Virtual Restoration of the Ghent Altarpiece Using Crack Detection and Inpainting

T. Ružić[1], B. Cornelis[2], L. Platiša[1], A. Pižurica[1], A. Dooms[2], W. Philips[1], M. Martens[3], M. De Mey[4], and I. Daubechies[5]

[1] Ghent University, TELIN-IPI-IBBT, Ghent, Belgium;
[2] Vrije Universiteit Brussel, ETRO-IBBT, Brussels, Belgium;
[3] Ghent University, Faculty of Arts and Philosophy, Dept. of Art, Music and Theatre, Ghent, Belgium;
[4] The Flemish Academic Centre for Science and the Arts (VLAC), Brussels, Belgium;
[5] Duke University, Mathematics Department, Durham, NC, USA;
`tijana.ruzic@telin.ugent.be`

Abstract. In this paper, we present a new method for virtual restoration of digitized paintings, with the special focus on the Ghent Altarpiece (1432), one of Belgium's greatest masterpieces. The goal of the work is to remove cracks from the digitized painting thereby approximating how the painting looked like before ageing for nearly 600 years and aiding art historical and palaeographical analysis. For crack detection, we employ a multiscale morphological approach, which can cope with greatly varying thickness of the cracks as well as with their varying intensities (from dark to the light ones). Due to the content of the painting (with extremely many fine details) and complex type of cracks (including inconsistent whitish clouds around them), the available inpainting methods do not provide satisfactory results on many parts of the painting. We show that patch-based methods outperform pixel-based ones, but leaving still much room for improvements in this application. We propose a new method for candidate patch selection, which can be combined with different patch-based inpainting methods to improve their performance in crack removal. The results demonstrate improved performance, with less artefacts and better preserved fine details.

Keywords: Patch-based inpainting, crack detection, digital restoration, Ghent Altarpiece

1 Introduction

Breaking of the paint layer, called *craquelure* or *cracks*, is one of the most common deteriorations in old paintings. The formation and the extent of cracks is influenced by factors such as ageing, drying of the paint, movement of the support (caused by changes in relative humidity) and physical impacts (such as vibrations during transport, etc.). These cracks form an undesired pattern that is, however, inherent to our appreciation of these paintings as old and valuable.

Yet, for specialists in visual perception for example, it is of interest how our perception of the painting is affected when observing it before the ageing process. Moreover, the crack patterns not only make art historical analysis more difficult, but also in the example of the inscribed text (see Section 3.2), the palaeographical deciphering. Therefore, an important task for the restoration of digitized paintings is the detection and removal of the cracks.

Different crack detection techniques include simple thresholding, the use of multi-oriented Gabor filters and various morphological filters [1]. Removal of cracks can be seen as a special class of a more general inpainting problem. While many general inpainting methods exist, like [2–5], further research is needed to improve their performance in the difficult problem of crack inpainting. Specialized methods for crack detection and inpainting include [6–9]. The method from [6] is semi-automatic: the user manually selects a point on each crack to be restored. In [7, 8] cracks are first detected by thresholding the output of the morphological top-hat transform and then the thin dark brush strokes, which have been misidentified as cracks, are separated through neural network classification with hue and saturation as features. Finally, the cracks are inpainted using order statistics filtering or anisotropic diffusion [7] or patch-based texture synthesis [8].

In this work, we focus on the difficult problem of crack inpainting in the *Ghent Altarpiece* (see Fig. 1). The polyptych consisting of 12 panels, dated by inscription 1432, was painted by Jan and Hubert van Eyck and is considered as one of Belgium's most important masterpieces known all over the world. It is still located in the Saint Bavo Cathedral in Ghent, its original destination. As in most 15^{th} century Flemish paintings on Baltic oak, fluctuations in relative humidity, acting over time on the wooden support, caused age cracks. These cracks are particular in a number of ways. First of all, their width ranges from very narrow and barely visible to larger areas of missing paint. Furthermore, depending on the painting's content, they appear as dark thin lines on a bright background or vice versa, bright thin lines on a darker background. Also, since this masterpiece contains many details and some of the brush strokes are of similar color as the cracks, it is difficult to make a distinction between them in some parts of the image. Finally, the bright borders that are present around some of the cracks cause incorrect and visually disturbing inpainting results. These borders are caused by either the reflection of light on the inclination of the paint caused by the crack on the varnish layer or the exposure of the underlying white preparation layer due to the accidental removal of the surface paint due to wear or after cleaning. In addition, the images we use are acquired under different conditions, i.e. different lighting circumstances, chemicals used to develop the negatives, as well as scanners to scan them. Therefore, the quality of the images varies significantly making a general, automatic crack detection and inpainting a very difficult problem.

