Signal Processing 93 (2013) 592–604

Contents lists available at SciVerse ScienceDirect

Signal Processing

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

Restoration of X-ray fluorescence images of hidden paintings

Anila Anitha [a], Andrei Brasoveanu [f], Marco Duarte [b], Shannon Hughes [a,*], Ingrid Daubechies [c], Joris Dik [d], Koen Janssens [e], Matthias Alfeld [e]

[a] Department of Electrical, Computer, and Energy Engineering, University of Colorado, Boulder, Boulder, CO 80309, USA
[b] Department of Electrical and Computer Engineering, University of Massachusetts, Amherst, Amherst, MA 01003, USA
[c] Department of Mathematics, Duke University, Durham, NC 27708, USA
[d] Department of Materials Science and Engineering, TU Delft, Delft, The Netherlands
[e] Department of Chemistry, University of Antwerp, Antwerp, Belgium
[f] Department of Mathematics, Princeton University, Princeton, NJ, 08544, USA

ARTICLE INFO

Article history:
Received 13 October 2011
Received in revised form
27 September 2012
Accepted 1 October 2012
Available online 13 October 2012

Keywords:
Art restoration
Image restoration
Artifact correction
Underdetermined source separation

ABSTRACT

This paper describes our methods for repairing and restoring images of hidden paintings (paintings that have been painted over and are now covered by a new surface painting) that have been obtained via noninvasive X-ray fluorescence imaging of their canvases. This recently developed imaging technique measures the concentrations of various chemical elements at each two-dimensional spatial location across the canvas. These concentrations in turn result from pigments present both in the surface painting and in the hidden painting beneath. These X-ray fluorescence images provide the best available data from which to noninvasively study a hidden painting. However, they are typically marred by artifacts of the imaging process, features of the surface painting, and areas of information loss. Repairing and restoring these images thus consists of three stages: (1) repairing acquisition artifacts in the dataset, (2) removal of features in the images that result from the surface painting rather than the hidden painting, and (3) identification and repair of areas of information loss. We describe methods we have developed to address each of these stages: a total-variation minimization approach to artifact correction, a novel method for underdetermined blind source separation with multimodal side information to address surface feature removal, and two application-specific new methods for automatically identifying particularly thick or X-ray absorbent surface features in the painting. Finally, we demonstrate the results of our methods on a hidden painting by the artist Vincent van Gogh.

1. Introduction

In recent years, imaging and digital image processing have increasingly been used to aid in art restoration and reconstruction. The primary motivation is that imaging can be applied without harm to a work of art, and digital image processing techniques can then be used on the resulting data in order to enhance features of interest, remove obstructive artifacts, repair damages, virtually undo the effects of aging, and/or merge data from multiple imaging methods (e.g. from visible, infrared, and ultraviolet light) to create a virtual restoration or reconstruction. The resulting virtual restoration can be valuable to art historians and the general public in its own right, or it can provide guiding information for art conservators intending to perform a physical restoration. See [1–3] for a review of this emerging area.

As examples of what has been achieved so far, such work may find and virtually repair specific areas of damage, such as cracks in the painting surface [4–12] or

* Corresponding author. Tel.: +1 303 492 7038; fax.: +1 303 492 2758.
E-mail address: smhughes@colorado.edu (S. Hughes).

0165-1684/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.sigpro.2012.09.027

lacunae [13,14]. It may alternatively virtually undo the effects of aging on the materials that the work contains, for example, restoring fading pigments to their original, more vibrant, colors [15–18] or virtually removing a layer of darkened varnish from the surface [19,11,12]. In extreme cases, small fragments of a destroyed work may be pieced back together [20,21] or information, e.g. color, about large missing sections of the work may be estimated from a few small surviving fragments [22]. In addition to paintings, many other types of artwork, including aging photographs [23], film sequences [24,25], frescoes [26,22], daguerrotypes [27], and mosaics [28] are treated as well. Indeed, some researchers have released suites of image processing tools to be used by art scholars so that they may conduct virtual restoration themselves [29,30].

One problem for which this strategy of non-invasive imaging combined with digital image processing is particularly well-suited, but that, to our knowledge, has not been previously investigated, is that of virtually reconstructing hidden paintings, i.e. paintings that have been painted over. Since these hidden paintings often lie beneath priceless works of art, they may only be reached via non-invasive imaging. Many such hidden paintings exist: for example, a recent X-ray analysis of 130 Van Gogh paintings at the Van Gogh Museum in Amsterdam showed that almost 20 of the 130, i.e. roughly 15%, contained some sort of hidden painting [31]. In some cases, these hidden paintings are finished works by the artist; in other cases, they may be preliminary studies for later projects. While some previous work has been done to enhance X-ray images of paintings [32], imaging of underdrawings [33] (e.g. preliminary sketches made on the canvas before painting), or overwritten texts such as are found in palimpsests [34], we are not aware of any prior work that attempts to restore images of hidden paintings themselves, as separate from the surface painting. Quality images of these hidden paintings can potentially aid art historical scholarship about the painter and his/her work. More broadly, they can also provide a look at these important lost pieces of our cultural heritage for the general public.

To reconstruct these hidden paintings, however, we must start with high quality imaging data of them. A number of non-invasive imaging methods have historically

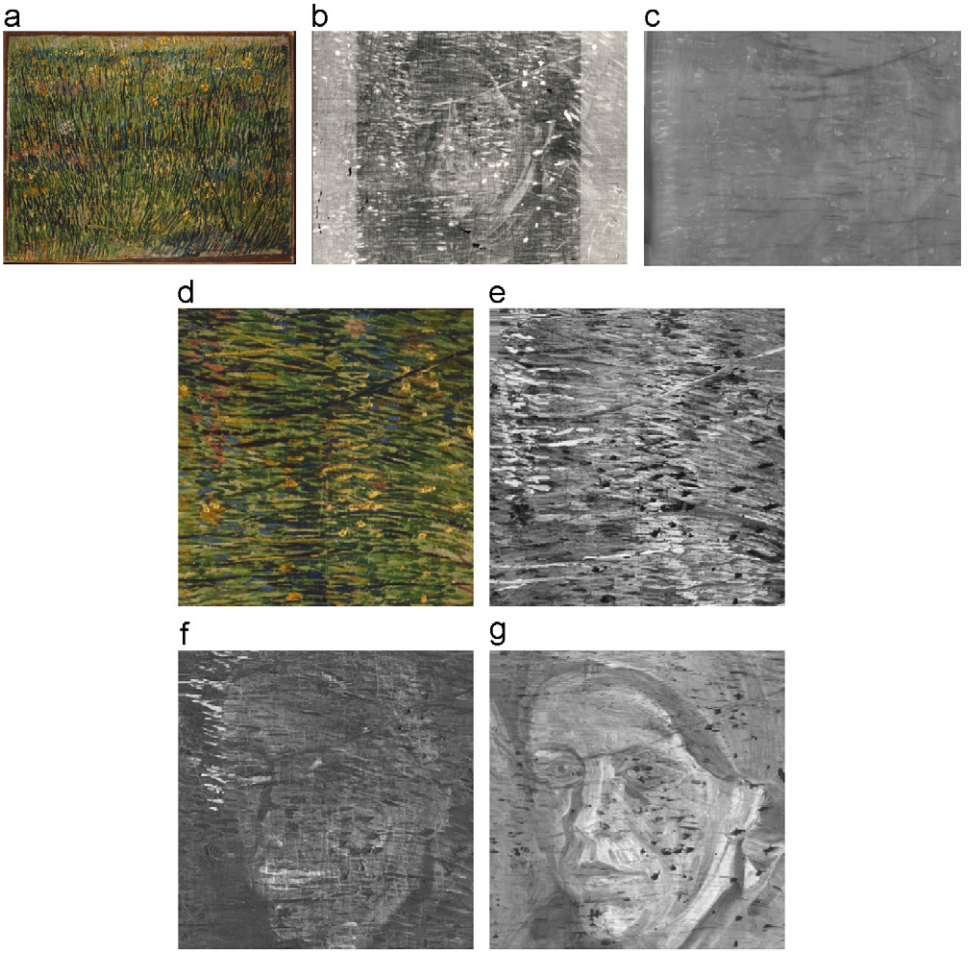

Fig. 1. Comparison of different types of non-invasive imaging on Vincent van Gogh's "Patch of Grass" including (a, d) the surface painting under visible light, (b) traditional X-ray imaging and (c) infrared imaging of portions of the painting, and (e)–(g) three different chemical channel images of the painting produced via the synchrotron X-ray fluorescence imaging technique of [37].

been used to gain information about underlayers of the painting, the most common of which are classic X-ray photography, infrared imaging, and multispectral imaging. X-ray photography produces a single image of the canvas showing the combined X-ray absorbance of all layers in the painting (see Fig. 1(b)). Infrared imaging uses infrared light to penetrate the top layers of the painting and reveal infrared-reflective/absorbent features beneath the surface (see Fig. 1(c)). It is particularly successful at imaging any underdrawings the artist made on the canvas before painting [35]. Multispectral imaging combines images under several different wavelengths of visible, infrared, and perhaps ultraviolet light [36]. However, the data these can provide is limited, as can be seen in Fig. 1.

