SCIENCE: Virtual Restoration and Mathematical Analysis of the Pearls in the *Ghent Altarpiece*

Aleksandra Pižurica
Ljiljana Platiša
Tijana Ružić
Bruno Cornelis
Ann Dooms
Maximiliaan P. J. Martens
Marc De Mey
Ingrid Daubechies

Abstract

This Chapter addresses virtual restoration and painter style characterization in digitized versions of The Adoration of the Mystic Lamb or The Ghent Altarpiece. Virtual restoration enables removal of the signs of aging, such as cracks. Today it is possible for the viewer to appreciate the finest details of the masterpiece by zooming into its high-resolution scans. Especially in these enlarged details, the appearance of inevitable signs of aging can become prominent. Virtual restoration enhances not only visual experience, which is important from the esthetical and psychological points of view, but in some cases it also facilitates deciphering the content (like text fragments) of the painting. For example, virtual restoration improves the legibility of the text in the Annunciation to Mary panel. Equally challenging is the development of new tools for painter style authentication. The second part of this Chapter reviews some of the metrics related to objective characterization of painted pearls, which are so beautiful and abundant in the works of Van Eyck. The potentials of these metrics to distinguish between painter styles are discussed as well as some side applications, such as bringing a painted object closer to the style of another painter.

Introduction

Digitization of art works has become a common practice. Museums are digitizing their collections mainly for the purposes of archiving and dissemination. This way the cultural heritage is protected and made accessible to a larger number of people. Digitization also enables the virtual restoration of art works and their mathematical analysis. For example, we can remove the signs of aging (such as cracks) from a digitized painting, visualize the effect of using different varnishes, discover patterns that would otherwise remain unnoticed or facilitate detection of forgeries.

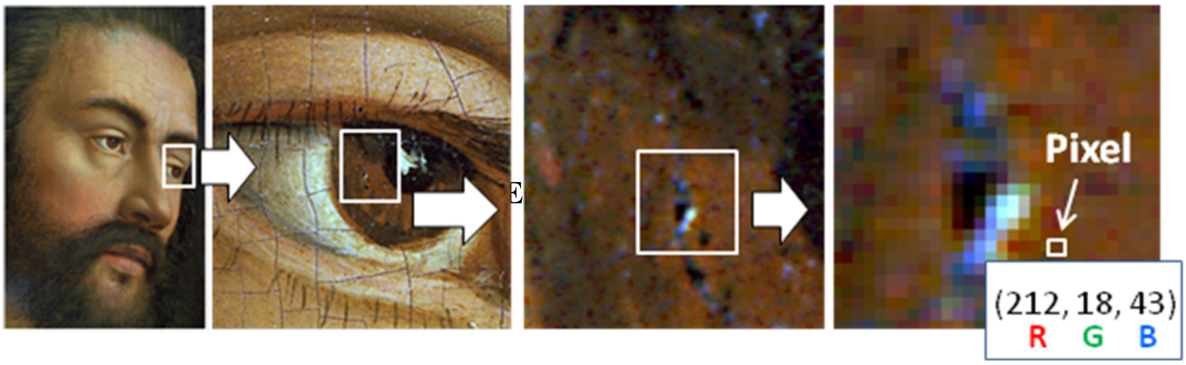

Fig. 1 Zooming into a digitized scan. We see gradually enlarged details of *Adam* up to a level where pixels become clearly visible. The color of each pixel is characterized by the values of its three components: Red (R), Green (G) and Blue (B).

In technical terms, digital images are composed of *pixels*. The name pixel was derived from *picture element*, which was shortened to pictel and evolved finally to pixel. Fig. 1 illustrates zooming into a digitized scan of *Adam* from the *Ghent Altarpiece* polyptych. After sufficiently enlarging the digitized scan, individual pixels become visible (see Fig. 1 on the right). In a color image, each pixel is characterized by the values of its Red (R), Green (G) and Blue (B) components. Typically, images are stored with 8 bits precision per color channel, meaning that there are 2 to the power 8 distinctive values possible per channel, i.e., the intensities in each channel are integer values ranging from 0 to 255. In certain applications, a larger "bit depth" is used, e.g. 16 bits per channel, resulting in a "deep color" representation. The actual process of digitizing paintings, as well as challenges and technologies for accurate color reproduction are described in (Berns, 2001) and (Larue et al, 2007).