To deal with these specific problems of cracks and the difficulties of the dataset, we present a novel method for crack detection and inpainting. We perform crack detection by thresholding the outputs of the multiple morphological
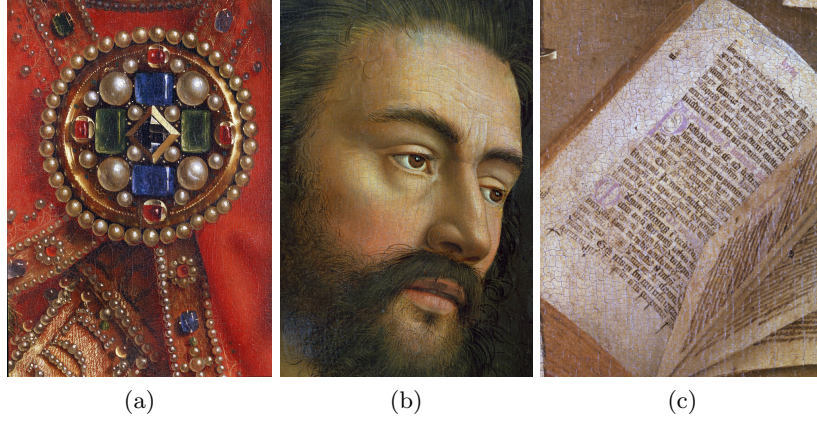
Fig. 1: Test images: (a) the piece of jewellery in the God the Father panel, (b) Adam's face in the Adam panel, (c) the book in the Annunciation to Mary panel.

top-hat transforms, applied by using structuring elements of different sizes. The resulting binary images at different scales are then combined in a novel fashion in order to deal with cracks of different properties and to separate them from other similar structures in the image. Due to the content of the painting (with extremely many fine details) and complex type of cracks (including inconsistent whitish clouds around them), the available inpainting methods do not provide satisfactory results on many parts of the painting. We show that patch-based methods outperform pixel-based ones, while still leaving much room for improvements in this application. We propose a new method for candidate patch selection, that we call *constrained candidate selection*, which can be combined with different patch-based inpainting methods to improve their performance in crack removal. The results demonstrate improved performance, with less artefacts and better preserved fine details. The comparison is for now performed visually, since there is no ground truth. However, we hope that the feedback from art historians on palaeographical deciphering (see Section 3.2) will help to further evaluate the effectiveness of our method.

The paper is organized as follows. Section 2 describes the proposed crack detection algorithm, with multiscale crack map formation and crack map refinement. Section 3 gives a brief overview of state-of-the-art inpainting methods and introduces our novel patch selection method. Finally, we present conclusions in Section 4.

2   Crack Detection

Cracks are typically elongated structures of low intensity, which usually appear as local intensity minima [7]. Recent studies [7–10] indicated the effectiveness of

the morphological *top-hat transformation* [11] for crack detection. As we generally need to detect dark cracks on a lighter background we use the *black top-hat* (or *closing top-hat*) transform $TH_B(A)$, defined as the difference between morphological closing $\varphi_B(A)$ of an input image A by a structuring element B and the image itself. The morphological closing operation is defined as *dilation* followed by *erosion*. The transformation yields a grayscale image with enhanced details, which is further thresholded to create a binary image of details, which are most likely to be cracks.

