Digital Image Processing of the Ghent Altarpiece: Supporting the painting's study and conservation treatment

Aleksandra Pižurica, Ljiljana Platiša, Tijana Ružić, Bruno Cornelis, Ann Dooms, Maximiliaan Martens, Hélène Dubois, Bart Devolder, Marc De Mey and Ingrid Daubechies

I. INTRODUCTION

Hanging in the Saint Bavo Cathedral in Ghent, Belgium, is The *Ghent Altarpiece*, also known as *The Adoration of the Mystic Lamb* (see Fig. 1). According to an inscription on the outer frames, it was painted by the brothers Hubert and Jan Van Eyck for Joos Vijd and his wife Elisabeth Borluut in 1432. It is one of the most admired and influential paintings in the history of art and has given rise to many intriguing questions that have been puzzling art historians to date [11]. Moreover, the material history of the panels is very complicated. They were hidden, dismantled, moved away, stolen and recovered during riots, fires and wars. The recovery of the panels by the American army in the Nazi hoards deep in the Altaussee salt mines has particularly marked memories. One panel was stolen in 1934 and never recovered. Besides varying conservation conditions, the panels underwent numerous restoration treatments and were even partially overpainted.

One of the most important unresolved questions related to this painting goes back to its creation: the division of hands between the two brothers and their respective workshops. The meticulous study of the painting technique, its different layers and materials, as well as the underdrawings and perhaps even numerous intriguing palm and finger prints could bring us closer to answering that question. However, the *Ghent Altarpiece* hides many other secrets, like the meaning of inscriptions that are difficult to decipher, such as the text in the book depicted in the panel *Virgin Annunciate*.

A major conservation and restoration campaign carried out by the Royal Institute for Cultural Heritage (KIK-IRPA) that is expected to take at least six years started in October 2012. One of the questions of the treatment, supported by an international commission of experts, concerns uncovering Van Eyck's original paint to the extent which can be safely carried out. Indeed, the paintings were covered over centuries with disfiguring retouchings, overpainting and varnishes. Certain decisions regarding the restorations benefit from multidisciplinary research and signal processing could help in this regard.

In this paper, we show progress in certain image processing techniques that can support the physical restoration of the painting, its art-historical analysis or both. We first introduce a multi-modal crack detection algorithm, which gives a clear improvement over earlier reported crack detection results on the *Ghent Altarpiece*. We then show how a relatively simple analysis of the crack patterns could indicate possible areas of overpaint, which may be of great value for the physical restoration campaign, after further validation. Next, we explore how digital image inpainting can serve as a simulation for the restoration of *losses* (missing areas in one or more layers of the painting, often caused by abrasion or mechanical fracture and revealed after the cleaning process). As a separate problem, we address crack inpainting, by outlining the main challenges and proposing a solution which improves upon earlier reported results on this painting [7]. Finally, we explore how the statistical analysis of the relatively simple and frequently recurring objects (such as pearls in this masterpiece) may characterize the consistency of the painter's style and thereby aid both art-historical interpretation and physical restoration campaign. We carry out our analysis on a recently released high-resolution data set and on some images taken during the current treatment of the altarpiece.

II. DATA SET "CLOSER TO VAN EYCK"

Until 2012, digitized scans of old photographic negatives, acquired by Alfons Dierick [12] and kept in the archives of Ghent University, were the only available high resolution data set of the *Ghent Altarpiece*. The development process of these negatives was mainly undocumented, which resulted in a data set where the images vary strongly in quality. Earlier reported results of digital image processing on the *Ghent Altarpiece*, such as crack detection, virtual crack inpainting [7], [24] and pearl analysis [23] were all based on images from that old data set.

In this paper, we report the results on extremely high resolution images that are publicly available on the website *Closer to Van Eyck: Rediscovering the Ghent Altarpiece*[1]. This data set is the result of an interdisciplinary research project that ran from April 2010 till June 2011, with the goal to assess the structural condition of the *Ghent Altarpiece* and determine whether a full restoration of Van Eyck's polyptych was necessary. The surfaces of the altarpiece were documented with the following imaging modalities: *digital macrophotography* (with a pixel size of

[1] http://closertovaneyck.kikirpa.be/

Fig. 1: The *Ghent Altarpiece*, open (left) and closed (right). Image copyright: Ghent, Kathedrale Kerkfabriek, Lukasweb.

7.2 μm; full panels, 140 extreme close-ups, and some cleaning tests), *infrared macrophotography* (in the same resolution), *infrared reflectography* and *X-radiography*. New acquisitions will be added to this data set in the scope of the current conservation-restoration campaign.

III. IMAGE PROCESSING IN SERVICE OF PAINTING RESTORATION

We address here two potential applications of image processing to support restoration of paintings: (i) detecting possible areas of overpaint based on the analysis of crack patterns and (ii) virtually inpainting losses.

A. Crack detection in paintings

Being able to accurately detect cracks can be very relevant to painting conservation since cracking is one of the most common forms of deterioration. Fluctuations in humidity, causing the wooden support to shrink or expand, are the main cause for crack formation. Because the way in which cracks develop and spread partly depends upon the choice of materials and methods used by the artist, assessing cracks is useful for judging authenticity [4]. Cracks can also assist conservators by providing clues to the causes of degradation of the paint surface. An in-depth study of the factors contributing to their formation can support preventive measures [1]. Furthermore, analysis of crack patterns provides *non-invasive* means of identifying the structural components of paintings [4].

Visually, cracks can be categorised into bright cracks on a dark background or dark cracks on a bright background. One can further distinguish between different types of cracks such as ageing cracks, premature cracking (generally due to drying defects related to the painting materials or

their application) or cracks formed only in the varnish layer when it becomes brittle through oxidation. The literature discusses mainly dark cracks; they are typically considered as having low luminance and being local (grayscale) intensity minima with elongated structure [14]. Different crack detection techniques include simple thresholding, line detectors and various morphological filters (see [1] for an overview). The method in [7] operates on a single image modality (visible image) and combines by means of a voting scheme three crack detection techniques: oriented elongated filters, a multiscale extension of the morphological top-hat transformation and a detection method based on dictionary learning [13].

