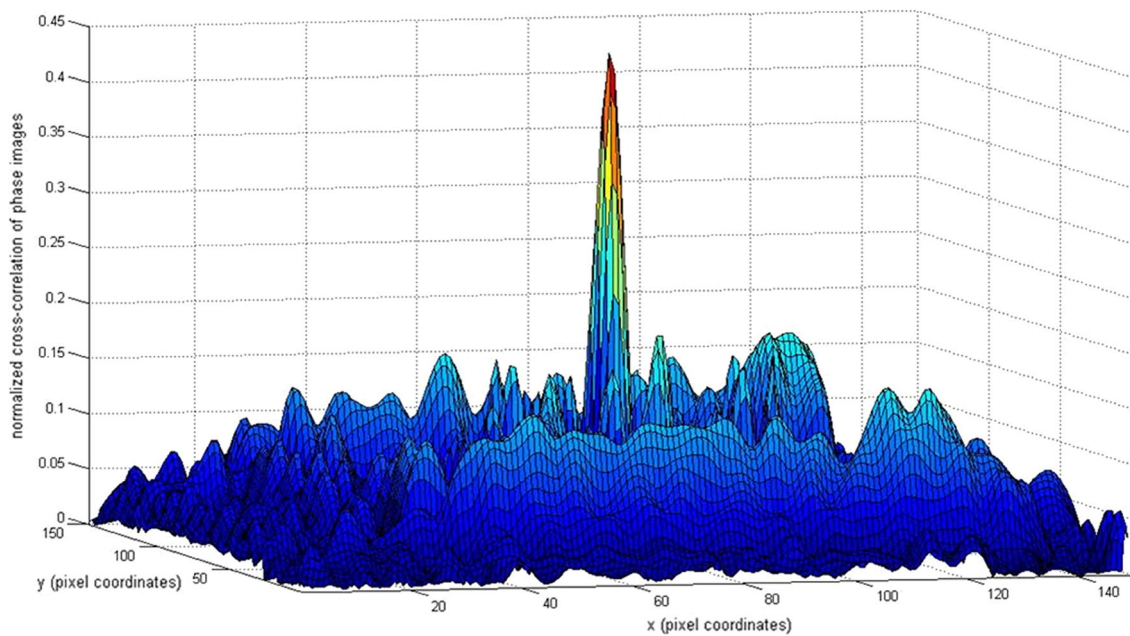


Fig. 4 Normalized cross-correlation of local regions of the phase image of the *red channel* of the color image and a corresponding phase image of an IR image (SD = 0.03)

to expect that the corresponding pairs are valid. Similarly, when there are outliers from that disparity relationship, it can be assumed that those outliers represent false pairs. To choose the best set of control-point pairs, pairs are removed based on the distances of their disparity points from their respective best-fit functions. Examination of the 3-D plot shows most points lie in a tilted plane that can be found by performing a least-squares fit of bilinear functions. Points that lie above and below this plane are identified as “bad pairs”. The largest distance is identified, and that pair is removed from the sets. The best-fit functions are re-computed, and the process is repeated until the largest distance is smaller than a threshold value. A threshold of $1/3$ pixel was used in most of the examples discussed below. The threshold has been increased, however, to as large as $4/3$ pixels to successfully register some of the more challenging X-radiographs⁵. In Fig. 5b, the blue points show control points that have been removed, and the red points show control points that are within the threshold distance to the best-fit function. The black points are the final best-fit function. In this orientation, it is clear that the red pairs have a consistent disparity relationship. If a substantial number of these pairs exist after the threshold has been met, there is reason to be confident that they can be used to produce an accurate transformation. For the examples described below, at least 64 pairs were required to produce

sufficient accuracy. By selecting the control-point regions more carefully, as described above, the method reduces the likelihood of obtaining an inaccurate transformation.