2.1   Multiscale Crack Map Formation

Since the size of the cracks in our application ranges from very small hairline cracks to larger areas of missing paint, we develop a *multiscale* morphological crack detection approach, which is the main difference with respect to crack detection in [7, 8]. A small structuring element will extract fine scale details while larger structuring elements will gradually extract coarser details. By thresholding these results we obtain different binary images, which we refer to as *crack maps*. The range of sizes of the structuring elements, depends on the resolution at which the painting was acquired. In particular, we used a square structuring element with size ranging from 2×2 to 10×10 pixels. The resulting crack maps are further cleaned up by bridging pixel gaps and removing isolated or small groups of pixels. Fig. 2b and Fig. 2c depict the crack maps obtained by choosing a very small structuring element (a 3×3 square) and a larger one (a 8×8 square). We obtain the final crack map, which we call *dark crack map* (Fig. 2d), by combining all the crack maps from different scales in a fine-to-coarse manner, as explained below.

In Fig. 2a we observe that in areas where dark colours are used, cracks can manifest themselves as thin bright lines. In order to detect these, we use the *white top-hat* (or *opening top-hat*) transformation, which is defined as the difference between the input image and its opening by some structuring element B. Recall that the morphological opening operation is defined as erosion followed by dilation. Just as for the closing top-hat transform the outputs of the white top-hat transform, constructed by using structuring elements of different sizes, are thresholded, cleaned and combined to form a single crack map called *bright crack map* (Fig. 2e). Finally, the dark and bright crack maps are combined as explained below and additional morphological dilation is performed to improve their connectivity. The resulting final crack map, that we call *inpainting map*, is shown in Fig. 2f and will be used as the input for the inpainting algorithm.

One possibility to combine the crack maps from multiple scales is to use logical OR operation. In our experiments, this simple approach works well in practice. However, in some parts of the painting it can happen that the detection misidentifies objects as cracks due to their similar size and structure. To avoid this problem, we construct the multiscale crack map in a fine-to-coarse manner, as follows. We first define a *base map* as a combination of cleaned crack maps at the two finest scales (obtained with the two smallest structuring elements). As we move to coarser scales, we gradually add objects (i.e. groups of connected
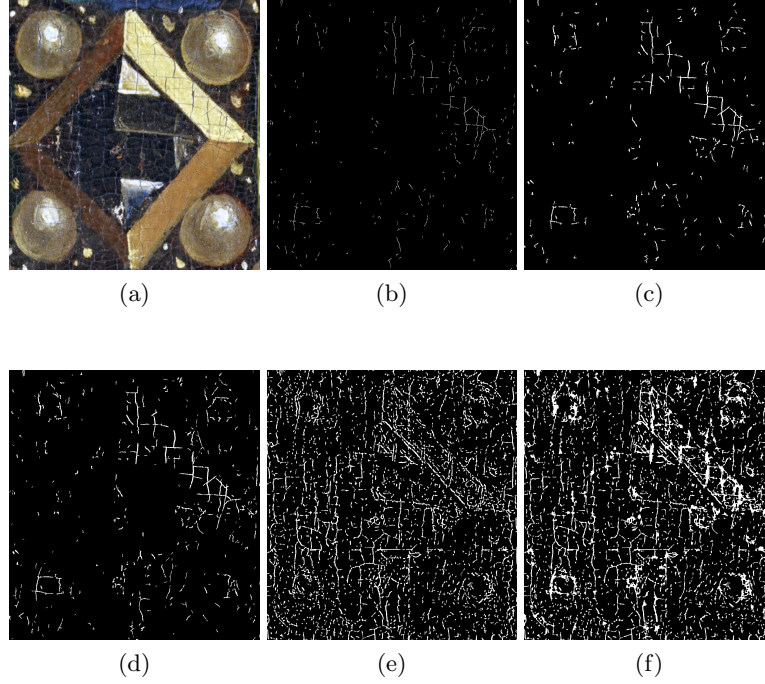
| (a) | (b) | (c) |

| (d) | (e) | (f) |

Fig. 2: Crack maps for the central part of Fig. 1a: (a) original, (b) 3×3 square, (c) 8×8 square, (d) dark crack map, (e) bright crack map, and (f) inpainting map.

pixels) from those scales that are connected to the base map. The advantage of this method is that in our final crack map most misidentified structures (that are detected at larger scales) are eliminated, while it still allows cracks to expand through scale. Since the letters of the book in Fig. 1c have similar properties as the larger cracks (i.e. dark elongated structures on light background) and misdetections can have a negative impact on the further art historical analysis (the palaeographical deciphering), we used this novel approach for the construction of the dark crack map.