B. BCTF method for multimodal data

The newly acquired multimodal data set (see Fig. 2 for an example) allows for new crack detection techniques that are able to make use of the information provided by each modality, yielding thereby a more reliable detection scheme. However, a pixel-perfect registration is required prior to using all modalities together. The panels of the *Ghent Altarpiece* were already roughly registered for adjacent viewing on the *Closer to Van Eyck* website but the spatial alignment of these pre-registered images is not sufficient in the current context as the images can be shifted by a few pixels or even exhibit local inconsistencies due to the different acquisition modalities. The nature of the different modalities and the stringent requirements for crack detection make direct registration a challenging task. However, since the cracks themselves are a more or less consistent component throughout all modalities we used them for the registration process. Crude crack maps are first obtained by filtering the unregistered images with elongated filters and subsequent thresholding (more details of the
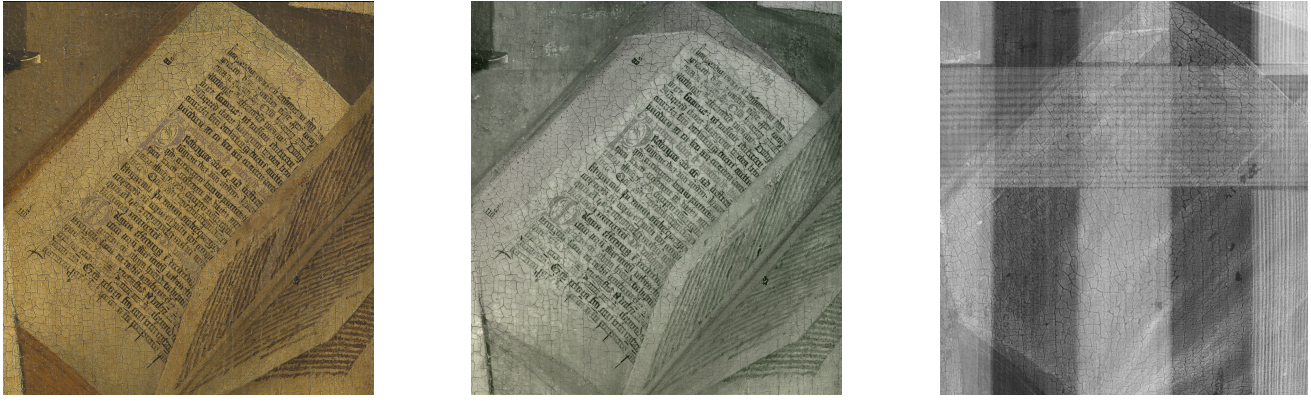
Fig. 2: Acquisitions of the *Ghent Altarpiece*: macro photography, infrared macro photography and X-ray radiography. Source: *Closer to Van Eyck*[1]. Image copyright: Ghent, Kathedrale Kerkfabriek, Lukasweb.

exact procedure can be found in [7]). It should be noted at this point that the presence of false positives is not a nuisance as long as the locations of most of the cracks in each modality are identified. The crude crack maps obtained from the X-radiograph, the visual and infrared images are mutually registered using the algorithm described in [5] (using the infrared crack map as the reference). The resulting transformation is then applied to the original images.

Simply applying the methods described in [7], which were designed for a single image, requires choosing an additional set of parameters per modality, which would be too cumbersome. Here we adopt a semi-supervised Bayesian approach that estimates for each pixel a posterior probability of belonging to the "crack" category, given a large set of feature vectors extracted over all modalities. These feature vectors are obtained by processing each image modality with a number of different filters, commonly used in image processing, ranging from morphological filters to multi-orientation filter banks, as described in [8]. The resulting feature vectors, hereafter denoted as *categorical predictors*, or briefly *predictors*, are quantized into an experimentally chosen number of bins. Let $X_1, ..., X_p$ denote p predictors at a given pixel location, and let Y denote a hidden random variable, taking values $y \in \{0, 1\}$, where the label "1" denotes a crack pixel and "0" a non-crack pixel. The conditional probability $P(Y|X_1, \ldots, X_p)$ is a $d_1 \times \ldots \times d_p$ dimensional tensor, with d_j the number of quantization bins of the j^{th} predictor X_j.

Using the Bayesian conditional tensor factorization (BCTF) of [30], inspired by *higher-order singular value decomposition* [10], the conditional probability tensor can be decomposed as:

$$P(Y = y | X_1 = x_1, \ldots, X_p = x_p)$$
$$= \sum_{j_1=1}^{k_1} \cdots \sum_{j_p=1}^{k_p} \lambda_{j_1,j_2,\ldots,j_p}(y) \prod_{m=1}^{p} \pi_{j_m}^{(m)}(x_m), \quad (1)$$

with the parameters subject to:

$$\sum_{j_m=1}^{k_m} \pi_{j_m}^{(m)}(x_m) = 1, \quad (2)$$

for every combination of (m, x_m). The factorization coefficients $\lambda_{j_1,j_2,\ldots,j_p}(y)$ can be seen as the latent class allocation probabilities and $\pi_{j_m}^{(m)}(x_m)$ as the response class probabilities, which control in a probabilistic manner how the levels of each predictor are clustered. The $k_m \in \{1, \ldots, d_m\}$ value impacts the number of parameters used to characterize the m^{th} predictor. In the special case where $k_m = 1$, (2) yields $\pi_1^{(m)}(x_m) = 1$, which means that $P(y|x_1, \ldots, x_p)$ will not depend on x_m and the m^{th} predictor can be excluded from the model. If $k_m = 1$ for most m's, the categorical predictor model becomes sparse. In practice, we do expect that only a few features have a significant impact on the classification results. More details on the exact posterior computation can be found in [30]. The resulting conditional probability tensor can be used as a lookup table where each entry contains a crack probability for a specific combination of predictor values. If this probability exceeds 0.5, we label the pixel as being part of a crack.