3.3 Generate the registered and mosaicked image

The pairs that remain after the filtering step are used to define a spatial transform that will bring the template image into alignment with the reference image. When transforming IR image frames and X-radiographs, a bilinear transformation is used, as described in [4, 21]. Once an independent transformation has been applied to each template image, a mosaic is formed. When forming the mosaic, the pixels in the overlap region are added to the mosaic based on their distances from the centers of their respective blocks. Since the distortion for any given block is most often greater toward the edge, we attempt to preserve the centers of the blocks and discard the edges when there is redundant information available. This is done by computing the distance for every pixel in the mosaic to the center of its respective block. When a new block is registered, the distances in the new block are computed and then compared to distances already in the mosaic. If a distance is smaller, then the corresponding pixel replaces the one already in the mosaic.

By capturing many small IR image frames (small relative to the color image), registering them to the color reference image, and forming a mosaic, we transform the resulting IR mosaic to approximately fit the distortion in the reference image. This type of fitting is a piecewise

⁵ X-radiographs are scanned at 500 pixels per inch, and thus the increase in disparity corresponds to a decrease in pixel size relative to the IR datasets.

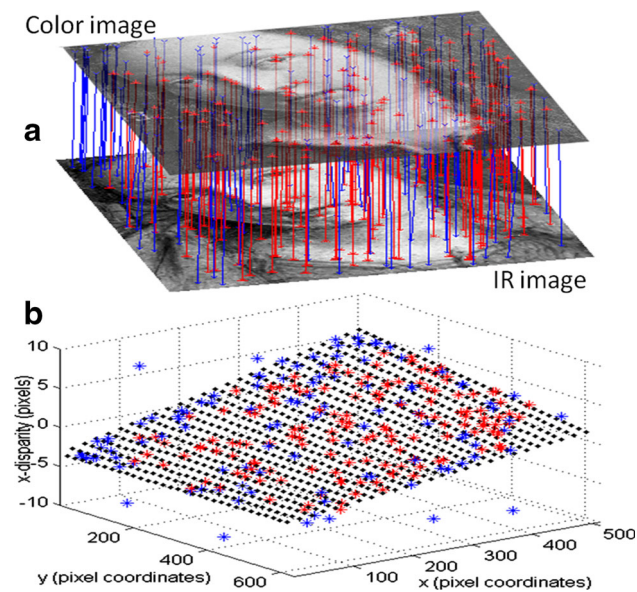


Fig. 5 **a** Matching control-point pairs between the color image (*top*) and infrared image (*bottom*) and **b** horizontal disparity map. *Blue* points and *lines* shows control points and pairs that have been removed, although they have good correlation peaks. *Red* points and *lines* show control points that are used in the transformation. The *black* points show the best-fit function. The points are fit to a bilinear function

approach. While a bilinear transformation cannot perfectly fit the distortion in the reference image, the independent transformation performed on each IR image frame (each piece) will better fit the local distortion. For example, as discussed above, a color camera capturing images at 565 pixels per inch can produce a shift of as much as 24 pixels at the edge of the image. But, by computing an independent bilinear transformation for each 512-by-640 pixel IR image frame, we constrain the maximum shift in the local region of the color image, relative to the spatial coordinates of the transformed IR image frame, to 0.72 pixels⁶.

4 Results—validation study

A painting was created to allow quantification of the accuracy of the registration and mosaicking algorithm. The painting contains an underdrawing, a painted surface layer, and cobalt blue lines that are used as fiducial markers. The algorithm then was used to generate an infrared (IR) mosaic that is in register with a reference color image (1280-by-1035 pixels, approximately 80 pixels per inch, 16 bits per pixel). The spatial distances between the fiducial

⁶ IR image frames are generally registered at a spatial sampling near 280 pixels per inch, so the maximum shift in that case would be closer to 0.36 pixels.