We apply anisotropic diffusion [12] as a preprocessing step in order to reduce the noise in the images while still preserving the edges. Furthermore, since the images we work on are acquired under different conditions, as mentioned earlier, the quality of the images varies significantly. Therefore, we choose the colour plane where the contrast between cracks and non-cracks is the highest for crack detection. For example, in the case of the jewel depicted in Fig. 2, we use the green plane of the RGB colour space for the detection of dark cracks and the blue plane for the detection of bright ones.
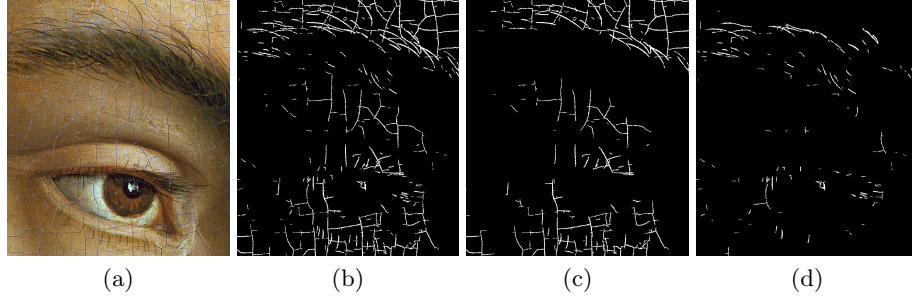
(a)   (b)   (c)   (d)

Fig. 3: Separation of brush strokes misidentifed as dark cracks on part of Fig. 1b: (a) original, (b) dark crack map, (c) dark crack map after separation of misidentified brush strokes, and (d) removed outliers.

2.2   Crack Map Refinement

Some of the digitized paintings contain whitish/bright borders along the cracks (see enlarged parts on Fig. 3a and Fig. 7a). The reason is that the paint is pushed upwards and forms a small inclination. Light reflections from the varnish and the ridges make them appear brighter than their immediate surroundings. Also, due to wear or during previous restorations the surface paint on these elevated ridges may have been accidentally removed, revealing parts of the underlying white preparation layer. While these white borders are problematic for the crack inpainting, they can serve as an additional feature for separating cracks from other structures in the image.

When observing the HSV and RGB colour spaces of most of the images, we notice that bright borders have high values in the blue plane of the RGB colour space. On the other hand, other dark elongated structures (e.g. Adam's eyebrows on Fig. 1b or letters in the book on Fig. 1c) have a very high saturation value in the HSV colour space. These two features are used to further filter the dark crack map obtained with the multiscale closing top-hat transform. For each crack pixel in the dark crack map (Fig. 3b) a weighted average (of the blue and saturation values) of the surrounding pixels is computed. Each object in the binary crack map is given a blue and saturation value. Based on these we are able to detect outliers and remove undesired false positives. The resulting crack map is shown in Fig. 3c and the removed outliers are shown on Fig. 3d.

3   Crack Inpainting

Crack inpainting methods considered in literature so far include order statistics filtering [7, 9], anisotropic diffusion [7] and interpolation [6] and remove cracks one pixel at the time. In our experiments, these pixel-wise methods do not always perform sufficiently well (see Fig. 5b and Fig. 6) and hence, we explore the use of
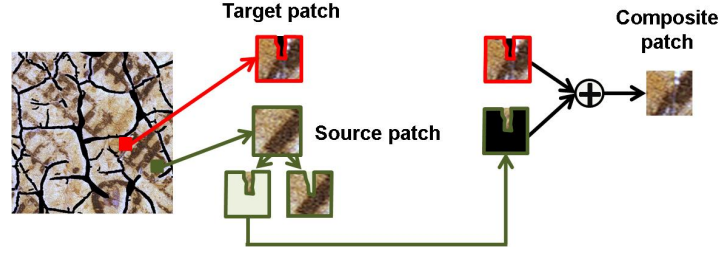
Fig. 4: Schematic representation of a patch-based inpainting algorithm.

patch-based techniques. These methods have recently demonstrated potentials in other inpainting applications, and here we shall adapt and improve them for crack inpainting.