A comparison between the multimodal BCTF method and the crack detection method introduced earlier in [7] on the same part of the painting is depicted in Fig. 3. It can be observed that the older method fails to detect some thin cracks, while it falsely labels some thin dark brushstrokes as cracks. It is clear that the multimodal BCTF method detects more cracks with fewer false positives.

C. Identifying overpaint from crack patterns

Some features of the detected crack patterns may have potential for guiding the restorers to places of interest such as retouchings or heavily damaged areas. As an example we applied the BCTF method, described above, on the upper left corner of the *Joos Vijd* panel. Fig. 4 shows that part of the painting and its corresponding crack map. A rather simple analysis consists of counting the number of crack pixels in a sliding window of 100×100. In doing so we obtain a *crack density map* (see Fig. 4) where we can identify low crack density zones as well as high crack density zones. The lowest (and highest) crack density zones are obtained automatically by hysteresis thresholding where the first chosen threshold is chosen to be very
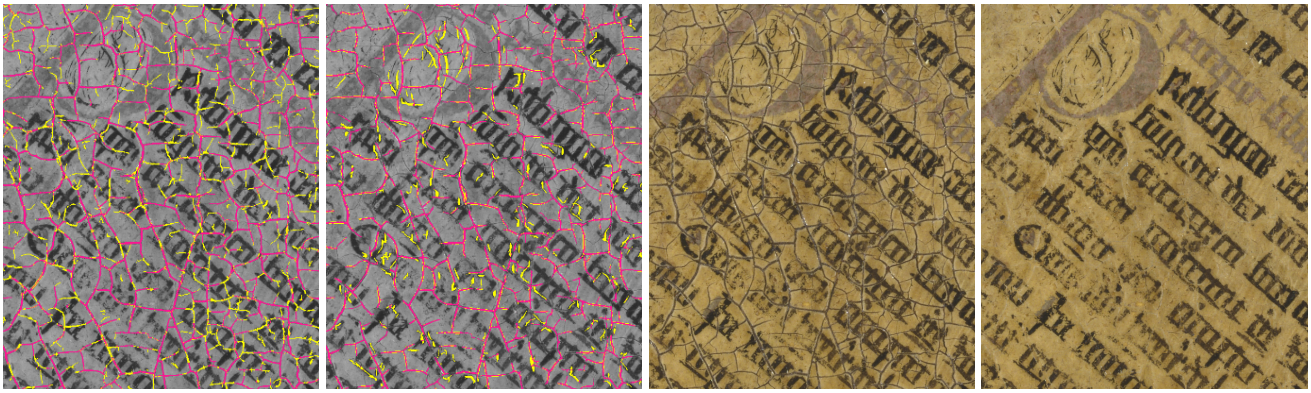
Fig. 3: Crack detection and inpainting. The two images on the left show comparison between BCTF and the method of [7], where cracks detected by *both* methods are marked in *red* and the *differences* in *yellow*. The first image: yellow marks cracks detected *only by BCTF*. The second image: yelow marks cracks detected *only by [7]*. The two images on the right show the original visible data and the result of our patch-based inpainting after BCTF-based crack detection.
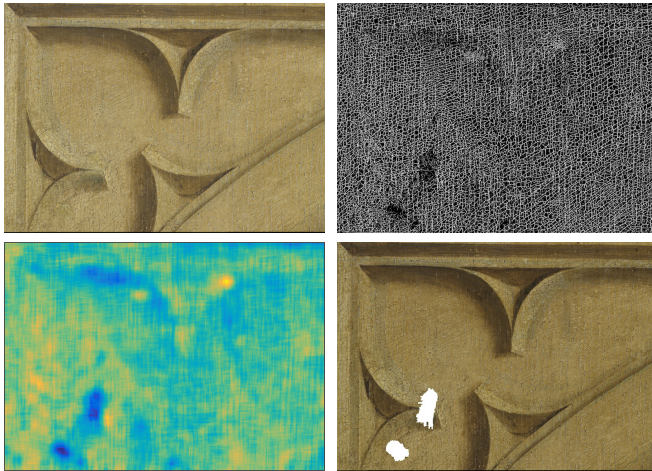

Fig. 4: Crack density analysis within upper left corner of the *Joos Vijd* panel. **Top left**: The high-resolution macro photograph[1]. **Top right**: Detected crack map. **Bottom left**: Crack density map (blue: low density; red: high density). **Bottom right**: Detected zones of lowest crack density.

close to the minimum (respectively maximum) value of the crack density. Painting conservators confirmed that the areas of lowest density marked in white in Fig. 4 are old losses covered with retouching that also overlap on the surrounding original paint. Other zones of low crack density however, such as the ones seen in the upper left corner of the density map, correspond to thinner original paint that developed a different crack pattern over time. The interpretation of the crack maps, like any other diagnostic tool in art conservation, needs to be checked by conservators using other examination techniques and linking the evidence by their critical and material skills.

D. Virtual inpainting

During the ongoing physical restoration of the *Ghent Altarpiece*, deteriorated retouching and overpaint are removed, revealing underlying losses in the original (see Fig. 5). Digital image inpainting can virtually fill in these areas and provide a "simulation" for the impact of certain actions to be taken during the physical restoration process.

A recent overview of inpainting methods is given in [15], and applications to virtual restoration of paintings include [21], [22]. *Patch*-based methods are capable of replicating both structure and texture by filling in the missing region patch-by-patch. In general, for each patch of the missing region (target patch), a well matching replacement patch is found in the available part of the image (source region) and copied to the corresponding location. Preserving structures is achieved by defining the *filling order* [9], which gives priority to the target patches containing object boundaries and fewer missing pixels. The so-called *global* methods, like [17], [25], allow the choice of multiple candidate patches (instead of choosing just one best match in a "greedy" manner) and define inpainting as a global optimization problem.