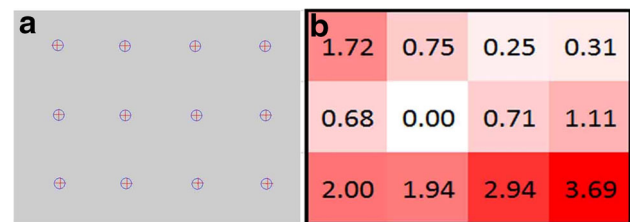


Fig. 6 **a** Extracted vertices of reference (*blue* circles) and easel-scanned cube (*red* crosses) grids, **b** distances, in pixels, between the reference and easel-scanned vertices

markers in the image and those in the mosaic then are used as a metric for the error of the registration.

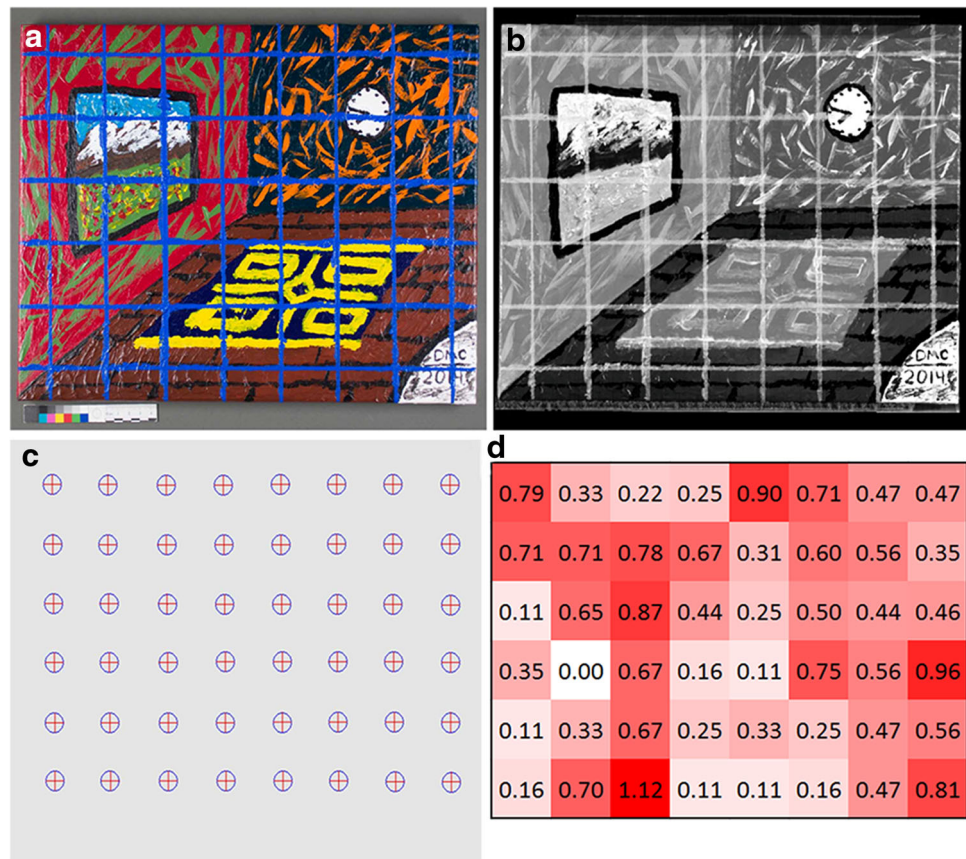
A hyperspectral camera [16], operating from the visible to near-infrared (VNIR), was used to collect a set of diffuse reflectance spectral image cubes of this painting. The set was captured using a high-precision (20 μm) computer-controlled easel to slowly move the painting in front of the camera. The easel produced a trigger that caused the camera to capture an image approximately every 300 μm , thus producing 300-by-300 μm pixels. The set initially consisted of two image cubes (512-by-1536 pixels) but was divided into six spatially overlapping cubes (512-by-640 pixels by 260 bands, 358–988 nm) prior to registration.

Prior to registration, the distortion of the first cube in the set was estimated. To do this, the central points of the intersections of blue lines for the color reference image (see Fig. 7a) and the 910-nm image was automatically extracted. This was done by filtering the reference image and portions of the cubes. The filter was constructed by intersecting vertical and horizontal Gaussian functions (SD = 2.5 pixels). For the reference image, the filtered red and green images were subtracted from the filtered blue image. For the VNIR sets, the filtered 499-nm image was subtracted from the filtered 399-nm image. The subtractions were performed to enhance the contrast of the blue lines relative to the background.