3.1   Patch-Based Inpainting

Present state-of-the-art general inpainting methods are typically patch-based [3–5]. These methods fill in the missing (*target*) region patch-by-patch by searching for similar patches in the known (*source*) region and placing them at corresponding positions. One can divide patch-based methods into two groups: "greedy" and global. The basic idea of the *"greedy"* method [3] is the following: for each patch at the border of the missing region (target patch), find *only* the best matching patch from the source region (source patch) and replace the missing pixels with corresponding pixels from that match, until there are no more missing pixels (Fig. 4). The matching criterion is usually the sum of squared differences between the known pixels in the target patch and the corresponding pixels in the source patch. In this way, both texture and structure are replicated.

Preserving structures, i.e. lines and object contours, is achieved by defining the *filling order*. Priority is given to the target patches that contain less missing pixels and object boundaries. In the case of digitized paintings, the object boundaries are usually difficult to determine due to painting technique (incomplete brush strokes), scanning artefacts, etc. Therefore, we define priority based only on the relative number of existing pixels within the target patch. Results of this method and its superior performance compared to anisotropic diffusion, that was used e.g. in [7], are shown in Fig. 5.

The *global* patch-based method [4] poses inpainting as a global optimization problem. For each target patch in the missing region several candidate patches are found based on the known pixels and/or neighbouring context. Then one of the candidates is chosen for each position so that the whole set of patches (at all positions) minimizes the global optimization function. We applied this method and the result is shown in Fig. 5d. We can see that this complex global method performs similarly to the simpler greedy one for this kind of problem.

In our experiments, the patch-based methods outperform pixel-based anisotropic diffusion (see Fig. 5), but still leave much room for improvement in
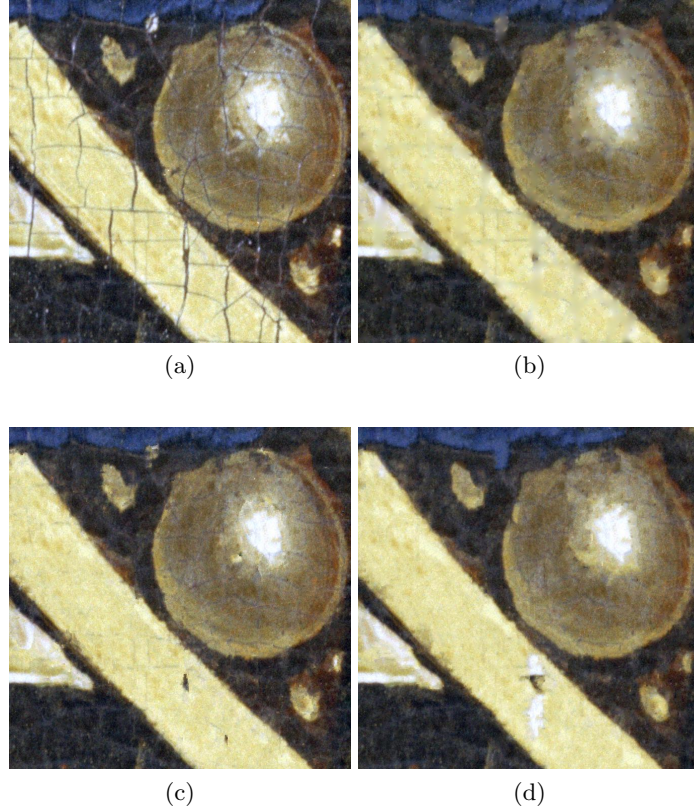
Fig. 5: Inpainting results for the central part of Fig. 1a: (a) original, (b) diffusion result, (c) result of "greedy" patch-based method, and (d) result of global patch-based method.