With the rapid development of imaging sensors and various imaging modalities, the interest in scientific analysis of paintings is growing. It is now possible to zoom in on the tiny details of the painting or the brushstrokes, revealing structures that could never have been noticed by the naked eye. Moreover, imaging in different parts of the electromagnetic spectrum (from infrared to X-ray) as well as simultaneous imaging in a multitude of narrow spectral bands (hyper-spectral imaging) can reveal other amazing aspects, such as underdrawings (Walmsley et al, 1994; Ibarra-Castanedo et al, 2010) and differentiation between the paint layers that would otherwise remain undiscovered. Digitization of paintings also enables a marvelous interaction between art and vision science: vision scientists can learn from works of art about features that are important for our visual perception of a scene (Livingstone, 2002). Digital painting analysis is a rapidly growing field, attracting an increasing interest in the signal processing community (Abry et al, 2015). Image processing techniques have already demonstrated potential in tasks such as characterization of painting style and forgery detection (Johnson et al, 2008; Platiša, et al. 2012) crack detection (Cornelis et al, 2013) and virtual inpainting (Spagnolo and Somma, 2010, Ružić et al, 2011, Pižurica et al, 2015).

In this Chapter, we demonstrate and discuss application of image processing and machine learning techniques in virtual restoration and stylistic analysis of the *Ghent Altarpiece*. We present some results of our research that was conducted with the aim to support the physical restoration of the painting, its art-historical analysis or both. We first introduce the problem of crack detection and illustrate some of our techniques that make use of multi-modal data. Next, we address a related problem of detecting automatically paint losses (missing areas in one or more layers of the painting, often caused by abrasion or mechanical fracture and revealed after the cleaning process). We then show how digital image inpainting can serve as a simulation for the restoration of losses and also as a tool for virtual crack inpainting. Our results show that virtual crack inpainting can facilitate interpretation of inscriptions like the text on the intriguing book in the *Annunciation to Mary* panel. Finally, we explore how the statistical analysis of the relatively simple and frequently recurring objects (such as pearls in this masterpiece) may serve in the attribution of the painter's style. Our statistical analysis of pearls in the Ghent Altarpiece offers also some insights into the consistency of the painter's style, which may be helpful both for art-historical interpretation and physical restoration of the painting. 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.

Multimodal high-resolution data set

Most of the earlier reported results of digital image processing on the *Ghent Altarpiece*, such as crack detection and inpainting in (Ružić et al, 2011; Cornelis et al, 2013) and pearl analysis in (Platiša, et al. 2011; Platiša, et al. 2012) were based on digitized scans of old photographic negatives, acquired by Alfons Dierick (Dierick, 1996). Until 2012, these old scans that are 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.

In this Chapter, 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].

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

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 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 are being added to this data set in the scope of the current conservation-restoration campaign.

Virtual restoration of the Ghent Altarpiece

The proces of virtual restoration starts with automatic detection of areas that need to be digitally corrected. We focus on digital inpainting of two types of deteriorations (i) cracks and (ii) paint losses. We describe next specifics of these degradations and in the following sections we present the corresponding detection methods and a virtual inpainting (i.e., *filling*) method.

Cracking of the paint layers (or *craquelure*) is one of the most common deteriorations in old paintings, arising inevitably with aging of the materials. The severity of this degradation is affected by many factors, from mechanical stress exposure to climate changes such as variations in temperature and relative humidity or pressurization (e.g., during air transport) (Abas and Martinez, 2003). In most 15th century Flemish paintings on Baltic oak, including the Ghent Altarpiece, age related cracks were caused mainly by fluctuations in relative humidity that make the wooden support shrink or expand. Age related or mechanical cracks can affect the entire composition of the paint layers, including both the preparation and the paint layers on top of it. Different from this is premature cracking (Mohen et al, 2006) being more dull-edged than the age-related cracks and originating in only one of the paint layers. Premature cracks arise from a defective technical execution at the painting stage, such as not leaving enough time for a layer to dry, or applying a layer that dries faster than the underlying one. These cracks should be further distinguished from varnish cracks, formed only in the varnish layer, when it becomes brittle through oxidation (Cornelis et al 2013).

The crack pattern can take different shapes: from nearly rectangular, circular or web-shaped to one-directional, tree-shaped or can even appear completely random (Bucklow, 1997; Abas, 2005). The appearance of the cracks and the whole crack pattern depends upon the choice of materials and methods used by the artist, and can hence be used for judging authenticity of the painting (Bucklow, 1997) and in general for non-invasive identification of its structural components (Bucklow, 1998). Analysis of crack patterns can reveal causes of degradation of the paint surface, which is of particular interest to conservators and which can help preventing/reducing further degradations (Abas, 2005) or even discover possible areas of overpaint and retouching (Pižurica, 2015).

Detection of paint losses is another problem where image processing techniques may be of great help to art conservators and restorers. Loss of paint in one or more layers can arise due to abrasion and mechanical fracture. In old oil paintings, paint losses were often overpainted during various restoration campaigns. Modern conservation treatments typically require not only removal of old varnish, but also removal of old retouches and overpaint, which may reveal paint losses underneath. Detection of such paint loss areas is needed for estimating the extent of the damaged area, which needs to be maintained for documenting purposes, but also as a crucial step for virtual inpainting to provide simulations for the actual restoration. While crack detection has been widely studied in the literature, the problem of automatic paint loss detection has received little attention so far, despite its importance. In the following, we give an overview of the state-of-the-art in relation to both of these problems and illustrate the corresponding results on the *Ghent Altarpiece*.