Hence, new types of imaging are being developed. In 2008, a team headed by coauthors Dik and Janssens developed a novel synchrotron-based X-ray fluorescence imaging method [37]. This non-invasive technique, when performed on a painting, provides a set of images showing the spatial distributions of specific chemical elements, including As, Ba, Cd, Co, Cr, Cu, Fe, Hg, Mn, Pb, Sb, Sr, and Zn. See Fig. 1(e)–(g) for examples. Furthermore, a more recent portable imaging method developed by coauthor Alfeld et al. [38] also uses X-ray fluorescence to produce chemical channel images, but eliminates the need for a synchrotron, allowing museums to produce (slightly noisier) images of this type in-house. This information regarding various chemical elements, as compared to other methods, provides the richest available data to use in studying a hidden painting.

However, even these new X-ray fluorescence images can come with significant artifacts and damages. It is our goal in this paper to develop image processing methods that can restore and repair these damages and artifacts in the acquired images. Specifically, three types of damages/artifacts are treated. First, as in any imaging process, this particular imaging technique can produce its own unique type of imaging artifacts: a set of pixel shifts, line by line, in the acquired image. This is illustrated in Fig. 2. We will aim to shift these pixels back to where they belong. Second, features of the surface and hidden paintings containing the same pigment are mixed together in the corresponding chemical element's image. An example can be seen in the mercury (Hg) channel shown in Fig. 1(f) where both the woman's lips from the hidden painting and a group of pink flowers from the surface painting (upper left) are visible. In order to obtain images that only reflect the hidden painting, we will aim to remove the surface painting features from such mixed images. Third, surface painting features can block signals from the

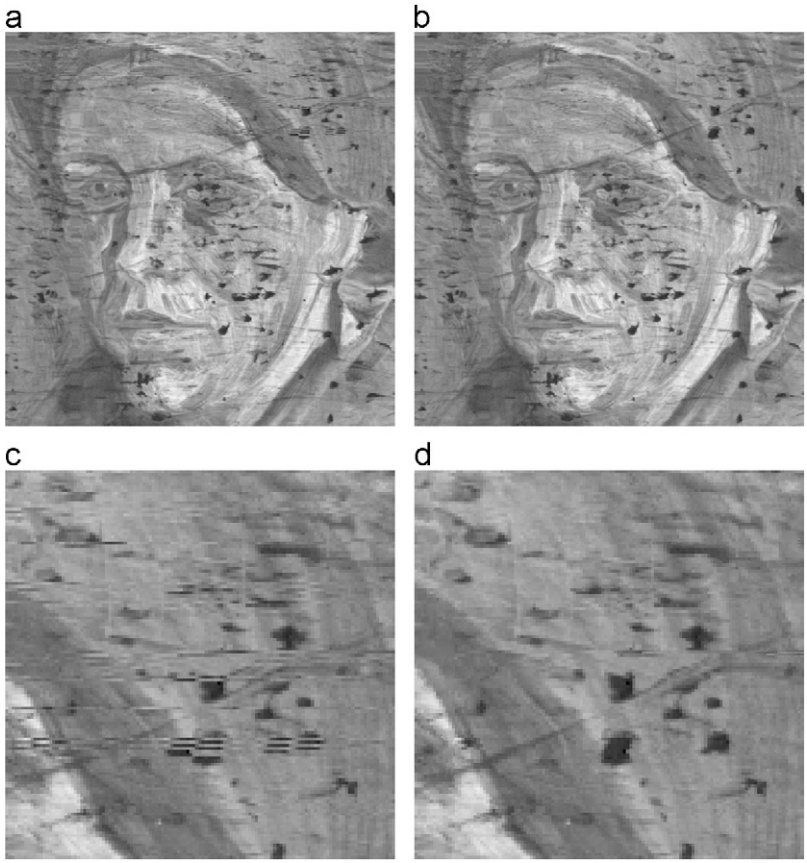

Fig. 2. Results of our total variation minimization technique for shifting line artifact correction. (a) Original acquired antimony channel image showing line alignment artifacts. (b) Corrected antimony channel image after applying our total variation minimization technique. The line alignment artifacts have almost all been repaired by our method. (c) Close-up of section of original acquired antimony image in (a), illustrating the line alignment artifacts. (d) Close-up of the same section of the corrected image.

underlayers from reaching the surface to be imaged, creating areas of information loss in the chemical element images. This can be observed in the similar pattern of black spots in each of Fig. 1(e)–(g). In some areas, the loss is complete; in others, the signal has merely been attenuated. We will aim to identify these areas of loss or attenuation, inpainting the areas of complete loss, and boosting signal strength in areas of attenuation.

In addition to serving the hidden painting recovery application, each of these methods represents a novel contribution to signal processing methodology as well. First, the line-shifting artifacts seen are unique in form to this application, distinct from anything previously seen in the literature, and hence require a tailor-made new approach. Second, while the problem of separating surface and hidden painting features in the images could be considered as an underdetermined blind source separation task and treated with existing approaches, to do so would be to fail to take advantage of all the side information that is available in our particular situation. RGB color photographs of one source, the surface painting, are readily available and provide a rich source of additional information that can be used to enhance the source separation result. However, since this side information is taken using a different imaging modality (color photography) than that used for the two sources (which are chemical element images), it is nontrivial to use this side information intelligently. To take advantage of this side information, we have thus developed a new method for underdetermined source separation with different-modality side information. This problem is potentially found in other source separation contexts, but has not to our knowledge been treated before in the literature, so our solution to it may potentially contribute to other application domains as well. Finally, while inpainting has been well-studied in the literature and hence no new inpainting methods are developed here, our methods for identifying where losses are likely to have occurred, using image analysis of the surface painting features, are new contributions. We provide two such novel loss-identification methods.

This paper is organized as follows. Section 2 will discuss our method for correcting the line shifting artifacts. Section 3 will detail our method for removing surface features from the images so that only the hidden painting is reflected in the resulting image. Section 4 will describe our methods for identifying and inpainting areas of information loss. Throughout, we will be demonstrating our methods on an X-ray fluorescence imaging dataset of the painting "Patch of Grass" by the artist Vincent van Gogh, which has an unknown woman's portrait beneath it (see Fig. 1). We believe that these methods together will allow for automated repair and restoration of X-ray fluorescence images of hidden paintings, improving their legibility for both art scholars and the general public. (Portions of these methods were described in our previous conference papers on this work; see [39,40].)

2. Scanning artifact correction

This section describes the first of three stages of hidden painting restoration that we propose in this paper: correction of scanning line artifacts. Chemical element images obtained through the X-ray fluorescence imaging technique may contain horizontal lines or portions thereof whose pixels are shifted horizontally left or right with respect to the lines above and below (see Fig. 2). This shifting behavior is the result of a timing problem in the acquisition. The analysis is performed with the aid of a scanning X-ray beam that irradiates the object pixel by pixel, line by line horizontally, switching direction with each line [37]. Occasionally, this scanning beam gets delayed, which results in some of the pixels getting captured later than scheduled and thus appearing shifted in their respective horizontal lines. In each distorted line, the amount of shift increases/decreases monotonically, although the direction varies by line. All chemical element images are acquired in one scan, so shifts are identical across all of them. Examination of the portrait under "Patch of Grass" reveals that approximately 70–80 of the 698 horizontal lines ($\sim 10\%$) are damaged, with a fairly uniform distribution through the entire picture.