this difficult application. To improve their performance, some specifics of the problem need to be tackled. We treat some of these in the next subsection. One more obvious problem is the presence of whitish clouds around the cracks. Most inpainting algorithms fill in gaps based on pixel values from their immediate surroundings, in this case the whitish borders around a crack. Hence, the missing regions will likely be filled with incorrect content and the positions of cracks remain visible after inpainting (see the results on the left of Fig. 6). We solve this problem by using the bright crack map described in Section 2 to extend the dark crack map with the corresponding bright regions. Additionally, this crack map also marks cracks that appear completely bright and results in an overall better detection performance. The benefit of using this map, with white borders and bright cracks included, is evident in all cases: in the results on the right of Fig. 6 more cracks are detected and inpainted and the appearance is visually more pleasing.
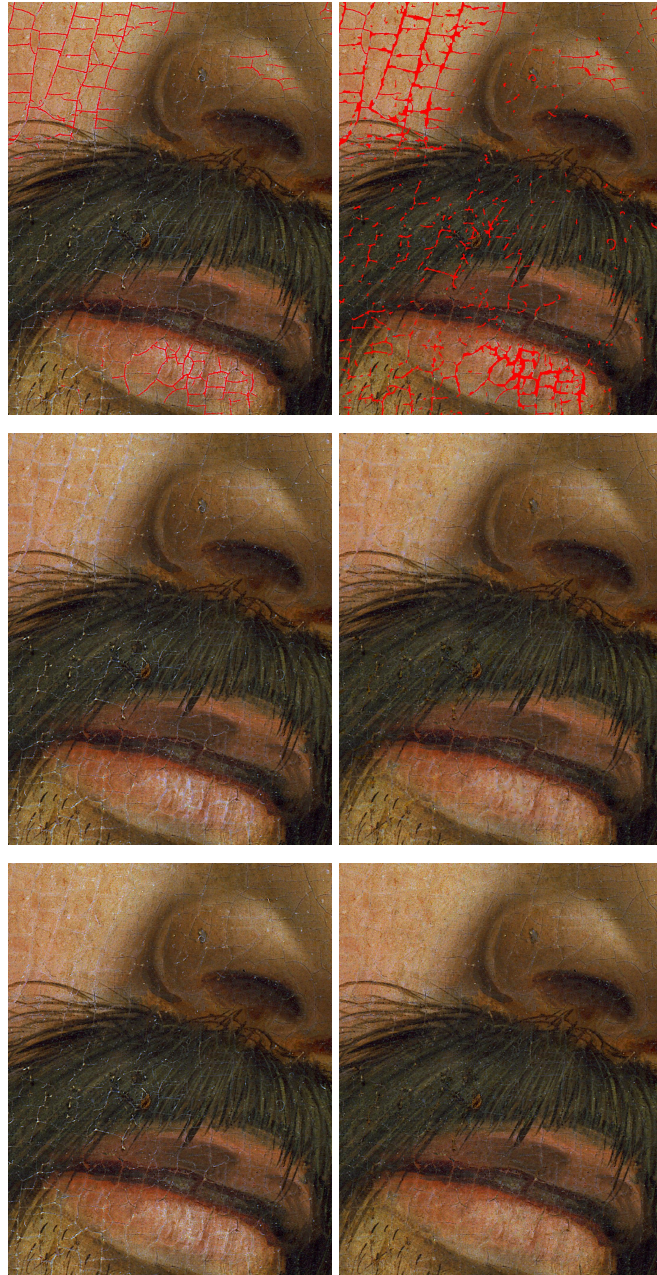
Fig. 6: Influence of white borders on inpainting result. Left: Dark crack map as an input. Right: Combined dark and bright crack map as an input. The first row shows the original image of Adam's nose and mouth overlapped with the crack maps. The second and third row show inpainting results with anisotropic diffusion and patch-based method, respectively.