Crack detection

Specific methods for crack detection in old paintings were reported e.g. in (Barni et al, 2000; Giakoumis et al, 2006; Solanki and Mahajan, 2009; Spagnolo and Somma, 2010; Cornelis al, 2013). Most of these are semi-automatic procedures, where users need to specify a location believed to belong to a crack network. The algorithm will then track other suspected crack points based on two main features, absolute gray-level and crack uniformity. Initial crack patterns are often detected by thresholding the output of a morphological top-hat transform (Meyer, 1979; Serra, 1984). Cracks are subsequently separated from brushstrokes either by feeding some color features to a neural network or by letting a user manually select seed points.

We found crack detection on the *Ghent Altarpiece* challenging especially due to the following: 1) strongly varying length and width of the cracks, ranging from hairline structures to larger areas of missing paint; 2) varying color (Fig. 2 shows examples of dark cracks on a lighter background and bright cracks on a dark background); 3) interference with image details (see Adam's eye in Fig. 1: some cracks are difficult to distinguish from fine details such as eye lashes) and 4) whitish clouds present around some cracks, caused by light reflections on the elevated ridges around the cracks. Also, during previous cleaning, the surface paint on these elevated ridges may have been accidentally removed, revealing parts of the underlying white preparation layer. The effect can be accentuated by the acquisition of the image data

Fig. 2 A detail of the broach from the *God the Father* panel. Cracks range from thin, hairline structures to wide lines and parts of missing paint. Also clearly visible is a strong variation in the crack color: from dark cracks on the golden part to very bright ones in dark regions.

To cope with these problems, a hybrid approach for crack detection from (Cornelis et al 2013) combines different crack detection techniques within a voting scheme. The employed methods are based on multiscale morphological transformations (Serra, 1989), oriented elongated filters (Poli and Valli, 1997) and dictionary learning principles (Elad et al, 2006). It was found that by fusing the results of these different detection methods, a more complete crack pattern can be extracted.

Although the hybrid crack detection method of (Cornelis et al, 2013) has been rather successful, its performance is limited by the fact that it has been designed for a single imaging modality (typically, visual data). Combining information from different imaging modalities, including X-ray and infrared can improve crack detection, especially in areas where cracks are visually difficult to distinguish from the background. The newly acquired multimodal data set, that we use in this study (see Fig. 3 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.

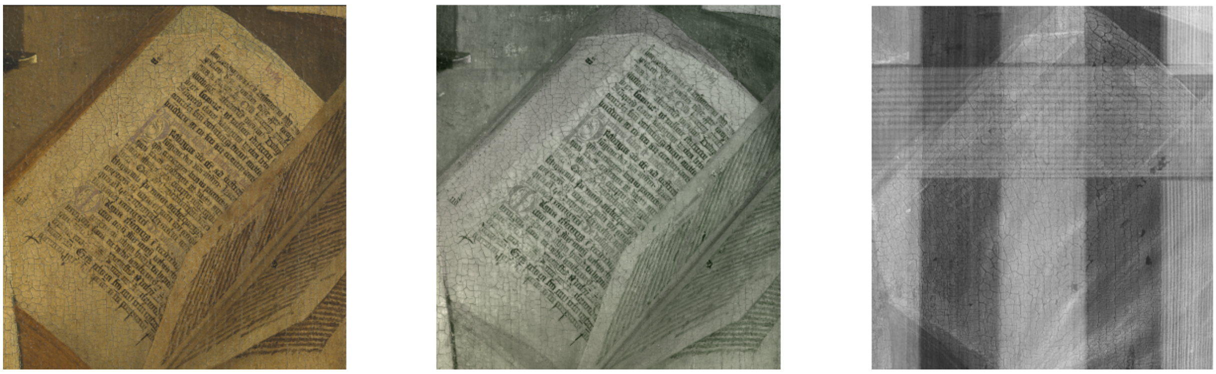

Fig. 3 Acquisitions of the Ghent Altarpiece: macro photography, infrared macro photography and X-ray radiography. Image copyright: Ghent, Kathedrale Kerkfabriek, Lukasweb.

A pixel-perfect registration is required prior to using different modalities together. This means applying the necessary transformations such that the content in different modalities is perfectly overlaid. 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. Since the cracks themselves are a more or less consistent component throughout all modalities, they can be used as reliable landmarks in the registration process. It is hence advisable to detect first crude crack maps from each modality separately, using some single-modal method, and then align the modalities using a registration method such as (Carreras et al, 2006).