Fig. 5 shows a part of the *John the Evangelist* panel that has been cleaned in the current restoration campaign of the *Ghent Altarpiece*. Removing overpaint revealed many losses that will be carefully inpainted by the conservators, using stable and reversible materials, in order to restore the visual coherence of the original image. In contrast to losses and abrasions, age cracks are not inpainted in actual conservation unless they severely interfere with the painted form. For this reason, the cracks are not inpainted intentionally in this experiment (unless they are inside a larger loss and assigned therefore to a missing region). We provide here a virtual inpainting simulation of some parts obtained with the patch-based algorithm of [25] (see Fig. 5 and enlarged part in Fig. 6). We have chosen to inpaint several figurative parts that contain structure and texture. The results show how challenging this problem is even for state-of-the-art inpainting methods. None of the methods tested so far produced a satisfying result that fully (albeit virtually) restored the painted form. Experienced conservators master a knowledge of the physical characterics of the paint layers and of the painted forms that call upon
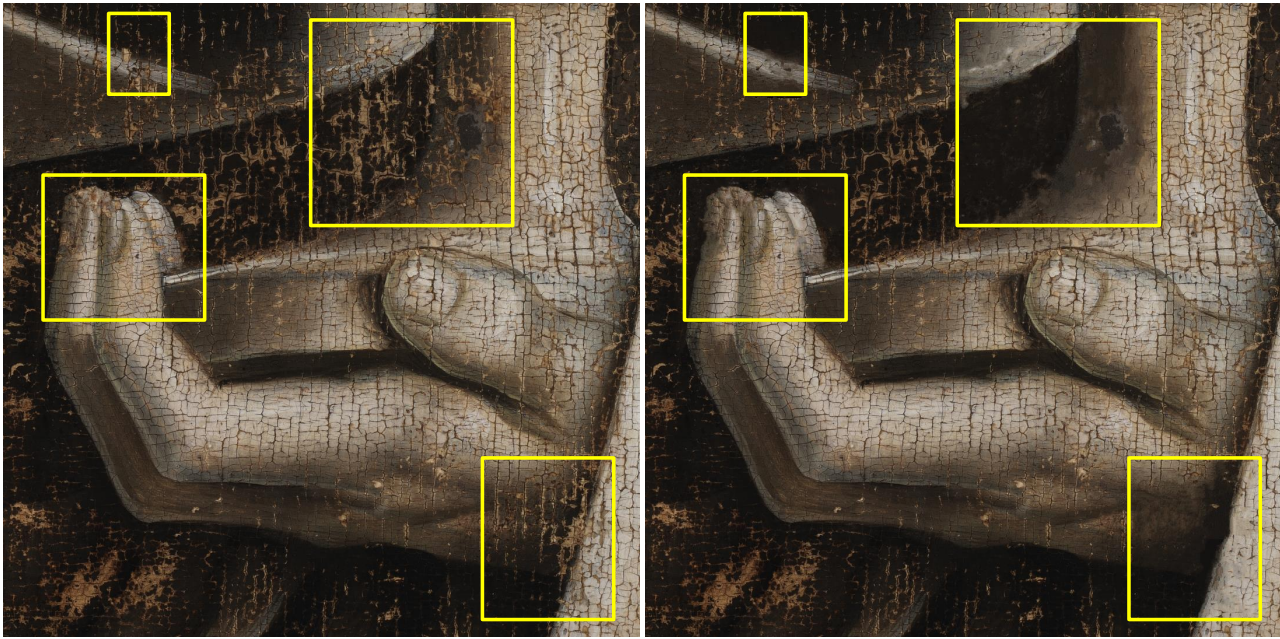
Fig. 5: A part of the cleaned *John the Evangelist* panel. The marked regions are inpainted on the right. (Image copyright Ghent, Kathedrale Kerkfabriek; photo courtesy of KIK-IRPA, Brussels.)
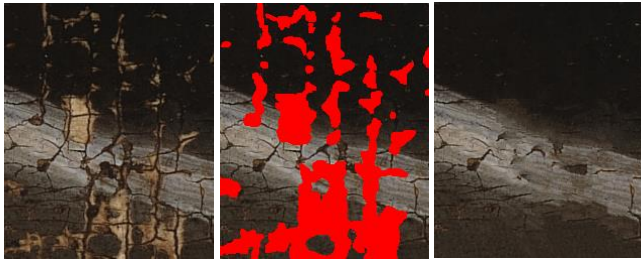

Fig. 6: Results on a small part of the cleaned *John the Evangelist* panel. Left to right: original image, overlay with damaged regions marked in red and the result of the algorithm from [25]. (Image copyright Ghent, Kathedrale Kerkfabriek; photo courtesy of KIK-IRPA , Brussels.)

complex visual perception and interpretation skills. Virtual inpaintings do not provide alternatives to their work on the original but test the potential of the methods that need to be further developed.

IV. DOES CRACK FILLING HELP READING THAT BOOK?

While cracks are rarely inpainted in the actual, physical restoration, virtual crack removal can be of interest in certain aspects. For example, crack inpainting may improve the legibility of the text present in parts of the polyptych, which can be of great importance to art-historical and iconographical studies. Virtual inpainting of the book in the *Virgin Annunciate* panel was reported in [7] on the scans from the Dierick collection. Here we identify some limitations of state-of-the-art inpainting techniques for this type of problem, we introduce an improved method and report the results on the new high-resolution scans from the *Closer to Van Eyck* data set.

A. Challenges in the book of Virgin Annunciate

Since cracks typically appear in images as very thin and elongated regions, crack inpainting methods are often based on rather simple, pixel-wise operations, including median filtering [14], [26], interpolation [2] and controlled anisotropic diffusion [14]. In cases where high-resolution scans are available, such that the width of some cracks spans multiple pixels, patch-based inpainting methods [9], [17], [24] typically yield better results [7], [27].

The book in the panel *Virgin Annunciate* (Fig. 2) is a very challenging case for virtual inpainting because the width of cracks varies a lot and some cracks are difficult to distinguish from parts of the letters. Moreover, as the cracks are typically surrounded by bright clouds of background matching colour, due to the lifting and abrasion of the surrounding paint and thereby imposed light reflections, the immediate areas around the cracks are unreliable as well. The crack inpainting method from [7] specifically tailored to this application already did a pretty good job in improving the legibility of the text and was shown to outperform some "general purpose" patch-based inpainting methods like [9], [17]. However, some problems are still present, e.g., parts of the letters through which wide cracks are passing are occasionally deleted after virtual inpainting. To alleviate this problem, a better approach to handling continuation of image structures is needed. We discuss a possible solution with encouraging initial results next.