3.2   Improved Patch-Based Inpainting for Crack Filling

To further improve the crack inpainting results we introduce a novel method that involves two contributions: a novel approach to patch candidate selection and use of adaptive patch size. This method, that we call *constrained candidate selection*, aims at performing context-aware inpainting by constraining the search to certain parts of the image, depending on the position of the current patch. For that we need to segment the image. The simplest case is foreground/background segmentation. The method is also beneficial in the parts of the panels with pronounced whitish borders around the cracks. There are three main steps:

1. **Exclusion of damaged pixels**
   Although we use the inpainting map to deal with the problem of whitish borders around the cracks, some damaged pixels still remain. These pixels are either too distant from the crack, belong to the non-detected cracks or appear in the source region not related to the cracks. We detect these pixels based on their high values in the blue plane and we treat them as missing ones. Additionally, we do not use the patches from the source region containing damaged pixels as possible matches.

2. **Label constrained matching**
   In the results from Fig. 7b it can be seen that patch-based inpainting occasionally introduces some artefacts. Small parts of letters appear erroneously in the background and the other way around, parts of letters get "deleted", i.e. replaced by background. This can happen when the known part of the target patch is not distinctive enough to find the right source patch. To minimize these errors we first segment the image in two classes: foreground and background. We use the k-means segmentation algorithm on the values of the red plane because the difference between the two classes is most visible there. Once we have the segmented image, we constrain the search for candidate patches accordingly. When inpainting a part of the background, i.e. when all the known pixels in the target patch are labelled as background, we only accept source patches that belong completely to the background as candidate patches. Otherwise, we search through all possible candidates.

3. **Adaptive patch size**
   Instead of using a fixed patch size, as most inpainting methods do, we make the patch size adaptive to the local content. We start from the maximal patch size and check if the target patch completely belongs to the background. If this is the case, we constrain the search to the background, as in the previous step. If not, we reduce the patch size and repeat the same procedure (e.g. we reduce the patch size from 17×17 to 9×9). Finally, if even this smaller patch only partially belongs to the background, we search for the match of the target patch of the maximal size at all possible locations.

The proposed constrained candidate selection approach can be applied together with any patch-based method (both "greedy" and global one). However, the global method results in a very high computational load due to the high resolution of the scans, making it impractical for processing of larger areas. On the
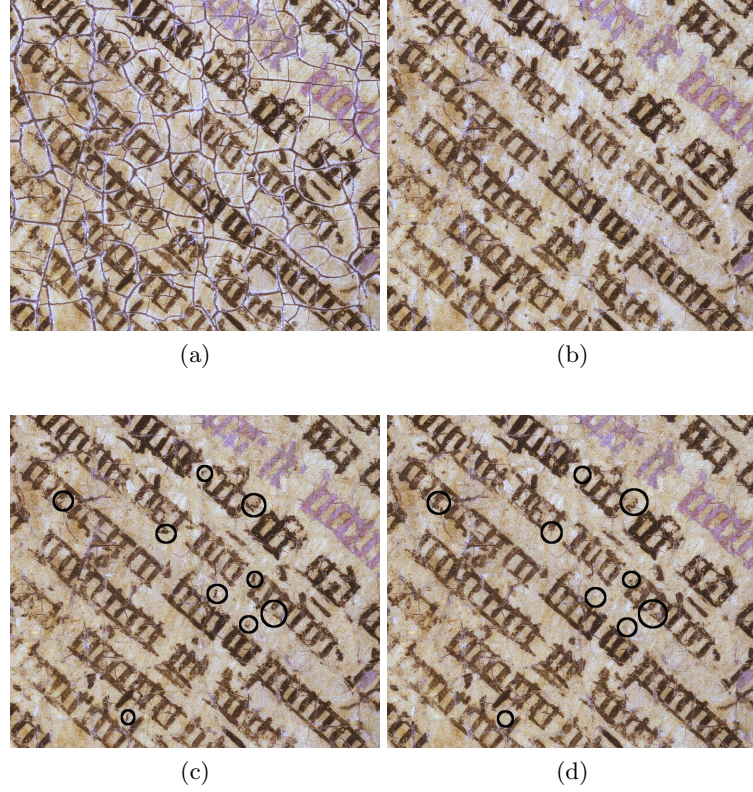
Fig. 7: Inpainting results of the part of the book: (a) original image, (b) result of patch-based method [3] with inpainting crack map, (c) improvement due to proposed patch-based inpainting (with constrained candidate selection and constant patch size), and (d) additional improvement due to adaptive patch size.