For each horizontal line, we model the shift on that line, i.e. the displacement of each pixel as a function of horizontal position, as a piece-wise constant function with at most d discontinuities. The shift function $S_{l,a}(n)$ thus has domain $\{1,\ldots,L\}$, where L is the length of each horizontal line in pixels, and range $\{-M,\ldots,M\}$, where M is the maximum allowable shift (we chose $M=25$). It has as parameters the locations $l_1,l_2,\ldots,l_d$ where a change in the shift size occurs due to a timing problem, and $a_1,a_2,\ldots,a_d$, which represent the amount of shifting introduced at each. The resulting piecewise constant function is then given by

$$S_{l,a}(n)=\begin{cases} 0, & 0\le n<l_1 \\ a_1, & l_1\le n<l_2 \\ a_1+a_2, & l_2\le n<l_3 \\ \quad\vdots & \quad\vdots \\ \sum_{k=1}^{d} a_k, & l_d\le n\le L \end{cases}$$

We found that allowing up to $d=2$ discontinuities in each line was enough to capture most of the occurring shift patterns.

We then determine the appropriate parameters $\{l_j,a_j\}_{j=1}^{d}$ through total variation minimization. The rationale is that introducing artificial horizontal line shifts into an otherwise normal image will reduce vertical continuity in the image and increase the total variation as measured in the vertical direction. We thus minimize the variation, as measured vertically, as a strategy for correcting these artifacts. Once the parameters are determined for each line, we can use them to repair the line, by shifting pixels back to their (supposed) original positions.

Our method has two variants, titled *single* and *dual* [41]. In the single method, the parameters are chosen to minimize the total variation of the tested line with respect to the line above. In the dual method, the parameters are chosen to minimize the total variation of the tested line with respect to both the lines above and below.

For the single method, this means that letting g_i and f_i be the values of the ith pixel of the line in question and the line

above it, respectively, we estimate the shift parameters as

$$\underset{l_1,\ldots,l_d,a_1,\ldots,a_d}{\operatorname{argmin}} \sum_{p=1}^{\min(L,L-\sum_{k=1}^{d} a_k)} |f_p - g_{p+S_{l,a}(p)}|^2 = \underset{l_1,\ldots,l_d,a_1,\ldots,a_d}{\operatorname{argmin}} \sum_{n=1}^{d+1} \sum_{j=l_{n-1}}^{l_n-1} |f_j - g_{j+\sum_{k=1}^{n-1} a_k}|^2$$

where we have set $l_0 = 1$, $l_{d+1} = \min(L, L - \sum_{k=1}^{d} a_k)$. We further require that $l_n - 1 + \sum_{k=1}^{n-1} a_k \le L$ for all n and $l_{n-1} + \sum_{k=1}^{n-1} a_k \ge 1$ for all n, so that the sum in Eq. (1) remains defined. In the case of $d=2$, this simplifies to

$$\underset{l_1,l_2,a_1,a_2}{\operatorname{argmin}} \sum_{j=1}^{l_1-1} |f_j - g_j|^2 + \sum_{j=l_1}^{l_2-1} |f_j - g_{j+a_1}|^2 + \sum_{j=l_2}^{\min(L,L-a_1-a_2)} |f_j - g_{j+a_1+a_2}|^2$$

Each line is corrected before moving on to the next. However, to avoid the correction of spurious shifts, only shifts whose pixel size in absolute value $|a_k|$ is greater than the threshold T (we used $T=5$) are actually corrected.

For the dual method, we perform a similar minimization, but incorporate both the pixel values of the line above f_i^a and those of the line below f_i^b into the minimization:

$$\underset{l_1,\ldots,l_d,a_1,\ldots,a_d}{\operatorname{argmin}} \sum_{n=1}^{d+1} \sum_{j=l_{n-1}}^{l_n-1} |f_j^a - g_{j+\sum_{k=1}^{n-1} a_n}|^2 + |f_j^b - g_{j+\sum_{k=1}^{n-1} a_n}|^2$$

where l_0 and l_{d+1} are again defined as above. No threshold is used here since the dual method seems to suffer less from spurious small shifts. In the $d=2$ case, this becomes

$$\underset{l_1,l_2,a_1,a_2}{\operatorname{argmin}} \sum_{j=1}^{l_1-1} (|f_j^a - g_j|^2 + |f_j^b - g_j|^2) + \sum_{j=l_1}^{l_2-1} (|f_j^a - g_{j+a_1}|^2 + |f_j^b - g_{j+a_1}|^2) + \sum_{l=l_2}^{\min(L,L-a_1-a_2)} (|f_j^a - g_{j+a_1+a_2}|^2 + |f_j^b - g_{j+a_1+a_2}|^2)$$

When using this method, we estimate the correction for each horizontal line based on the lines above and below it, but save this information for later, rather than correcting the line immediately. Only after we finish estimating the corrections for *all* lines in the image do we proceed to correct all the found shift patterns. However, we may find afterwards that lines are still misaligned since each line's correction has been based on the original positions of the lines above and below, which may have changed in the meantime. Hence, the method is applied iteratively, with the repaired channel hopefully converging to the true one.

Since the single method assumes the top horizontal line in each image to be correct and then corrects each subsequent line once based on the one above, this method is fast and works well when the erroneously shifted lines are uniformly sparse across the image, but problems can arise when an area of the image has many consecutive lines with shifts, because inexact fixes can propagate to the lines below. Meanwhile, the more intensive iterative dual method seems to be more robust and can be used in areas of high density acquisition errors without worry that errors will propagate from one line to the next. Hence, we use a combination of the two methods. We use the single method with $T=5$ if the image area is not particularly dense with shifts. However, in images or areas that are dense with shifts, we use the single method with a threshold of $T=15$ first, followed by 30 iterations of the dual method. See Fig. 2 for the result of this approach on the "Patch of Grass" hidden painting. As can be seen, this strategy corrected all but a few of the shift patterns in the image. Hence, it seems to be a successful method for artifact repair in X-ray fluorescence images.

3. Removal of surface features

We now move on to the second of three stages of hidden painting restoration that we propose in this paper: removal of surface painting features from the images. As explained previously, each chemical element image obtained through X-ray fluorescence imaging gives the total amount of the element at each x–y location in the canvas, without regard to the element's depth under the surface of the painting. Each chemical channel image is thus a mixture of both hidden painting and surface painting features containing this element. However, since only the hidden painting is of interest to us, we would ideally like to remove those features corresponding to the surface painting from each chemical channel image, leaving only the features of the hidden painting behind.

The problem of separating surface and hidden painting features in each chemical channel image can be viewed as a source separation problem. We wish to split each chemical channel image C into two, an image C_{hidden} that contains only the hidden painting's features and an image C_{surf} containing only the surface painting's features. In the absence of other information, this might be considered as a typical underdetermined blind source separation problem [42–44] and be treated as such.

However, in this situation, we have additional information that we would like to take advantage of. Color visible light photographs of the surface painting can easily be taken and used to provide additional information to aid in the source separation. This side information, while potentially useful, is nontrivial to incorporate since it comes from a different type of imaging than the sources themselves. While the sources show concentrations of chemical elements, a color photograph of the surface painting shows only the RGB color of the final pigment mixture. Hence, pigments show up very differently in the chemical element and color visible light images. The problem is thus one of underdetermined source separation with different-modality side information.

While one can easily imagine this problem arising in other fields, we are not aware of any previous work that tackles this particular type of problem. Hence, in this section, we will propose a new method for this type of source separation.

In solving this problem, we are inspired by the metrics for measuring relatedness between images acquired through different imaging modalities that are popular in the multimodal image registration literature (see e.g. [45,46]). In this prior work, the goal is to align two images of the same subject from different imaging modalities (e.g. a color and an infrared image of the same object). Quality of alignment is evaluated via a *similarity metric* between the two images that is designed to reflect similarity of image content, while being robust to differences of modality. Most commonly, each pair of pixel values obtained by drawing a value from the same location in each of the two images is modeled as an independent realization of a pair of unknown random variables. Entropic measures, such as the (negative of) estimated joint entropy or the estimated mutual information for this pair of random variables, are then used as the similarity metric between the images. The hypothesis underlying these metrics is that, even if they come from different imaging modalities, images are most likely to be of the same subject if the corresponding pixel values in the two images are highly dependent.