A multi-modal crack detection method of (Cornelis et al, 2013a) extracts a large number of features from each modality by applying various filters, and combines these features within a Bayesian framework. In particular, *a posteriori* probability of crack presence is inferred, given the information supplied by the various features from the registered modalities. A comparative analysis between crack detection approaches reported in (Pižurica et al, 2015) shows that the multi-modal approach is able to detect more cracks, while at the same time reducing the false detections. An example in Fig. 4 demonstrates this clearly, on a challenging example, where cracks interfere with the text and painted features. Note that relatively many actual cracks are detected only by the multi-modal approach (see detections marked yellow in the second image), and that much more false alarms are produced by the single-modal approach (see detections marked in yellow in the third image in Fig. 4).

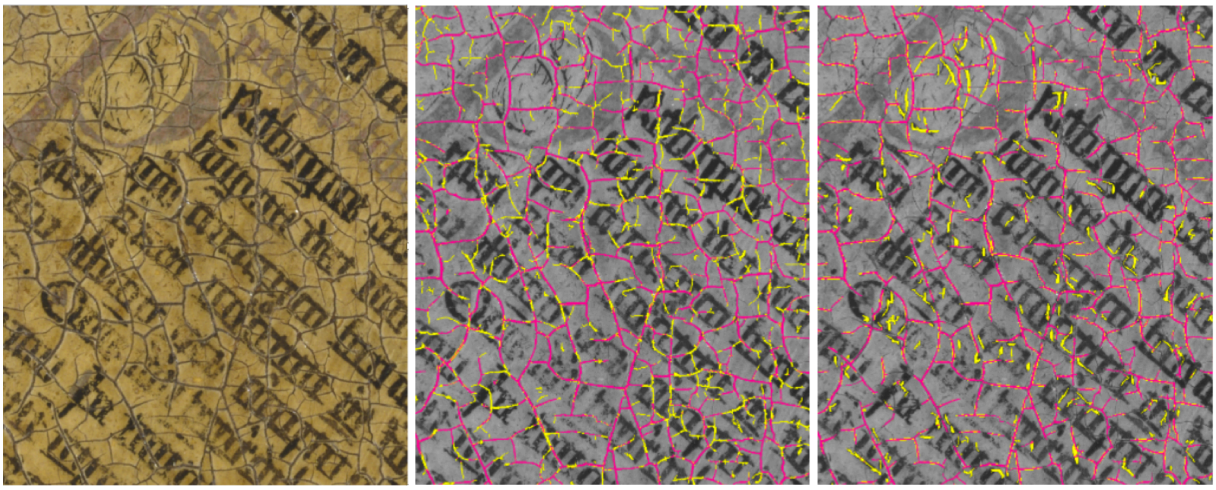

Fig. 4 Comparison between single-modal and multi-modal crack detection on a part of the book from Fig. 3. The first image shows original visual data and the other two images show crack maps. Cracks detected by *both* methods are marked in *red* and the *differences* in *yellow*. The image in the middle: yellow marks cracks detected *only by* the multi-modal approach. The image on the right: yellow marks cracks detected *only* by the single-modal approach.

Emerging deep learning tools (LeCun et al, 2015; Jacobsen et al, 2017) offer great potentials for improving further crack detection. Two main advantages of such a framework are the following. Firstly, there is no need to hand-engineer features and tune the corresponding parameters. The discriminative features and the corresponding weights are instead learned automatically from the provided examples. Secondly, the processing speed is drastically improved: once the network is trained, the processing is rapid and memory-efficient. Deep convolutional neural networks (CNN) are currently widely used in computer vision providing state-of-the-art in various image classification, and object categorization tasks (Krizhevsky et al, 2012; Simonyan and Zisserman, 2014; He et al, 2016). Two recent works applied deep networks to crack detection in roads (Lei et al 2016; Cha et al, 2017). Our work reported in (Sizyakin et al, 2018) is to the best of our knowledge the first one to apply deep learning to crack detection in paintings. A detailed analysis in (Sizyakin et al, 2018) indicates a huge potential of multi-modal deep convolutional neural networks in crack detection compared to the more traditional multi-modal approach from (Cornelis et al, 2013a; Pizurica et al, 2015) described above, both in terms of detection accuracy and speed.

A deep convolutional neural network consists of two main stages: feature extraction and classification, as illustrated in Fig. 5. In the feature extraction stage, the input images are passed through a sequence of filters followed by nonlinearities and (optionally) pooling operations to yield a sequence of feature maps. Mathematically, filtering is expressed as convolution with a given kernel, hence the name convolutional network. A nonlinearity called Rectified Linear Unit (ReLU): f(x) = max(x,0) is commonly applied as the activation function to the convolution outcomes, yielding the corresponding feature maps. Optionally, pooling is performed by taking, e.g., the maximum value from a square window of pixels. This operation progressively reduces the spatial size of the representation, which is of interest in object categorization but less so in pixel-wise labelling such as crack detection. The weights with which various feature maps are fed into the subsequent layers of convolutions as well as the parameters of the convolutional kernels are learned based on the supplied examples of the desired input-output pairs. This is called supervised learning. State-of-the-art deep networks, like those employed for face recognition, often have 16 or more layers and hundreds of millions of these adjustable weights, and hundreds of millions of labelled examples with which to train the machine (LeCun et al, 2015). For the purpose of crack detection, our network from (Sizyakin et al, 2018) employed three convolutional layers and was efficiently trained on manually labelled cracks from several details of the *Ghent Altarpiece*.