other hand, limiting the method on small areas can jeopardize finding the right match. Therefore, we apply our method only together with the "greedy" patch-based method from [3]. Note also that our label constrained matching is applicable on more complex images containing more than two segments. However, in that case, more sophisticated segmentation techniques are required (e.g. [13]).

The effects of the proposed constrained candidate selection are illustrated in Fig. 7c for the constant patch size and in Fig. 7d for the adaptive patch size. Fig. 7d has less artefacts in the background meaning that adaptive patch size approach can better locate the target patches belonging to the background (see parts circled with black). Also, some letters are better inpainted. In comparison with the results of the method from [3] in Fig. 7b, the letters are better inpainted and the whole image contains less visually disturbing white regions. More results can be found on http://telin.ugent.be/~truzic/Artwork/.

4   Conclusion

In this paper, we explored the use of patch-based inpainting of cracks in digitized paintings and highlighted some specific problems using the Ghent Altarpiece as a case study. We introduced a *multiscale morphological filtering* approach for the detection of dark and bright cracks and we presented improvements to the patch-based inpainting methods for crack filling. The results demonstrated improved performance with less artefacts and overall visually more pleasing results.

Acknowledgements

References

1. Abas, F.: Analysis of craquelure patterns for content-based retrieval. Ph.D. thesis, University of Southampton (2004)
2. Bertalmio, M., Sapiro, G.: Image inpainting. In: SIGGRAPH, 417–424 (2000)
3. Criminisi, A., Perez, P., Toyama, K.: Region filling and object removal by exemplar-based image inpainting. IEEE Trans. Image Proc., vol. 13, no. 9, 1200–1212 (2004)
4. Komodakis, N., Tziritas, G.: Image completion using efficient belief propagation via priority scheduling and dynamic pruning. IEEE Trans. Image Proc., vol. 16, no. 11, 2649–2661 (2007)
5. Xu, Z., Sun, J.: Image inpainting by patch propagation using patch sparsity. IEEE Trans. Image Proc., vol.19, no.15 (2010)
6. Barni, M., Bartolini, F., Cappellini, V.: Image processing for virtual restoration of artworks. IEEE MultiMedia, vol. 7, issue 2, 34–37 (2000)
7. Giakoumis, I., Nikolaidis, N., Pitas, I.: Digital image processing techniques for the detection and removal of cracks in digitized paintings. IEEE Trans. Image Proc., vol. 15, no. 1, 178–188 (2006)
8. Spagnolo, G., Somma F.: Virtual restoration of cracks in digitized image of paintings. J. of Physics: Conference Series, vol. 249, no. 1 (2010)
9. Solanki, S. V., Mahajan, A. R.: Cracks inspection and interpolation in digitized artistic picture using image processing approach. Int. Journal of Recent Trends in Eng., vol. 1, no. 2, 97–99 (2009)
10. Abas, F., Martinez, K.: Classification of painting cracks for content-based analysis. In: IS&T/SPIE Elect. Imag. 2003: Mach. Vis. App. in Ind. Inspection XI (2003)
11. Meyer, F.: Iterative image transforms for an automatic screening of cervical smears. J. Histoch. Cytochem. 27, 128–135 (1979)
12. Perona, P., Malik, J.: Scale-space and edge detection using anisotropic diffusion. IEEE Trans. on Pattern Anal. and Machine Intel., vol. 12, no. 7, 629–639 (1990)
13. Malik, J., Belongie, S., Leung, T., Shi, J.: Contour and Texture Analysis for Image Segmentation. Int. J. Comput. Vision, vol. 43, issue 1, 7–27 (2001)