For example, consider the case of attempting to align a grayscale $M \times N$ 256-grayscale-level image I_1 (perhaps representing an infrared image) with an $M \times N$ RGB color image I_2. To estimate the entropy of the underlying random variable associated with I_1, we can define $p_{I_1}(b)$, our estimate of the probability of this underlying random variable taking on the value b, as the fraction of pixels in I_1 taking this value

$$p_{I_1}(b) = \frac{\left|\{(m,n)|I_1(m,n)=b\}\right|}{MN}$$

where the notation $|\cdot|$ represents set cardinality. We would then estimate the entropy of the random variable underlying image I_1 as usual via

$$H(I_1) = -\sum_{b\in\mathbb{Z}_{256}} p_{I_1}(b)\log(p_{I_1}(b))$$

Similarly, we would estimate the conditional entropy of the random variable underlying image I_1 given the random variable underlying image I_2 by defining

$$p_{I_2}(b_2) = \frac{\left|\{(m,n)|I_2(m,n)=b_2\}\right|}{MN} \quad \text{for all } b_2 \in \mathbb{Z}_{256}^3$$

$$p_{I_1,I_2}(b_1,b_2) = \frac{\left|\{(m,n)|I_1(m,n)=b_1, I_2(m,n)=b_2\}\right|}{MN}$$

for all $b_1 \in \mathbb{Z}_{256}, b_2 \in \mathbb{Z}_{256}^3$

and then letting

$$H(I_1|I_2) = -\sum_{b_1\in\mathbb{Z}_{256}}\sum_{b_2\in\mathbb{Z}_{256}^3} p_{I_1,I_2}(b_1,b_2)\log\frac{p_{I_1,I_2}(b_1,b_2)}{p_{I_2}(b_2)}$$

We might then expect that out of all possible transformations (e.g. translations, scalings, rotations, etc.) $T(I_1)$ of I_1, the one that is optimally aligned with image I_2 is the one that minimizes $H(T(I_1)|I_2)$ (or maximizes the mutual information $H(T(I_1))-H(T(I_1)|I_2)$.

Inspired by this previous work, we propose that after our source separation procedure, we would expect that the pixel values of the "surface" source C_{surf} should be highly dependent on those in the corresponding location in the color surface image I_{surf}. Meanwhile, the pixel values of the "hidden painting" source C_{hidden} should be completely independent of those of the color surface image I_{surf}. An initial attempt at an algorithm would thus try to encourage $H(C_{hidden}|I_{surf})$ to be large and $H(C_{surf}|I_{surf})$ to be small.

However, it turns out that such a choice leads to a trivial solution. By setting $C_{hidden} = C$ and $C_{surf} = 0$, the algorithm can maximize $H(C_{hidden}|I_{surf})$ while keeping $H(C_{surf}|I_{surf})$ at 0, simply as a result of maximizing the entropy of C_{hidden} while minimizing that of C_{surf}. For this reason, we instead choose to work with a normalized conditional entropy function; instead of aiming to maximize $H(C_{hidden}|I_{surf})$, we aim to maximize the normalized conditional entropy $H(C_{hidden}|I_{surf})/H(C_{hidden})$, which reflects only the proportion of C_{hidden}'s entropy explained by I_{surf}, rather than the overall total entropy of C_{hidden}. This choice incorporates the same intuition that one source should be dependent on the surface image and the other independent of it, while avoiding the trivial solutions obtained by merely maximizing $H(C_{hidden})$ while minimizing $H(C_{surf})$.

Finally, we also try to ensure that both sources after source separation behave like smooth real-world images with sharp edges by encouraging each to have a low total variation. The total variation of an $M \times N$ image I is typically estimated as

$$TV(I) = \sum_{m=1}^{M-1}\sum_{n=1}^{N-1}\sqrt{(I(m+1,n)-I(m,n))^2+(I(m,n+1)-I(m,n))^2}$$

To summarize, for each chemical element, this translates into minimizing over C_{surf} and C_{hidden}:

$$\frac{H(C_{surf}|I_{surf})}{H(C_{surf})} - \frac{H(C_{hidden}|I_{surf})}{H(C_{hidden})} + \lambda TV(C_{hidden}) + \lambda TV(C_{surf}) \tag{1}$$

subject to the constraint

$$C = C_{surf} + C_{hidden}$$

where C is the original chemical element image, I_{surf} is the RGB-color visible light surface image, C_{surf} and C_{hidden} are the two sources we are splitting into, $\lambda \in \mathbb{R}$ is a positive constant, and the functions $H(\cdot)$, $H(\cdot|\cdot)$, and $TV(\cdot)$ are as defined above.

We note that, as in most image processing applications involving minimization or maximization of an estimated entropy (see e.g. [45,46]), it is unnecessary to obtain a completely accurate and unbiased entropy estimate. Since we only aim to find values of the arguments resulting in minimum (maximum) entropy, the entropy estimator only needs to serve as a monotonic function of the true entropy to ensure that the correct solution is found.

We solve this optimization problem via gradient descent, but we perform this gradient descent on the wavelet coefficients of C_{surf} and C_{hidden} which helps to maintain coherent image features and avoid pixelwise overfitting. We use a multiresolution approach, allowing the wavelet coefficients at the coarsest scales to converge before starting to work on the wavelet coefficients of the next finest scale as well. The two sources are each initialized to $\frac{1}{2}C$.

If desired, standard histogram estimators (see [47]) can be used to more coarsely estimate the entropy values, speeding the computation.

Fig. 3 shows the result on this procedure on a synthetically generated toy example. Here, the classic "peppers" and "baboon" images have been mixed to create a model chemical image. An RGB color "peppers" image models side information of a different modality. We see that this method seems to produce a reasonably good separation of the two underlying sources. Indeed, the mean-squared error between the recovered sources and the true "ground truth" sources is found to be 173.7.

Fig. 4 shows a result of this procedure on a real painting: separating the Hg channel of "Patch of Grass," shown in Fig. 4(a). Fig. 4(b) shows the color visible light image of the corresponding part of the surface painting, registered to the Hg channel image. Fig. 4(c) and (d) then show the recovered hidden-painting and surface sources, respectively. Here, we see that the pink flowers in the surface painting have mostly been removed from the hidden-painting source; they remain after our source separation process, but are quite faint. Meanwhile, the woman's face remains approximately as legible in the hidden-painting source as it was to start, but is now quite difficult to make out in the surface source. Hence, we see that, in the hidden-painting source, our method has preserved the hidden painting's features, while diminishing the effect of the surface features, exactly as desired. While not perfect, this result serves as a clear proof-of-concept for our new source separation approach.

In future work, we hope to explore this approach further. In particular, we wish to develop better algorithms for minimizing the objective function in (1), since the gradient descent approach is quite slow, and does not necessarily reach the global optimum. Indeed, even the example results shown here continued to slowly improve with more iterations run, seemingly for as long as we cared to continue. This suggests that future efforts at streamlining our approach's computational efficiency will likely allow even better source separation results to come from it.

4. Identification and inpainting of areas of attenuation

Finally, in this section, we will discuss the third stage of hidden painting restoration that we propose in this paper: identification and inpainting/repair of areas of information loss. As previously noted, X-ray fluorescence imaging data of the hidden painting may contain areas of information loss or attenuation where a particularly thick or X-ray-absorbent surface feature has blocked signal