Fig. 6 illustrates the process of supplying examples of cracks and non-cracks to the system and the corresponding result of crack detection using the network architecture from (Sizyakin et al, 2018). Another example of crack detection using this new methodology is employed in the next section as input to virtual restoration of cracks in the book from Fig. 3.

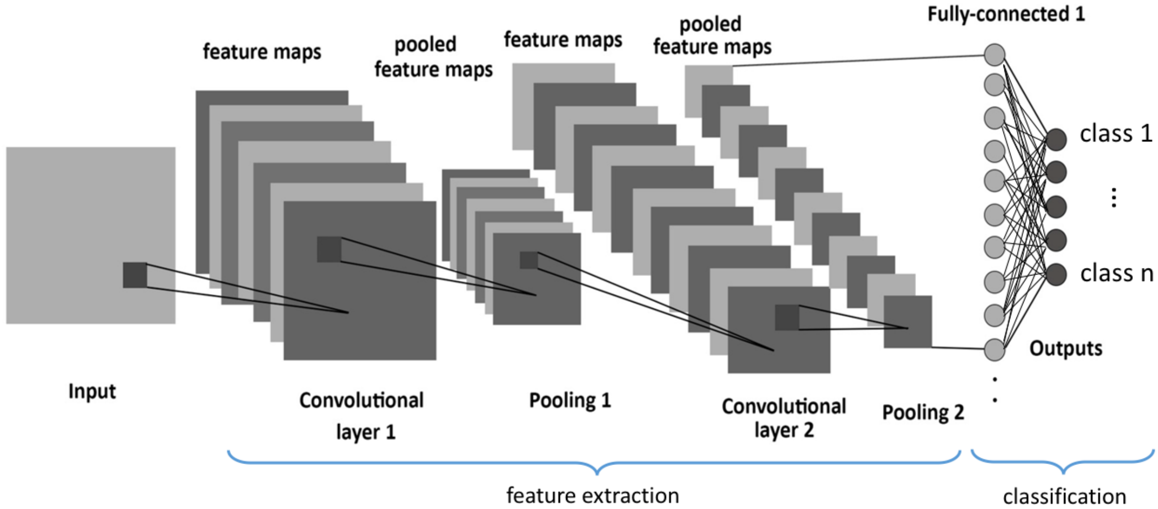

Fig. 5 A conceptual scheme of a deep convolutional neural network (see text).