B. Crack inpainting

In cases where painted structures, like the characters in Fig. 3, are relatively small compared to the crack width, it is very difficult for the inpainting algorithm to infer the correct structure locally. Patch-based inpainting methods typically
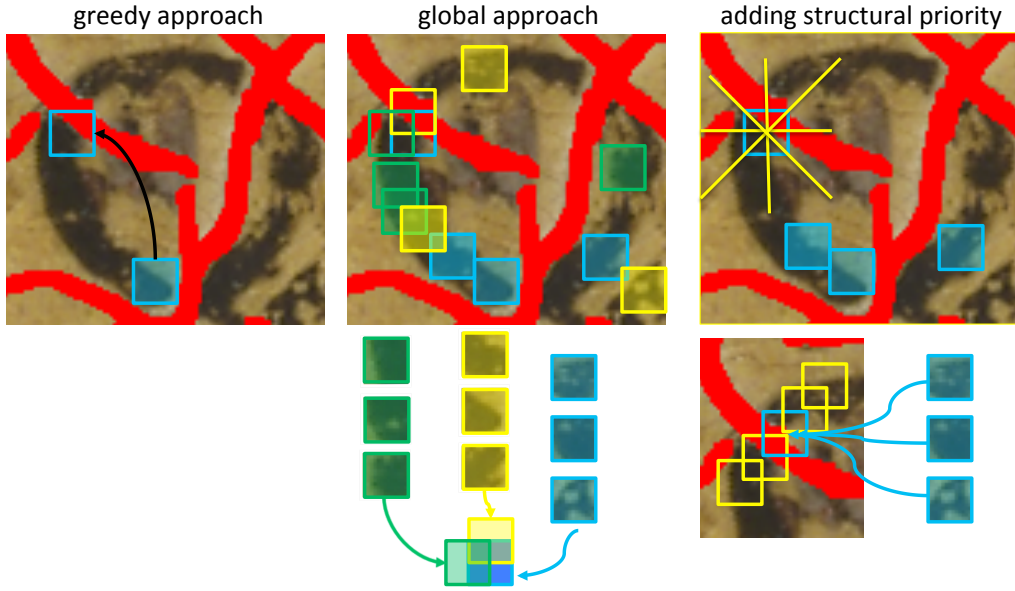
Fig. 7: Patch based inpainting and the proposed improvement. **Left**: A greedy approach chooses one replacement patch based on the known part of the damaged patch. **Middle**: A global approach considers multiple candidates and their mutual agreement. **Right**: Our approach in addition adapts the candidate selection according to the locally detected structures.

Fig. 8: Parts of the original image (first row), the results of our method (second row) and the method of [7] (third row).

handle *structure propagation* by defining the right filling order [9], [15]. Once the filling order is determined, most of the methods choose plausible candidates for replacement patches based *solely* on the agreement with the undamaged part of a single target patch (see Fig. 7, left), and concentrate on defining effective distance metrics between the known portion of the target and the candidates [20]. Matching only against a small part within the target patch increases the risk of propagating wrong textures and wrong colors into the missing region. Global methods, like [17], [25] allow multiple candidates and optimize their mutual agreement in the overlap regions (Fig. 7, middle), but even this cannot ensure agreement with surrounding *undamaged* structures: the optimization which takes care that neighbouring replacement patches mutually agree cannot "undo" the damage done by selecting wrong candidates in the first place. Ideally, undamaged areas around the target patch should be taken into account in the candidate selection as well, ensuring that plausible candidates agree with true structures (see Fig. 7, right).

Guiding the selection of candidate patches by the agree-ment with undamaged areas is related to the idea of *global visual coherence* introduced in [29] for video completion. The approach of [29] is very effective for replicating larger missing areas and textures, but it does not treat continuation of curvilinear structures. Alternative solutions that propagate structures along user specified lines [19], [28] showed excellent results in photo editing, but for our application the amount of the user intervention required by such methods would be prohibitive.

The main idea of our approach is to simultaneously detect directions of local structure propagation and adapt the candidate selection accordingly. We propose a fully automatic and low-complexity method for selecting the candidate replacement patches based on their agreement with the undamaged part of the target patch *and* with the neighbouring undamaged areas, along directions where the structures are likely to propagate. Let ϕ_i denote an image patch centered at position i, $\mathcal{S}(\phi_i, \phi_j)$ a certain measure of *similarity* between ϕ_i and ϕ_j, and denote by $N_{i,k}$ a directional neighbourhood of ϕ_i along direction k. We define *prior preference* $P_{i,j}$ for selecting a source patch ϕ_i

as a candidate replacement for the damaged target patch ϕ_j as follows:

$$P_{i,j} = \mathcal{S}(\phi_i, \phi_j) + \max_k \sum_{l \in N_{j,k}} \mathcal{S}(\phi_i, \phi_l) \qquad (3)$$

The first term measures as usual the similarity with the known part of the target patch. The novelty is in the second term, which takes care of the agreement with the wider context around the target patch. In particular, the more the candidate patch ϕ_i fits with the neighbourhood of the damaged patch ϕ_j in any direction where structures of interest are likely to propagate, the more preference it will get in the selection process. Common measures of patch similarity are defined in terms of the *sum of squared differences* (SSD) among the patches $\mathcal{D}(\phi_i, \phi_j) = ||\phi_i - \phi_j||^2$, calculated over the *known* pixels. We used $\mathcal{S}(\phi_i, \phi_j) = -\mathcal{D}(\phi_i, \phi_j)$ for the candidate selection in (3). With this improved candidate selection process in combination with simple greedy inpainting (selection of one replacement patch at each position), we already obtain a clear improvement over the earlier method from [7], as is visible in Fig. 8 (notice in particular that the effect of deleting parts of letters is less severe). One can also select multiple candidates, with several largest values of $P_{i,j}$ in (3) and subsequently solve the resulting "puzzle" using a global optimization method like in [17], [25]. It would be interesting to explore also alternative solutions, like the statistics of patch offsets [16] or hierarchical, super-resolution based inpainting [18].