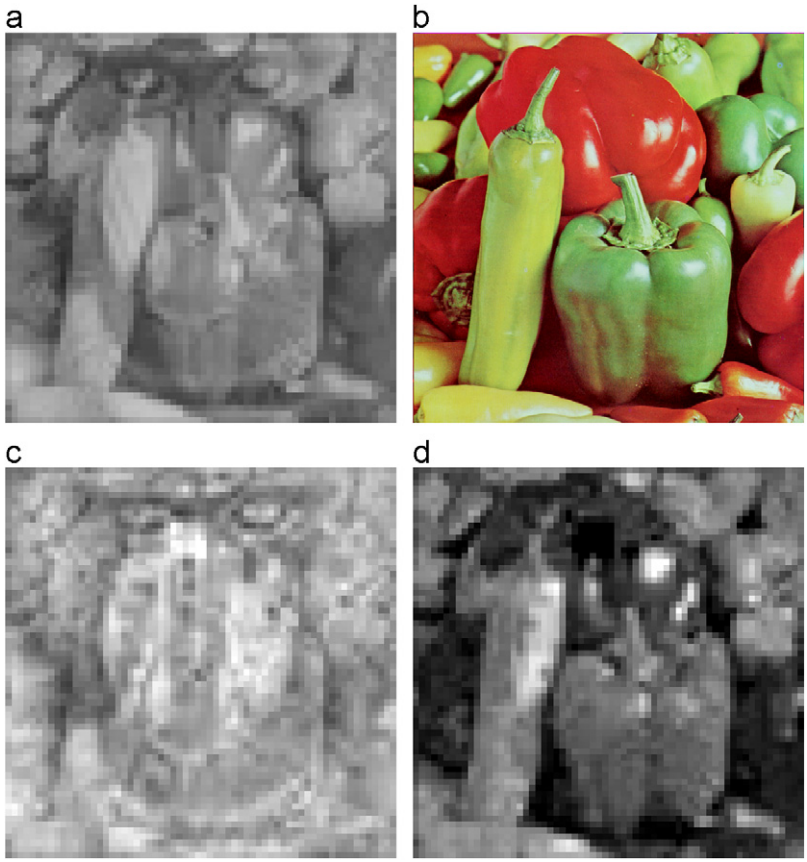

Fig. 3. Results of the procedure for underdetermined source separation described in Section 3 on a synthetic example. (a) The mixture to be separated into sources consists of one half "peppers" image with one half "baboon" image. (b) An RGB color version of "peppers" is used as different-modality side information. (c,d) The two sources after separation. The hidden-painting source shown in (c) is mostly "baboon", while the surface source shown in (d) is mostly "peppers". Note that the legibility of the "baboon" image has greatly increased after applying the method.

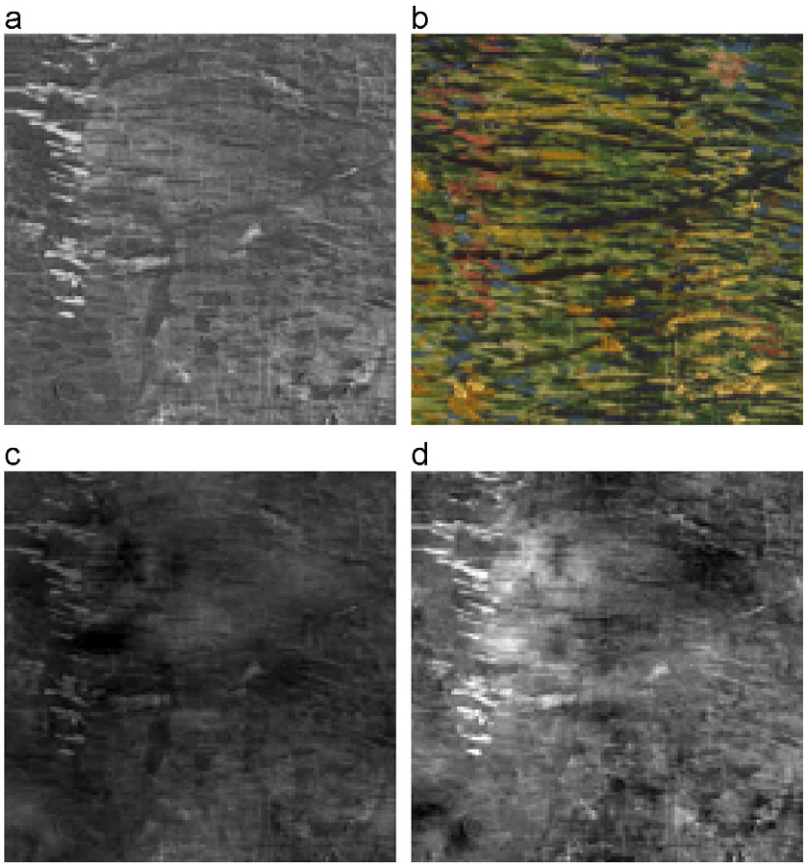

Fig. 4. Results of the procedure for underdetermined source separation described in Section 3 on a real painting example. (a) The Hg channel of "Patch of Grass", to be separated into surface and hidden-painting sources. (b) An RGB color image of the surface painting, to be used as side information in the source separation. (c,d) The (c) hidden-painting and (d) surface sources after separation. Notice that after separation, the pink flowers from the surface are bright in the surface source and greatly diminished in the hidden-painting source. Meanwhile, the face from the hidden painting is clear in the hidden-painting source and hard to make out in the surface source.

from the underlayers from reaching the sensors to be imaged. This is clearly seen in Fig. 1, in which all three channels of X-ray fluorescence imaging show a similar pattern of black spots, and we observe additional darkened streaks across the images, most notably in the antimony channel.

Two main attributes of the surface painting features determine the extent of these losses. First, some pigments are more X-ray absorbent that others and are thus more likely to obstruct signals passing through them. Second, particularly thick layers of paint on the surface are more likely to block signal than thinner features are. Together, these two attributes determine whether signal will be completely lost in an area or merely attenuated.

In this section, we develop two methods to automatically identify areas of probable loss or attenuation, one based on the thickness of surface painting features (Section 4.1) and one based on the composition of these features (Section 4.2). Taking the union of the potentially problematic areas returned by these two methods as the set of areas to potentially correct, we can then distinguish between areas of complete loss and areas of attenuation only (Section 4.4). This allows us to boost signal appropriately in areas of partial attenuation (Section 4.3) and to inpaint in areas of complete loss (Section 4.4).

4.1. Identification of thick surface features and estimation of their height

Our first method to identify areas of potential information loss attempts to do so based on the thickness of the paint layer from the surface painting. The procedure uses raking light photography of the painting, which is commonly used to study the painting's texture. In raking light photography, a light source is placed at a shallow angle at the border of the painting. This results in a photograph (Fig. 5(a)) in which the ridges present in the painting are highlighted by their brightness; additionally, these ridges cause shadows to appear opposite to the location of the light source.

To identify the locations of ridges, we use the saturation channel of the raking light picture, shown in Fig. 5(b). The saturation of a color is a metric of its light intensity and its distribution across the spectrum of different wavelengths/colors. In the figure, high values of saturation