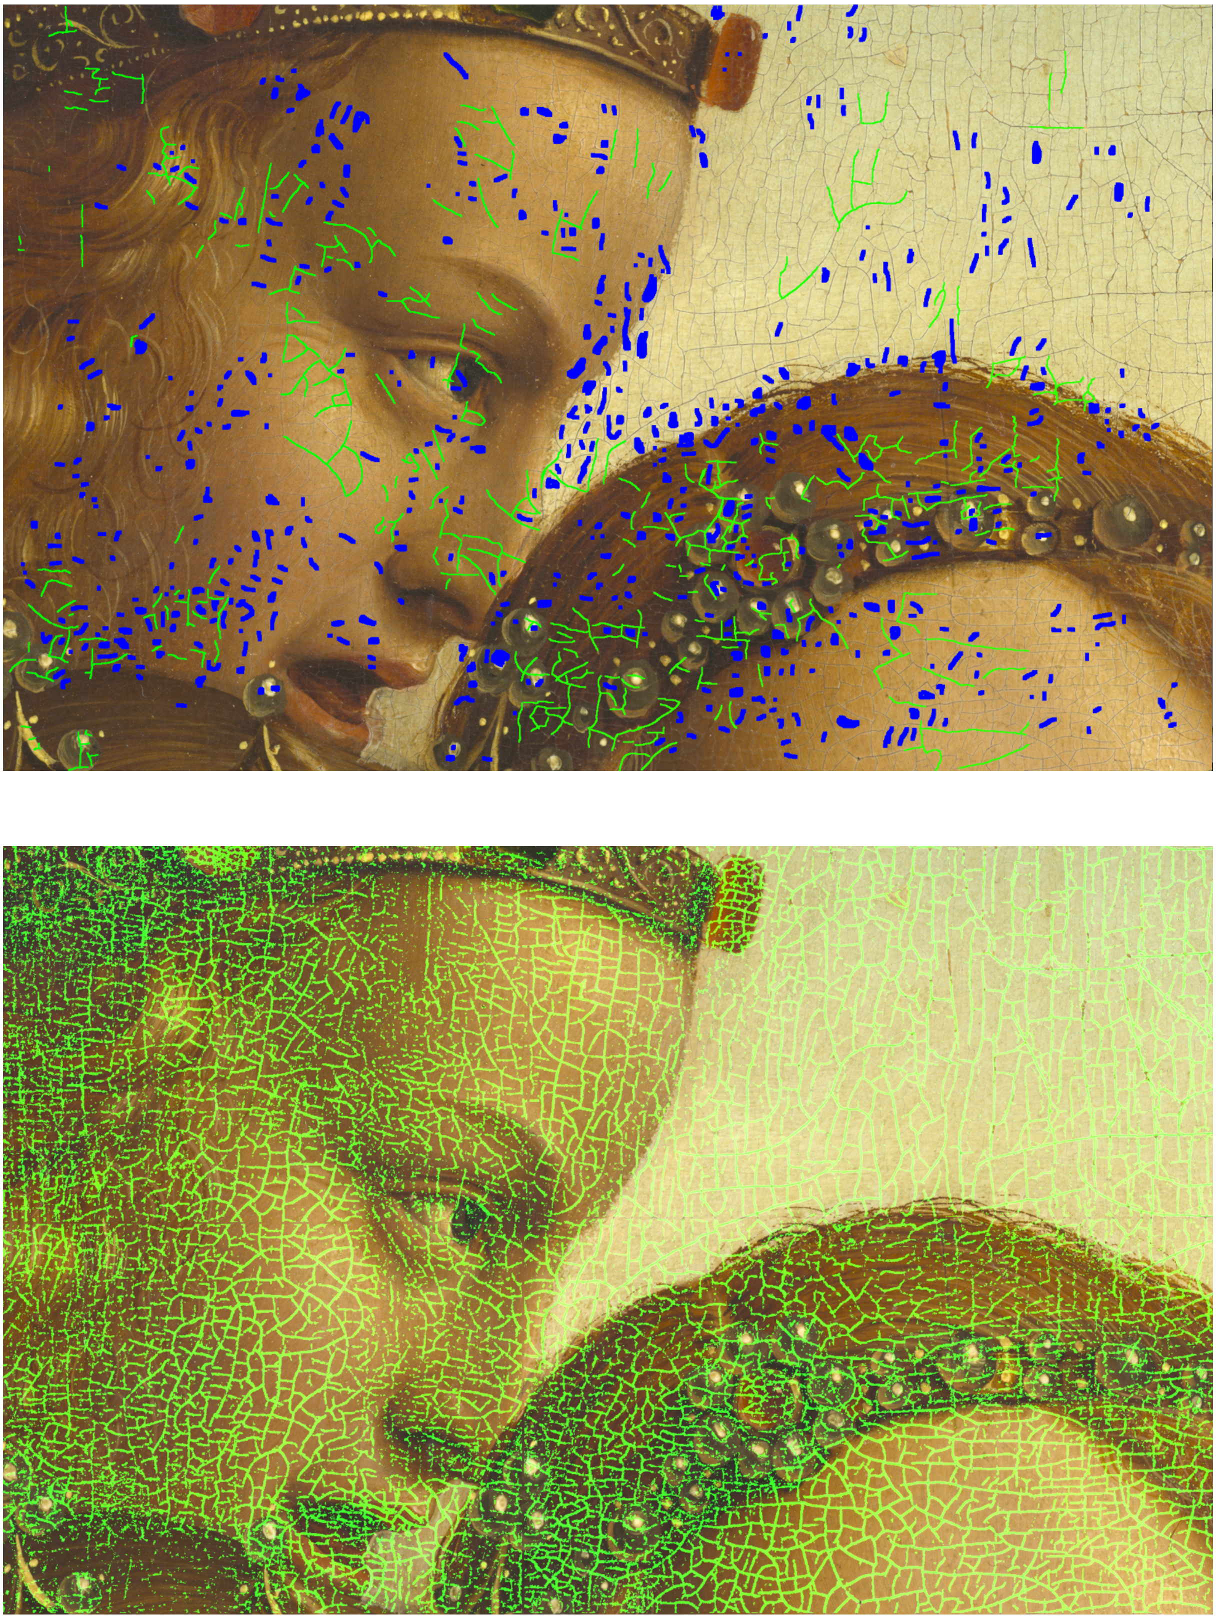

Fig. 6 An illustration of crack detection using a deep convolutional neural network (CNN) from (Sizyakin et al, 2018). **Top**: examples of cracks (green) and non-cracks (blue) provided by the user in a training stage. **Bottom**: automatic crack detection result using the trained network.

Fig. 7 A detail of *Prophet Zachariah* in four modalities infrared macrophotography, infrared reflectography and X-radiography before treatment, and the macro photograph during treatment. Image copyright: Ghent, Kathedrale Kerkfabriek, Luksaweb.