After filtering and computing the differences, the algorithm scaled up the resulting matrices by a factor of nine to increase the precision of the resulting intersection coordinates, using bicubic interpolation, and the local maxima were identified. Finally, a pixel region approximately 1.5 times the grid spacing (2169-by-2169 pixels) was shifted through each matrix, discarding all but the largest maxima in that region. In Fig. 6a, the vertices of the reference color image are shown as blue circles, and the vertices of the two original hyperspectral cubes are shown as red crosses. Fig. 6b shows the differences, in pixels, between the corresponding vertices for one easel-scanned cube.

The dataset was registered and mosaicked using the algorithm described above. The 910-nm image from each of the cubes was registered to the color reference image, shown in Fig. 7a, and mosaics were formed, each 1280-by-

Fig. 7 **a** Color reference image, **b** mosaic of easel-scanned 910-nm image, **c** extracted vertices of reference (*blue circles*) and mosaic (*red crosses*) grids, **d** distances, in pixels, between the reference and mosaic vertices



1035 pixels. The mosaic for the easel-scanned dataset is shown in Fig. 7b. The same technique as above was used to identify the vertices for the color reference image and IR mosaic and to compute the differences. The 48 vertices are shown in Fig. 7c. The vertices for the reference image are displayed as blue circles, and the vertices for the easel-scanned VNIR dataset are shown as red crosses. Fig. 7d shows the difference, in pixels, between the two sets of vertices. The mean error of the 48 fiducials for the easel-scanned dataset was 0.47 pixels (standard deviation of 0.27 pixels). To put this error measurement into context, the widths of the blue lines vary from 5 to 10 pixels throughout the images. While computationally intensive, the process has reasonable run times. For example, registering/mosaicking of 150 IR images (each 512-by-640 pixels) to a color image of a panel painting took approximately two hours on an 8-core desktop computer and produced a final IR mosaic of 4700-by-6000 pixels.

5 Application example—multispectral infrared images/X-radiographs

The algorithm was designed to allow for the registration of individual image frames, acquired from a near-infrared (NIR) imaging array and X-radiographs, to a reference

color image [4, 22]. In the first example, IR image frames taken of Jan van Eyck's *The Annunciation* (Fig. 8a) were registered. This painting was painted on a wood panel support (transferred to a coarsely woven canvas in the nineteenth century) and consists of an underdrawing, to plan the composition, over which an image was built up with layers of paint and transparent glazes [23]. The IR dataset consists of 220 partially overlapping IR image frames, taken in the 1100–1400-nm spectral range using an InSb area array [23]. The set was registered and mosaicked with a color image of similar resolution.

Comparison of the color image and the registered IR mosaic shows several regions where there are differences between the underdrawing and the final composition. One such region is in the figure of the Virgin. A detail from the registration results of that figure is shown in Fig. 8, where Fig. 8b is the detail color image and Fig. 8c is the corresponding region of the registered IR mosaic. The registration results clearly show compositional changes in the neckline of the Virgin's dress and a change in her hair. What is not clear without alternately viewing the overlaid color and IR images is that the Virgin's eyes, ear, and fingers have been repositioned. Online Resources 1–3 show animations of the overlay of the color and IR images of the Virgin's face, dress, and hand. As the animations alternate between color and IR, the objects that have changed