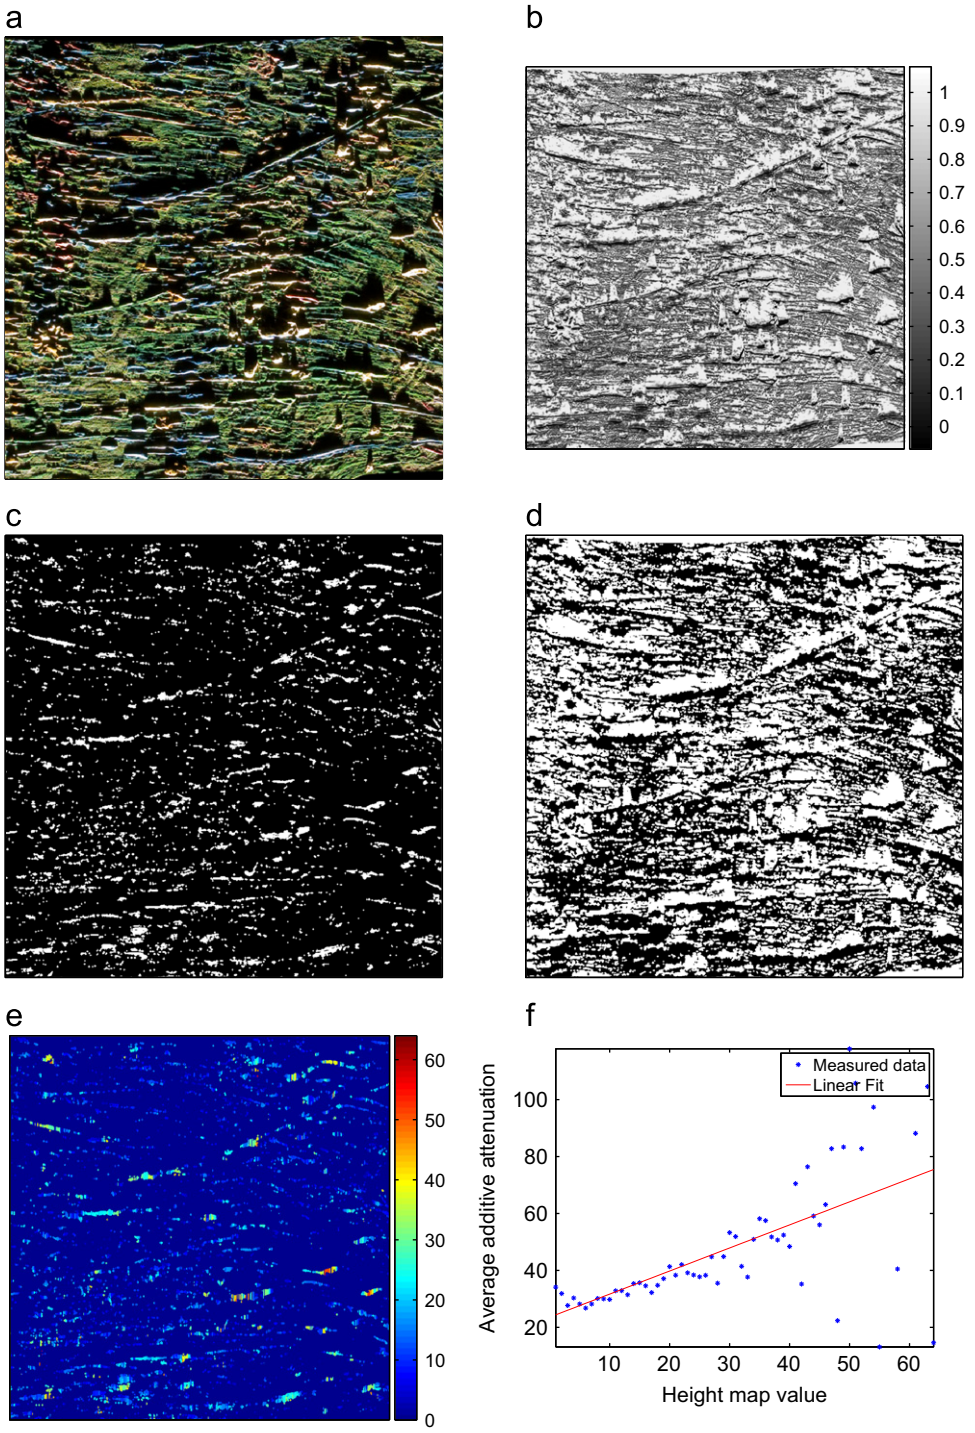

Fig. 5. Use of raking light photograph of "Patch of Grass" to estimate the thickness of paint ridges on the surface of the painting and identify probable areas of attenuation: (a) Raking light photograph demonstrating thickness of various pigments on surface. (b) Saturation channel of raking light photograph. (c) Ridge map from raking light photograph. (d) Shadow map from raking light photograph. (e) Height metric calculated from shadow map. (f) Attenuation function (linear approximation) estimated from height map and antimony image.

correspond to colors that are diffused in the spectrum, i.e. grayscale colors, while low values of saturation correspond to pure, bright colors. Thus, one can identify ridges that reflect the raking light by masking low saturation values and identify dark shadows by masking high saturation values. Fig. 5(c) and (d) shows masked images where

the upper threshold and lower thresholds were set to 0.3 and 0.5, respectively. These ridge and shadow maps have been processed morphologically via opening (erosion and dilation in sequence) to remove spurious detections consisting of single pixels in the image.

After identifying ridges and shadows, the ridge height can be estimated by exploiting the shadow information. We measure the length of the shadow in the direction of the raking light and label the corresponding ridge (which is directly adjacent to the shadow in the same direction) with the calculated shadow length, providing us with a ridge height estimate shown in Fig. 5(e).

Ridges that have a large estimated height represent thick surface features that are likely to be attenuating or obstructing to signal beneath. Hence, we can identify all pixels for which the estimated surface height is above a threshold value as areas of potential information loss or attenuation. We will use the height information that we have estimated to correct attenuation later in Section 4.3.

4.2. Identifying X-ray absorbent surface features

To identify surface features that likely block or attenuate signal based on their composition, we will also want to use an image of the surface painting. For this, it will be easiest to identify chemical composition with RGB surface color, so that we can use the RGB color image of the surface painting to map the distribution of various pigments. However, alternatively, we could choose to use instead the "surface" sources for each chemical element found in Section 3, along with the method that follows.

In any case, we wish to identify surface hues that are highly correlated with darkened areas in multiple chemical channels. Thus, for each RGB color $\mathcal{C}$, we use the pixels of this color in the surface painting as a mask applied to the chemical channel images. We then compare the average grayscale value of the masked-off region with the average grayscale value in the hidden-painting sources in a small surrounding area. Surface colors for which the masked region is significantly darker than its surroundings are likely attenuating. More precisely, for each color $\mathcal{C}$ and chemical element $\mathcal{E}$, we compute the mean of the masked region

$$\tau_{\mathcal{C},\mathcal{E}} = \text{average}(\{I_{\mathcal{E}}(x,y) \,|\, I_{surface}(x,y) = \mathcal{C}\})$$

where $I_{\mathcal{E}}$ is the image for chemical element $\mathcal{E}$ and $I_{surface}$ the surface image, as well as the mean for the surrounding area

$$\rho_{\mathcal{C},\mathcal{E}} = \text{average}(\{I_{\mathcal{E}}(x,y) \,|\, I_{surface}(x,y) \neq \mathcal{C} \quad \text{and} \quad \exists x', y' \text{ s.t. } \sqrt{(x'-x)^2+(y'-y)^2} \leq R, \quad I_{surface}(x',y') = \mathcal{C}\})$$

for a given distance threshold R (we used 50 pixels). We then compare $\tau_{\mathcal{C},\mathcal{E}}$ and $\rho_{\mathcal{C},\mathcal{E}}$ to determine whether the pixels under this surface color tend to be darker than the surrounding pixels. If the ratio ρ/τ exceeds a specified threshold (we used 1.1) and ρ is sufficiently greater than 0, we consider this color $\mathcal{C}$ attenuating.

An experiment on the Van Gogh selected shades of yellow, dark green, and occasionally pink as those likely to have attenuated the hidden painting's signal. These first two colors, yellow and dark green, clearly appear in a pattern mimicking that of the attenuated regions (see Fig. 1(d)). The pink is initially surprising, but on more careful inspection, we see darkened areas in the antimony channel corresponding to the pink flowers in the upper left corner of the surface patch. Using the chosen colors, we then produce a mask for areas of probable loss or attenuation.

4.3. Undoing the effects of partial attenuation

The thickness map obtained in Section 4.1 gives us a means to estimate the extent of and thus potentially correct partial attenuation. We assume a linear relationship between thickness and attenuation in areas where thickness is the key contributor to attenuation. We can then estimate the slope and offset of this linear relationship, thereby creating an attenuation function or map that we can invert for correction purposes.

To estimate the attenuation function, we merge information from the antimony layer image (see Fig. 6(a)) with information from the raking light photograph. First, we perform basic registration of the raking light photograph in Fig. 5(a) with respect to the image in Fig. 6(a). Denote by $I_{\mathcal{E}}(x,y)$ the intensity of the antimony image at pixel (x,y). We estimate the amount of attenuation for each pixel in the ridge map as

$$a(x,y) = \max_{x',y' \text{ s.t. } |x'-x| \leq n, |y'-y| \leq n} I_{\mathcal{E}}(x',y') - I_{\mathcal{E}}(x,y)$$

i.e., the difference between the intensity of the pixel in question and the maximum pixel intensity within a certain n-neighborhood. Once this attenuation has been estimated for every pixel in the ridge map, we average all attenuation values for each different height value to create a map between height estimates and attenuation estimates. Fig. 5(f) shows the estimated map between ridge height and estimated attenuation, together with a best linear fit to the obtained mapping.

The linear functional approximation is then used to correct the attenuation, by scaling and offsetting the attenuated intensity values accordingly, with the results shown in Fig. 6(c). We see that the algorithm has made an attempt to correct values lying in many of the darkened streaks of the image. Areas of complete loss often remain uncorrected however, since there is little signal in these areas to boost.

4.4. Distinguishing and inpainting areas of complete loss

Applying the attenuation correction method above has little impact on areas of complete loss, because in these, the appropriate correction is not proportional to the intensity value. Fortunately, however, the persistence of these areas of complete loss after attenuation correction readily allows us to distinguish them from areas of only partial loss.