V. WHAT CAN PEARLS TELL US?

Painted pearls, which are abundant in the *Ghent altarpiece*, provide a nice case study for the statistical analysis of the consistency of the painterly execution. Spatial histograms, or *spatiograms* [3] were employed in [23] as digital signatures of painted pearls and showed potential to distinguish pearls painted by different artists. In particular, the pearls in the copy of the panel *Just Judges*, made by J. Van der Veken between 1939 and 1951 to replace the panel stolen in 1934, showed clearly different spatiograms than those from other panels of the altarpiece. Similar conclusions were drawn when comparing the spatiograms of the recent reconstructions by other artists. However, this earlier analysis in [23] was performed on the old scans of the altarpiece, with varying resolutions, which may have affected to some extent the numerical findings. It is important to verify the main conclusions of this earlier analysis in the light of the new high quality photographic material. Here we also go a step further, extending the study to different panels of the altarpiece and making a hypothesis that the consistency of the painted pearls could provide an additional support for the division of hands between the painters or within the workshop, as well as for detecting possible areas of former restorations and overpaintings campaigns.
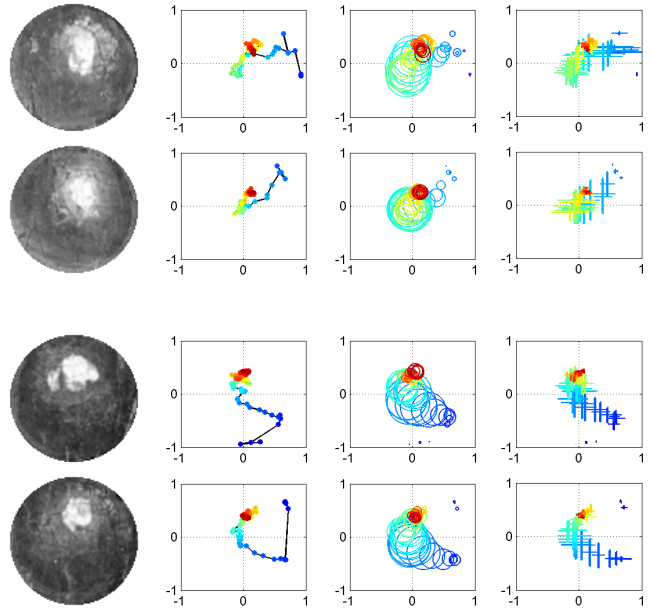

Fig. 9: Two pairs of pearls from the *Ghent Altarpiece* and their corresponding spatiogram triple-plots (S1, S2, S3).

A. Digital pearl signatures from spatiograms

An image spatiogram [3] is a generalized histogram, with second-order spatial moments. Suppose an image consists of N pixels and denote the spatial position of the n-th pixel $n \in \{1, ..., N\}$ by $\mathbf{p}_n = (x_n\ y_n)^{\mathrm{T}}$ and its intensity by I_n. Let b denote a histogram bin, being a range of pixel intensities and let $\mathbf{1}_b(x)$ denote the indicator function (returning 1 if $x \in b$ and zero otherwise). The spatiogram triplet for bin b is then computed as follows:

$$c_b = \eta \sum_{n=1}^{N} \mathbf{1}_b(I_n)$$

$$\boldsymbol{\mu}_b = \eta c_b^{-1} \sum_{n=1}^{N} \mathbf{p}_n \mathbf{1}_b(I_n) \qquad (4)$$

$$\boldsymbol{\Sigma}_b = \eta c_b^{-1} \sum_{n=1}^{N} (\mathbf{p}_n - \boldsymbol{\mu}_b)(\mathbf{p}_n - \boldsymbol{\mu}_b)^T \mathbf{1}_b(I_n)$$

The normalizing constant η is chosen such that $\sum_{b=1}^{B} c_b = 1$. For bins with $c_b = 0$ also the values of $\boldsymbol{\mu}_b$ and $\boldsymbol{\Sigma}_b$ are set to zero (not of interest). To enable comparison between images of different sizes, we normalize all spatial coordinates to the range [-1,1]. For the purpose of visualization of the highly dimensional spatiogram data, [23] proposed spatiogram triple-plots (S1, S2, S3) illustrated in Fig. 9:

S1: connected centers of bins, $\boldsymbol{\mu}_b = (\bar{x}_b\ \bar{y}_b)$;

S2: $\boldsymbol{\mu}_b$–positioned counts of bins (the radii of the circles are proportional to bin counts);

S3: $\boldsymbol{\mu}_b$–positioned variances in the x- and y-direction.