Paint loss detection

Localization of paint losses is an important part of documenting the restoration process of a painting. With current commercial tools, the restores can either do a tedious manual job or annotate the damage areas only roughly. We do not know of any reported signal/image processing techniques that address this specific problem except for the preliminary works of some of the authors, reported in (Huang et al, 2016) and (Meeus et al, 2018). Both of these techniques make use of multi-modal data, including visible images before and after treatment, infrared macrophotography, infrared reflectography and X-radiography. An example of a multi-modal data set (consisting of a subset of these available modalities) is illustrated in Fig. 7.

The approach reported in (Huang et al, 2016) makes use of the so-called sparse representation classification (SRC) framework, which was first introduced in machine learning and computer vision areas for the purpose of face recognition (Wright et al, 2009). The main idea behind this approach is to construct a dictionary of prototype functions (atoms) that can represent well signals/images from a certain class. An image is represented well in a given dictionary if each of its local regions can be reproduced as a linear combination of relatively few atoms from the dictionary. Then we say that the particular image (or class of images) has a sparse representation in the given dictionary. In the training phase, the SRC method constructs structured dictionaries, where each sub-dictionary is composed of representatives of a given class (such as, e.g. ‘paint loss’, ‘cracks’, ‘non-damaged’, etc.). In the classification stage, each test sample (a small image region or a pixel with its values across the available modalities) is assigned to the class for which the corresponding sub-dictionary gives the smallest reconstruction error. Fig. 8 shows an