To start, we form a new mask for areas of probable loss by taking the union of those masks obtained in Sections 4.1 and 4.2. We process this mask morphologically to remove single pixels and then segment this mask into

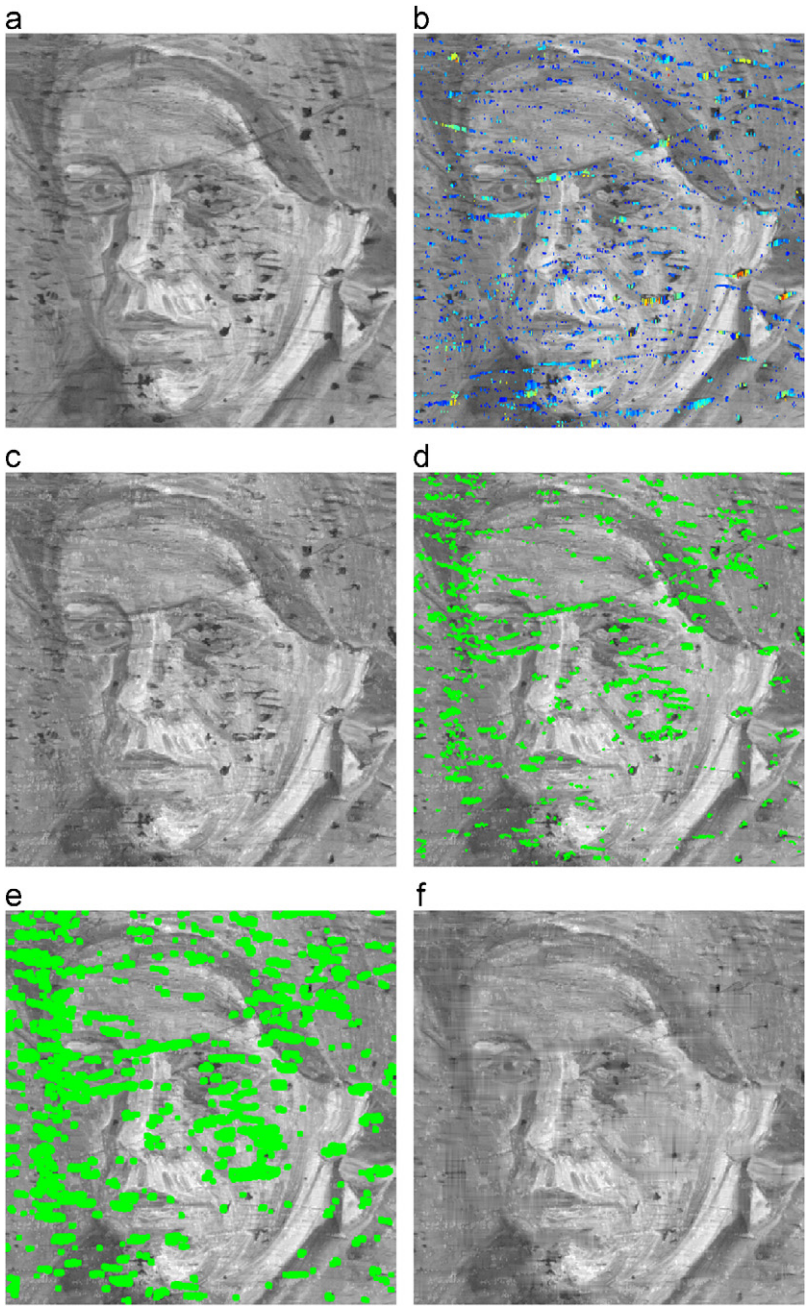

Fig. 6. Results of the two attenuated area identification and inpainting procedures. (a) Original antimony channel image showing attenuated regions. (b) The same antimony image but with the estimated thickness map from Fig. 5(e) superimposed on top of it to show locations of estimated attenuation using the thickness identification method. (c) The antimony channel after application of the thickness-based attenuation correction method. Darkened streaks are reduced, but areas of complete loss remain. (d) Image from (c) with the probable areas of loss identified by the attenuating hue identification method superimposed in green. (e) Same mask as in (d) but dilated to ensure that areas of loss are completely covered. (f) Final result after inpainting using the mask in (e). Most losses have been corrected.

contiguous areas. For each such area, we then measure the average grayscale level underneath it after the attenuation correction of Section 4.3. If this grayscale level after correction remains below a certain threshold, then we identify this as an area of complete information loss that will need to be inpainted.

Finally, we construct a mask of all the contiguous areas we decided to inpaint above, and we convolve this mask with a small point spread function to ensure that it completely covers the edges of all areas of loss. One may then use any one of a number of standard inpainting tools [48–50] to inpaint these areas of complete loss.

We chose the method of Bertalmio and Sapiro [48] since it aims to continue contours as smoothly as possible through the inpainted region, and is thus well-suited for inpainting brushstrokes. As in all inpainting applications, we emphasize that the final inpainted result is just an idea of what the lost regions might look like, based on the content of the surrounding region, and may or may not exactly reflect the true content of these lost areas.

5. Conclusions

In conclusion, we have presented a series of methods for repairing and restoring images of hidden paintings that have been obtained through X-ray fluorescence imaging. We first showed how to correct line shifting errors produced by timing glitches in the image scanning during acquisition via a total variation minimization approach. We found that this approach was able to correct the vast majority of all lines. Second, we gave a method to solve the underdetermined source separation problem with different-modality side information that is presented when we wish to separate each chemical channel image into its component surface and hidden painting features. We found that this method did well at separating the sources in both a synthetic ground truth and a real painting example. Finally, we presented two methods for identifying surface features that were likely to have blocked or attenuated signal from the hidden painting from getting through to be imaged. One method was based on estimating the thickness of the surface features based on raking light photography of the surface painting. The other was based on trying to identify hues on the surface painting that were highly correlated with darkened areas in the hidden painting images. After identifying potentially lost or attenuated areas through these methods, we corrected for estimated attenuation based on the height of the surface features above, and we inpainted those remaining areas of complete signal loss. As a result, we were able to produce chemical channel images of vastly improved quality.

It is our hope that these methods can be of use in eventual imaging of the hundreds or even thousands of hidden paintings currently residing beneath other paintings. As mobile systems like the one of [38] develop further, large-scale imaging projects of this type become more and more possible. We hope that these methods will offer a greatly improved look at these hidden paintings for art scholars and conservators. Moreover, we hope in future work to be able to build upon this work still further, eventually providing full color visual reconstructions of these lost pieces of our cultural heritage for study and for the general public to enjoy.

Acknowledgments

Marco F. Duarte was supported during this research by NSF Supplemental Funding DMS-0439872 to UCLA-IPAM, PI: R. Caflisch. Matthias Alfeld receives support in the form of a Ph.D. fellowship of the Research Foundation Flanders (FWO). This research was also supported by the Interuniversity Attraction Poles Programme—Belgian Science Policy (IUAP VI/16). The text also presents results of GOA "XANES meets ELNES" (Research Fund University of Antwerp, Belgium) and from FWO (Brussels, Belgium) projects no. G.0704.08 and G.01769.09.