Fig. 9 demonstrates that mutually similar pearls produce similar spatiogram triple-plots.
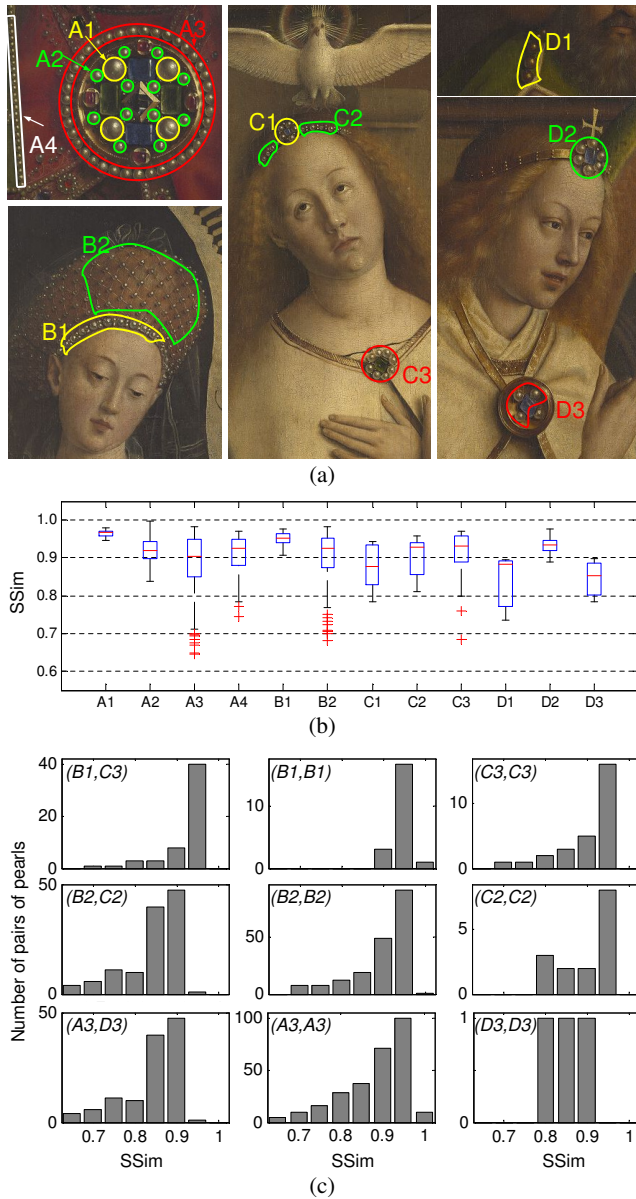
Fig. 10: Exploring consistency of the painted pearls in different panels. a) The selected sets of pearls: A - clothing decoration of *God the Father*, B - the hat of *Cumaean Sibyl*, C - diadem and brooch of *Virgin Annunciate*, D1 - decoration in *prophet Zachary* and D2, D3 - diadem and brooch in *Archangel Gabriel*. b) The corresponding $SSim$ values. The central marks show the medians, the boxes indicate 25th and 75th percentiles, the extreme vertical lines (whiskers) are 1.5 times the interquartile range, and + marks denote the outliers c) Some cross-set and within-set $SSim$ histograms.

B. Consistency of the pearls in the altarpiece

Here, we evaluate the consistency of painted pearls in the *Ghent Altarpiece* on the *Closer to Van Eyck* data set. We measure similarity between two painted pearls with the spatiogram similarity index $SSim$ [6], where $0 \leq SSim \leq 1$. In particular, we select twelve sets of

pearls from five different panels, as marked in Fig. 10(a). The size of each set and the average radius of the pearls (in pixels) are given in brackets: A1(4, 576), A2(12, 265), A3(24, 273), A4(12, 144), B1(7, 177), B2(20, 129), C1(5, 138), C2(6, 111), C3(8, 180), D1(3, 276), D2(6, 224), and D3(3, 239). We make a comparative analysis of the similarity of painted pearls within each set (*within-set* similarity) and between different sets (*cross-set* similarity).

Fig. 10(b) shows *within-set* similarity for the twelve pearl sets, computed from 64-bin spatiograms. Clearly, the largest pearls (A1) are the most similar, which agrees with the findings of [23], but now we can also see that the consistency within each set depends also on the panel and the position of the pearls. For example, C2 shows much higher within-set similarity than D1, even though D1 pearls are twice as large. This could be attributed to the position of D1 pearls (very high in the altarpiece, where they are less well visible to viewers). Similarly, $SSim$ values are less consistent for B2 than for B1 of the *Cumaean Sibyl*, in line with the fact that B2 pearls are smaller, as well as less densely and less neatly arranged than B1 pearls. Sets C1-C3 (from the *Virgin Annunciate*) and D2, D3 (from the *Archangel Gabriel*) are on equally important places in the altarpiece: two outer panels in the middle row of the closed view, see Fig. 1. Our analysis suggests that C3 pearls are more consistent than C1 and C2, which agrees with the fact that they are larger, more visible and have a more central place in the panel. Further on, D2 pearls are almost twice as large as C1 (which appear in the same relative position in the other panel) and indeed give more consistent spatiograms. However, there are some interesting exceptions that deserve extra attention. D3 pearls seem less consistent than D2, even though they are larger and at a more visible place (closer to the viewer). It is still unclear whether this could be (partly) attributed to less careful execution, the state of varnish, possible retouching or overpaint, or simply to having a rather small set (three pearls).

It is also interesting to assess consistency of the painted pearls of similar sizes in different panels. Fig. 10(c) shows *cross-set* $SSim$ for three different combinations of pearl sets in comparison to within-set $SSim$ for the same sets. The $SSim$ histograms show a high similarity between B1 and C3 sets, and much less between B2 and C2. This can be attributed to the fact that B2 and C2 pearls are relatively small. Moreover, B2 is a large set of pearls decorating a hat, and executing them more consistently would not change the overall visual impression much. Also cross-set similarity between A3 and D3 is much smaller than the similarity within A3, which is not surprising since A3 pearls are in the central panel *God the Father* in the open view, brimming with jewels and reflections, while the D3 ones are in the closed view, figuring different type of lighting, possibly executed by different hands within the workshop and have quite likely undergone different conservation treatments in the past.

VI. CONCLUSION

Signal processing shows promise in helping in the decision making process that is involved in a painting's conservation and restoration. Our initial results show that analysis of crack patterns could indicate certain areas of overpaint, even though the processed crack maps still need to be interpreted by conservators using other examination techniques. State-of-the-art inpainting techniques still do not succeed in fully restoring the painted form at a level that would match the criteria of art conservators, but the interaction between the two communities provides already a crucial feedback for improving virtual inpainting techniques in this challenging application. The use of statistical analysis to assess the consistency of the painting style can be of interest for art-historical interpretation of the content. Our analysis of the consistency of the painted pearls in the *Ghent Altarpiece* points to some instances (specific painted objects) that might be of interest to art-historians and conservators to examine in more detail.

REFERENCES

[1] F. S. Abas. *Analysis of Craquelure Patterns for Content-Based Retrieval*. PhD thesis, University of Southampton, 2004.

[2] M. Barni, F. Bartolini, and V. Cappellini. Image processing for virtual restoration of artworks. *IEEE Multimedia*, 7(2):34–37, 2000.

[3] S. T. Birchfield and S. Rangarajan. Spatiograms versus histograms for region-based tracking. In *Proc. IEEE Conf. on Computer Vision and Pattern Rec. (CVPR)*, pages 1158–1163, 2003.