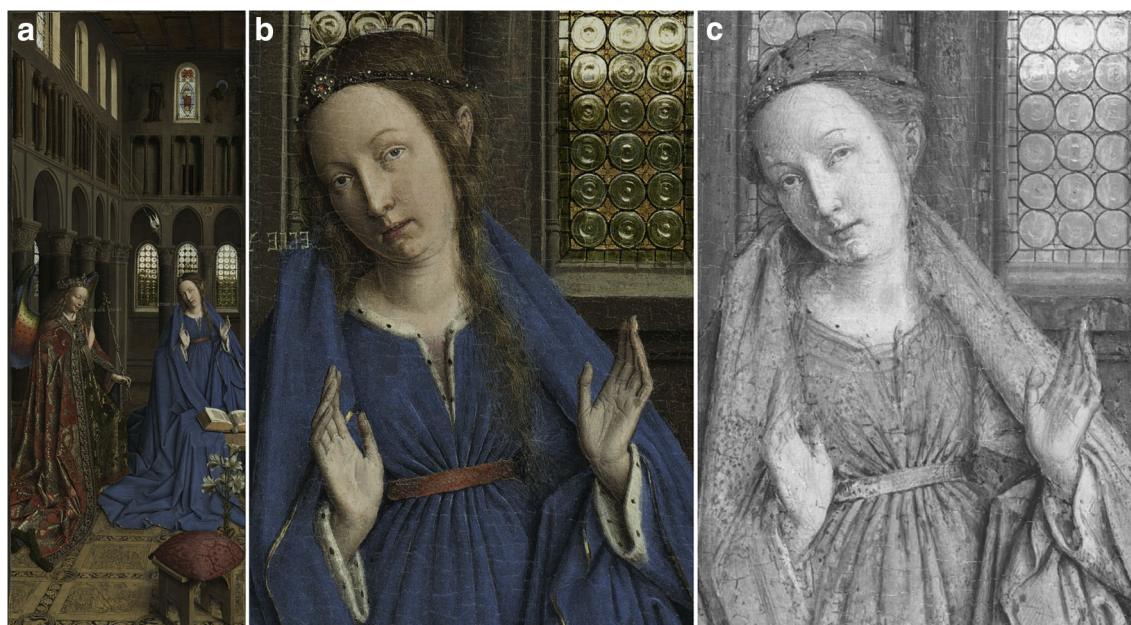


Fig. 8 Jan van Eyck's *The Annunciation* (1434/1436): **a** color image, **b** detail from color image, and **c** registered infrared (1100–1400 nm) image (see Online Resources 1–3). Andrew W. Mellon Collection, 1937.1.39, National Gallery of Art, Washington, D.C.

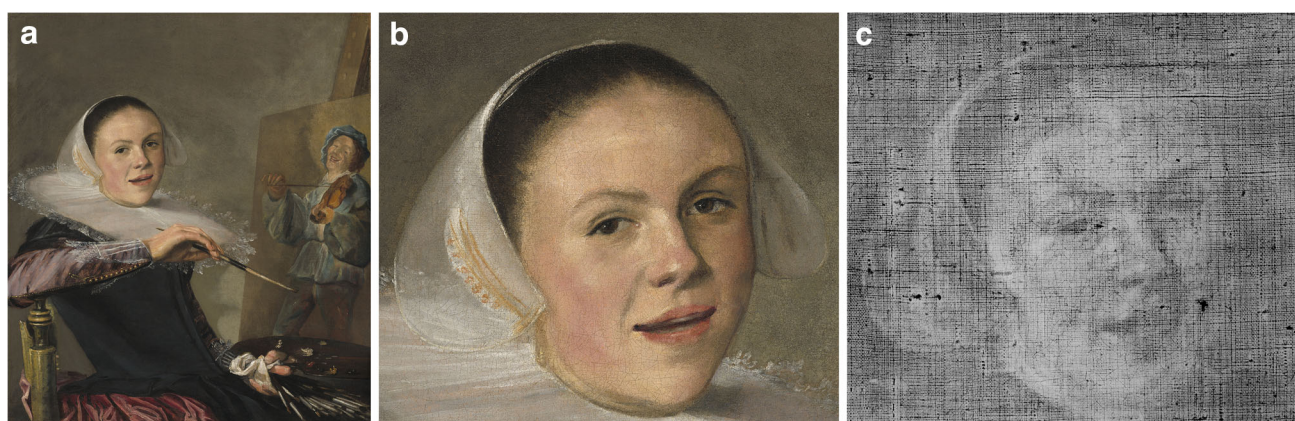


Fig. 9 Judith Leyster's *Self-Portrait* (1630): **a** color image, **b** detail from color image, and **c** registered X-radiograph (see Online Resource 4). Gift of Mr. and Mrs. Robert Woods Bliss, 1949.6.1, National Gallery of Art, Washington, D.C. [24]

between the underdrawing and final composition appear to be moving. Note, however, that in the animations other objects, such as cracks, remain stationary.