References

[1] H. Maitre, F. Schmitt, C. Lahanier, 15 years of image processing and the fine arts, in: IEEE International Conference on Image Processing (ICIP), vol. 1, 2001, pp. 557–561.
[2] M. Barni, A. Pelagotti, A. Piva, Image processing for the analysis and conservation of paintings: opportunities and challenges, IEEE Signal Processing Magazine 22 (2005) 141–144.
[3] V. Capellini, A. Piva, Opportunities and issues of image processing for cultural heritage applications, in: European Signal Processing Conference (EUSIPCO), 2006.
[4] F. Abas, K. Martinez, Craquelure analysis for content-based retrieval, in: International Conference on Digital Signal Processing (DSP), 2002, pp. 111–114.
[5] I. Giakoumis, I. Pitas, Digital restoration of painting cracks, in: IEEE International Symposium on Circuits and Systems (ISCAS), 1998, pp. 269–272.
[6] F. Abas, K. Martinez, Classification of painting cracks for content-based retrieval, in: Proceedings of SPIE 5011: Machine Vision Applications in Industrial Inspection, 2003.
[7] A. Hanbury, P. Kammerer, E. Zolda, Painting crack elimination using viscous morphological reconstruction, in: International Conference on Image Analysis and Processing (ICIAP), 2003, pp. 226–231.
[8] T. Ruzic, B. Cornelis, L. Patisa, A. Ptzurica, A. Dooms, W. Philip, M. Martens, M.D. Mey, I. Daubechies, Virtual restoration of the Ghent altarpiece using crack detection and inpainting, in: Advances Concepts for Intelligent Vision Systems, Lecture Notes in Computer Science, vol. 6915, Springer, Berlin, 2011, pp. 417–428.
[9] S.V. Solanki, A. Mahajan, Cracks inspection and interpolation in digitized artistic picture using image processing approach, International Journal of Recent Trends in Engineering 1 (2009) 97–99.
[10] M. Agarwal, S. Gahlaut, V. Khandelwal, M.C. Srivastava, Exemplar based method to remove cracks in digitized paintings, in: Proceedings of the Second International Conference on Information Processing, 2008, pp. 227–233.
[11] M. Barni, F. Bartolini, V. Cappellini, Image processing for virtual restoration of artworks, IEEE Multimedia 7 (2000).
[12] N. Nikolaidis, I. Pitas, Digital image processing in painting restoration and archiving, in: IEEE International Conference on Image Processing (ICIP), 2001.
[13] S.-C. Pei, Y.-C. Zeng, C.-H. Chang, Virtual restoration of ancient chinese paintings using color contrast enhancement and lacuna texture synthesis, IEEE Transactions on Image Processing 13 (2004) 416–429.
[14] A.D. Rosa, A. Bonacchi, V. Cappellini, Image segmentation and region filling for virtual restoration of art-works, in: IEEE International Conference on Image Processing (ICIP), 2001, pp. 562–565.
[15] R. Berns, S. Byrns, F. Casadio, I. Fiedler, C. Gallagher, F. Imai, A. Newman, L. Taplin, Rejuvenating the color palette of George Seurat's A Sunday on La Grande Jatte-1884: a simulation, Color Research & Application 31 (2006) 278–293.
[16] Y. Zhao, R.S. Berns, L.A. Taplin, J. Coddington, An investigation of multispectral imaging for the mapping of pigments in paintings, in: Proceedings of SPIE Electronic Imaging: Computer Image Analysis in the Study of Art, 2008.
[17] X. Li, D. Lu, Y. Pan, Color restoration and image retrieval for Dunhuang fresco preservation, IEEE Multimedia Magazine 7 (2000) 38–42.
[18] B. Wei, Y. Liu, Y. Pan, Using hybrid knowledge engineering and image processing in color virtual restoration of ancient murals, IEEE Transactions on Knowledge and Data Engineering 15 (2003) 1338–1343.
[19] F. Drago, N. Chiba, Locally adaptive chromatic restoration of digitally acquired paintings, International Journal of Image and Graphics 5 (2005) 617–637.
[20] M. Fornasier, D. Toniolo, Fast, robust, and efficient 2D pattern recognition for re-assembling fragmented images, Pattern Recognition 38 (2005) 2074–2087.
[21] C. Papaodysseus, T. Panagopoulos, M. Exarhos, C. Triantafillou, D. Fragoulis, C. Doumas, Contour-shape based reconstruction of fragmented, 1600 BC wall paintings, IEEE Transactions on Signal Processing 50 (2002) 1277–1288.
[22] M. Fornasier, R. Ramlau, G. Teschke, The application of joint sparsity and total variation minimization algorithms to a real-life

art restoration problem, Advances in Computational Mathematics 31 (2009) 157–184.
[23] F. Stanco, G. Ramponi, A. de Polo, Towards the automated restoration of old photographic prints: a survey, EUROCON 2003: Computer as a Tool 2 (2003) 370–374.
[24] T. Saito, T. Komatsu, T. Ohuchi, T. Seto, Image processing for restoration of heavily-corrupted old film sequences, in: International Conference on Pattern Recognition (ICPR), vol. 3, 2000, pp. 13–16.
[25] S. Tilie, L. Laborelli, I. Bloch, Blotch detection for digital archives restoration based on the fusion of spatial and temporal detectors, in: International Conference on Information Fusion, 2006, pp. 1–8.
[26] J. Blazek, B. Zitova, M. Benes, J. Hradilova, Fresco restoration: Digital image processing approach, in: European Signal Processing Conference (EUSIPCO), 2009.
[27] X. Tang, P.A. Ardis, R. Messing, C.M. Brown, R.C. Nelson, P. Ravines, R. Wiegandt, Digital analysis and restoration of daguerreotypes, in: Proceedings of SPIE Electronic Imaging: Computer Vision and Image Analysis of Art, 2010.
[28] B. Zitova, J. Flusser, F. Lroubek, An application of image processing in the medieval mosaic conservation, Pattern Analysis and Applications 7 (2004) 18–25.
[29] V. Cappellini, M. Barni, M. Corsini, A. Rosa, A. Piva, Artshop: an art-oriented image processing tool for cultural heritage applications, Journal of Visualization and Computer Animation 14 (2003) 149–158.
[30] A.D. Mastio, V. Cappellini, R. Caldelli, A.D. Rosa, A. Piva, Virtual restoration and protection of cultural heritage images, in: International Conference on Digital Signal Processing (DSP), 2007, pp. 471–474.
[31] S. van Heughten, Radiographic images of Vincent van Gogh's paintings in the collection of the Van Gogh museum, Van Gogh Museum Journal (1995) 63–85.
[32] F. Heitz, H. Maitre, C. de Couessin, Application of autoregressive models to fine arts painting analysis, Signal Processing 13 (1987) 1–14.
[33] P. Kammerer, E. Zolda, R. Sablatnig, Computer aided analysis of underdrawings in infrared reflectograms, in: Proceedings of the Fourth International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage, 2003, pp. 19–27.
[34] E. Salerno, A. Tonazzini, L. Bedini, Digital image analysis to enhance underwritten text in the Archimedes palimpsest, International Journal on Document Analysis and Recognition 9 (2007) 79–87.
[35] S. Webber, Technical Imaging of Paintings, Williamstown Art Conservation Center Technical Bulletin, 2008.
[36] A. Pelagotti, A.D. Mastio, A.D. Rosa, A. Piva, Multispectral imaging of paintings, IEEE Signal Processing Magazine 25 (2008) 27–36.
[37] J. Dik, K. Janssens, G. van der Snickt, L. van der Loeff, K. Rickers, M. Cotte, Visualization of a lost painting by Vincent van Gogh using synchrotron radiation based X-ray fluorescence elemental mapping, Analytical Chemistry 80 (2008) 6436–6442.
[38] M. Alfeld, K. Janssens, J. Dik, W. de Nolf, G. van der Snickt, Optimization of mobile scanning macro-XRF systems for the in situ investigation of historical paintings, Journal of Analytical Atomic Spectrometry 26 (2011) 899–909.
[39] A. Anitha, S. Hughes, Attenuating hue identification and color estimation for underpainting reconstruction from X-ray synchrotron imaging data, in: SPIE 7798: Applications of Digital Image Processing XXXIII, 2010.
[40] A. Anitha, A. Brasoveanu, M.F. Duarte, S.M. Hughes, I. Daubechies, J. Dik, K. Janssens, M. Alfeld, Virtual underpainting reconstruction from X-ray fluorescence imaging data, in: European Signal Processing Conference (EUSIPCO), 2011.
[41] A. Brasoveanu, Uncovering a lost painting of Vincent van Gogh, Undergraduate Senior Thesis, Princeton University, 2009.
[42] P. Bofill, M. Zibulevsky, Underdetermined blind source separation using sparse representations, Signal Processing 81 (2001) 2353–2362.
[43] M. Zibulevsky, B. Pearlmutter, Blind source separation by sparse decomposition in a signal dictionary, Neural Computation 13 (2001) 863–882.
[44] Y. Li, S. Amari, A. Cichocki, D. Ho, S. Xie, Underdetermined blind source separation based on sparse representation, IEEE Transactions on Signal Processing 54 (2006) 423–437.
[45] P. Viola, Alignment by Maximization of Mutual Information, Ph.D. Dissertation, Mass. Inst. of Technology, 1995.
[46] F. Maes, A. Collignon, D. Vandermuelen, G. Marchal, P. Suetens, Multimodality image registration by maximization of mutual information, IEEE Transactions on Medical Imaging 16 (1997) 187–198.
[47] P. Thevenaz, M. Unser, Optimization of mutual information for multiresolution image registration, IEEE Transactions on Image Processing 9 (2000) 1083–1100.
[48] M. Bertalmio, G. Sapiro, V. Caselles, C. Ballester, Image inpainting, in: SIGGRAPH, 2000, pp. 417–424.
[49] A. Criminisi, P. Perez, K. Toyama, Object removal by exemplar-based inpainting, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2003.
[50] T. Chan, J. Shen, Nontexture inpainting by curvature-driven diffusions, Journal of Visual Communication and Image Representation 12 (2001) 436–449.