[4] S. Bucklow. A stylometric analysis of craquelure. *Computers and Humanities*, 31, 1998.

[5] I. A. Carreras, C. O. S. Sorzano, R. Marabini, J. M. Carazo, C. O. de Solorzano, and J. Kybic. Consistent and elastic registration of histological sections using vector-spline regularization. In *Computer Vision Approaches to Medical Image Analysis*, volume 4241 of *Lecture Notes in Computer Science*, pages 85–95. Springer Berlin / Heidelberg, May 2006.

[6] C. O. Conaire, N. E. O'Connor, and A. F. Smeaton. An improved spatiogram similarity measure for robust object localisation. In *Proc. IEEE Int'l Conf. ICASSP*, pages I–1069 –I–1072, 2007.

[7] B. Cornelis, T. Ružić, E. Gezels, A. Dooms, A. Pižurica, L. Platiša, J. Cornelis, M. Martens, M. De Mey, and I. Daubechies. Crack detection and inpainting for virtual restoration of paintings: The case of the Ghent Altarpiece. *Signal Processing*, 93(3):605–619, 2013.

[8] B. Cornelis, Y. Yang, J. Vogelstein, A. Dooms, I. Daubechies, and D. Dunson. Bayesian crack detection in ultra high resolution multimodal images of paintings. In *18th Int'l Conf. on Digital Signal Process. (DSP 2013)*, Santorini, Greece, 2013.

[9] A. Criminisi, P. Perez, and K. Toyama. Region filling and object removal by exemplar-based image inpainting. *IEEE Trans. Image Process.*, 13(9):1200–1212, September 2004.

[10] L. De Lathauwer, B. De Moor, and J. Vandewalle. A multilinear singular value decomposition. *SIAM J. Matrix Anal. Appl.*, 21:1253–1278, 2000.

[11] M. De Mey, M. Martens, and C. Stroo. *Vision and material: Interaction between art and science in Jan van Eyck's time*. KVAB Press, 2012.

[12] A. L. Dierick. *The Ghent Altarpiece: Van Eyck's masterpiece revisited - In colour and life size*. 1996.

[13] M. Elad, M. Aharon, and A. M. Bruckstein. K-SVD: An algorithm for designing of overcomplete dictionaries for sparse representation. *IEEE Trans. Signal Process.*, 54(11):4311–4322, 2006.

[14] I. Giakoumis, N. Nikolaidis, and I. Pitas. Digital image processing techniques for the detection and removal of cracks in digitized paintings. *IEEE Trans. Image Process.*, 15(1):178–188, 2006.

[15] C. Guillemot and O. Le Meur. Image inpainting: Overview and recent advances. *Signal Processing Magazine*, 31(1), January 2014.

[16] K. He and J. Sun. Statistics of patch offsets for image completion. In *Proc. ECCV*, pages 16–29, 2012.

[17] N. Komodakis and G. Tziritas. Image completion using efficient belief propagation via priority scheduling and dynamic pruning. *IEEE Trans. Image Process.*, 16(11):2649–2661, November 2007.

[18] O. Le Meur, M. Ebdelli, and C. Guillemot. Hierarchical super-resolution-based inpainting. *IEEE Trans. Image. Process.*, 22(10):3779–3790, 2013.

[19] J. Lee, D.-K. Lee, and R.-H. Park. Robust exemplar-based inpainting algorithm using region segmentation. *IEEE Trans. Consum. Electron.*, 58(2):553–561, 2012.

[20] R. Martínez-Noriega, R. Roumy, and G. Blanchard. Exemplar-based image inpainting: Fast priority and coherent nearest neighbor search. In *IEEE Int'l Workshop on Machine Learning for Signal Processing*, 2012.

[21] G. Papandreou, P. Maragos, and A. Kokaram. Image inpainting with a wavelet domain hidden Markov tree model. In *Proc. IEEE Int'l Conf. on Acoust., Speech, and Signal Process. (ICASSP)*, pages 773–776, 2008.

[22] S.-C. Pei, Y.-C. Zeng, and C.-H. Chang. Virtual restoration of ancient chinese paintings using color contrast enhancement and lacuna texture synthesis. *IEEE Trans. Image Process.*, 13(3):416–429, March 2004.

[23] L. Platiša, B. Cornelis, T. Ružić, A. Pižurica, A. Dooms, M. Martens, M. De Mey, and I. Daubechies. Spatiogram features to characterize pearls in paintings. In *Proceedings of IEEE Int. Conf. on Image Processing (ICIP)*, pages 801–804, 2011.

[24] T. Ružić, B. Cornelis, L. Platiša, A. Pižurica, A. Dooms, W. Philips, M. Martens, M. De Mey, and I. Daubechies. Virtual restoration of the Ghent Altarpiece using crack detection and inpainting. In *Proc. Advanced Concepts for Intelligent Vision Systems (ACIVS)*, pages 417–428, 2011.

[25] T. Ružić and A. Pižurica. Context-aware patch-based image inpainting using Markov random field modelling. *IEEE Trans. Image Process.*, 24(1):444–456, 2015.

[26] S. V. Solanki and A. R. Mahajan. Cracks inspection and interpolation in digitized artistic picture using image processing approach. *International Journal of Recent Trends in Engineering (IJRTE)*, 1(2):97–99, May 2009.

[27] G. S. Spagnolo and F. Somma. Virtual restoration of cracks in digitized image of paintings. *Journal of Physics: Conference Series*, 249(1):012059, 2010.

[28] J. Sun, L. Yuan, J. Jia, and H-Y. Shum. Image completion with structure propagation. *ACM Trans. Graphics*, 24(3):861–868, 2005.

[29] Y. Wexler, E. Shechtman, and M. Irani. Space-time completion of video. *IEEE Trans. Pattern Anal. Mach. Intell.*, 29(3):463–476, 2007.

[30] Y. Yang and D. B. Dunson. Bayesian Conditional Tensor Factorizations for High-Dimensional Classification. *arXiv preprint arXiv:1301.4950*, January 2013.