In the next example, four partially overlapping X-radiographs of Judith Leyster's *Self-Portrait* (Fig. 9a) were scanned at 500 pixels per inch. The scanned X-ray images then were divided into a series of small subimages to minimize the effect from the difference of the imaging geometries between them and the color image (this permits a bilinear transformation to be used in the registration), and the subimages were registered with a color image. These X-ray films contain a repetitive pattern due to the painting being executed on relatively coarse canvas, as shown in

Fig. 9c. This pattern adds an additional challenge in the registration process, because it dominates the X-ray images but is not visible in the final composition. One method to minimize the contribution of the weave would be to filter the image. In this algorithm, however, by filtering the control-point pairs instead, it is generally not necessary to take this step. What remain after our filtering step are the control-point pairs, related to texture and cracks, that are visible in both the color and X-ray images. In this image, the remaining control-point pairs are sufficient to produce well-registered images. Online Resource 4 shows an animation of the overlay of color and X-ray detail images (Fig. 9b–c). As the animation alternates between color and

X-ray the large degree of difference between the images is apparent, while the success of the registration can be seen in the close alignment of the cracks.

6 Conclusions

We have described an automatic registration and mosaicking algorithm that is capable of registering images acquired using various imaging modalities that are used in the cultural heritage field. The algorithm is capable of performing registration while accommodating various forms of optical distortion. To date, the algorithm has been used to register and produce mosaicked IR images and X-radiographs of more than 100 paintings and works on paper. This includes many datasets from the National Gallery of Art's (Washington, D.C.) online catalogue for its collection of seventeenth century Dutch Paintings [25] and the Art Institute of Chicago's digital collections: Monet Paintings and Drawings [26] and Renoir Paintings [27].

Acknowledgments The authors acknowledge funding from the National Science Foundation (award 1041827). J.K.D. and D.M.C. acknowledge funding from the Andrew W. Mellon and Samuel H. Kress Foundations, and D.M.C. acknowledges funding from the ARCS Foundation. The authors are grateful to Kimberly Schenk, Elizabeth Walmsley, Catherine Metzger, Melanie Gifford, Barbara Berrie, and Mervin Richards of the National Gallery of Art (Washington, D.C.), and Kelly Keegan of the Art Institute of Chicago.

Conflict of interest The authors declare that they have no conflicts of interest.

References

1. E. Walmsley, C. Metzger, J. Delaney, C. Fletcher, *Stud. Conserv.* (1994). doi:10.1179/sic.1994.39.4.217
2. J. Delaney, E. Walmsley, B. Berrie, C. Fletcher, (*Sackler NAS Colloquium*) *Scientific Examination of Art: Modern Techniques in Conservation and Analysis* (National Academies Press, Washington, D.C., 2005)
3. A. Wheelock Jr. Johannes Vermeer/Girl with the Red Hat/c. 1665/1666 (NGA Online Editions, 2014), <http://purl.org/nga/col/lection/artobject/60>. Accessed 21 May 2014
4. D. Conover, J. Delaney, P. Ricciardi, M. Loew, *Proc. SPIE* (2011). doi:10.1117/12.872634
5. K. Martinez, J. Cupitt, *IEEE Proc.* (2005). doi:10.1109/ICIP.2005.1530120
6. P. Ravines, K. Baum, N. Cox, S. Welch, M. Helguera, *J. Am. Inst. Conserv.* (2014). doi:10.1179/1945233013Y.0000000014
7. The Bosch Research and Conservation Project (BRCP), Image processing and digital infrastructure (BRDP, 2014), <http://boschproject.org/research.html#digital>. Accessed 21 May 2014
8. D. Bertani, M. Cetica, P. Poggi, G. Puccioni, E. Buzzegoli, D. Kunzelman, S. Cecchi, *Stud. Conserv.* (1990). doi:10.1179/sic.1990.35.3.113
9. C. Daffara, E. Pampaloni, L. Pezzati, M. Barucci, R. Fontana, *Acc. Chem. Res.* (2010). doi:10.1021/ar900268t
10. D. Saunders, H. Liang, R. Billinge, J. Cupitt, N. Atkinson, *Stud. Conserv.* (2006). doi:10.1179/sic.2006.51.4.277
11. Opus Instruments, Specification of Osiris infrared reflectography cameras (Opus Instruments, 2014), <http://www.opusinstruments.com/specification/>. Accessed 21 May 2014
12. M. Alfeld, J. Pedroso, M. van Eikema, Hommes, G. Vander Snickt, G. Tauber, J. Blass, M. Haschke, K. Erler, J. Dik, K. Janssens, *J. Anal. At. Spectrom.* (2013). doi:10.1039/C3JA30341A
13. J. Delaney, P. Ricciardi, L. Glinsman, M. Facini, M. Thoury, M. Palmer, E.R. de la Rie, *Stud. Conserv.* (2014). doi:10.1179/2047058412Y.0000000078
14. K. Dooley, S. Lomax, J. Zeibel, C. Miliari, P. Ricciardi, A. Hoenigswald, M. Loew, J. Delaney, *Analyst* (2013). doi:10.1039/c3an00926b
15. P. Ricciardi, J. Delaney, M. Facini, J. Zeibel, M. Picollo, S. Lomax, M. Loew, *Angew. Chem. Int. Ed.* (2012). doi:10.1002/anie.201200840
16. K. Dooley, D. Conover, L. Glinsman, J. Delaney, *Angew. Chem. Int. Ed.* (2014). doi:10.1002/anie.201407893
17. L. Brown, *ACM Comput. Survey.* (1992). doi:10.1145/146370.146374
18. B. Zitová, J. Flusser, *Image Vis. Comput.* (2003). doi:10.1016/S0262-8856(03)00137-9
19. G. Hong, Y. Zhang, *Comput. Geosci.* (2008). doi:10.1016/j.cageo.2008.03.005
20. R. Gonzalez, R. Woods, *Digital Image Processing*, 3rd edn. (Pearson Prentice Hall, New Jersey, USA, 2008)
21. G. Wolberg, *Digital ImageWarping* (IEEE Computer Society Press, Washington, D.C., USA, 1990)
22. D. Conover, J. Delaney, M. Loew, *Proc. SPIE* (2013). doi:10.1117/12.2021318
23. E.M. Gifford, C. Metzger, J. Delaney, *Facture* **1**, 128–153 (2013)
24. A. Wheelock Jr., Judith Leyster/Self-portrait/c. 1630, (NGA Online Editions, 2014), <http://purl.org/nga/collection/artobject/37003>. Accessed 21 May 2014
25. A. Wheelock Jr., Dutch paintings of the seventeenth century (NGA Online Editions, 2014), <http://purl.org/nga/collection/catalogue/17th-century-dutch-paintings>. Accessed 21 May 2014
26. G. Groom, J. Shaw, eds., Monet paintings and drawings at the Art Institute of Chicago (Art Institute of Chicago, 2014), <http://publications.artic.edu/reader/monet-paintings-and-drawings-art-institute-chicago>. Accessed 21 May 2014
27. G. Groom, J. Shaw, eds., Renoir paintings and drawings at the Art Institute of Chicago (Art Institute of Chicago, 2014), <http://publications.artic.edu/reader/renoir-paintings-and-drawings-art-institute-chicago>. Accessed 21 May 2014