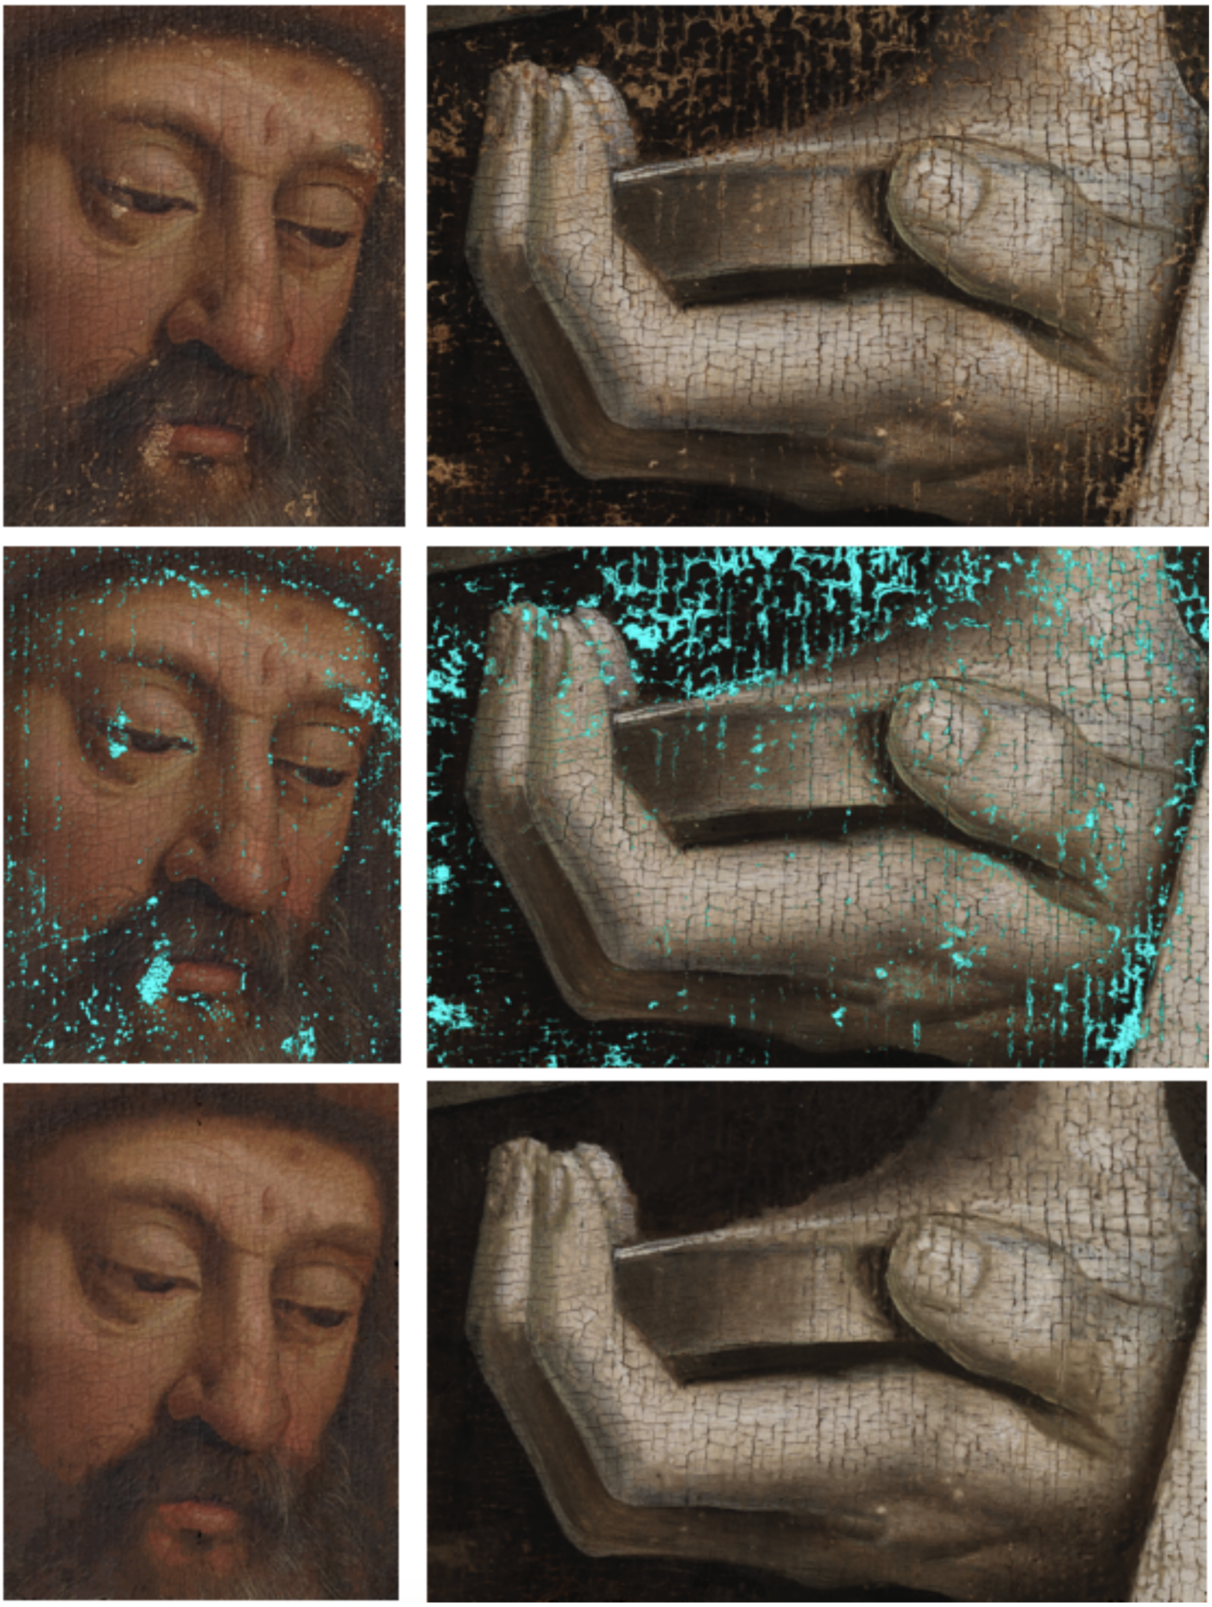

Fig. 8 Examples of automatic paint loss detection during treatment of the *Ghent Altarpiec*e, using the method of (Huang et al, 2016), followed by virtual inpainting using the method of (Ružić et al, 2015).

example of binary classification into ‘paint loss’ and ‘non paint loss’ areas obtained by this technique. The results of virtual inpainting starting from these detections are shown as well.

A deep learning method for paint loss detection described in (Meeus et al, 2018) employs a similar architecture as the crack detection method of (Sizyakin et al, 2018) presented in the previous Section. The initial results indicate similar performance in terms of detection accuracy as the sparse representation method of (Huang et al, 2016), but with the potential to process much faster large panels. Fig. 9 illustrates the application of this technique.

Fig. 9 An example of paint loss detection on the panel *John the Evangelist* using a deep convolutional neural network architecture from (Meeus et al, 2018).

Virtual inpainting

After detecting damaged parts (cracks and/or paint losses), we need to fill in the missing regions. In technical terms, this process is called *inpainting*. Digital image processing literature contains a vast number of general inpainting methods which can be roughly separated into two groups: *pixel-based* and *patch-based* methods. Pixel-based methods aim at replacing one missing pixel at the time (Bertalmio et al, 2000, Chan and Shen, 2001), while patch-based methods (Criminisi et al, 2004; Sun et al, 2005; Komodakis and Tziritas, 2007, Xu and Sun, 2010, Le Meur et al, 2011, Ružić et al, 2015) fill in the missing region patch-by-patch. Patch-based processing is typically more intensive computation-wise but is also able to reconstruct underlying textures more faithfully. Crack inpainting methods reported so far are mostly pixel-based, employing principles like order statistics filtering (Solanki and Mahajan, 2009), controlled anisotropic diffusion (Giakoumis et al, 2006) and interpolation (Barni et al, 2000). Patch-based crack inpainting was reported in (Ružić et al, 2011; Pižurica et al, 2015).

Fig. 10 illustrates the problem to be solved by inpainting. The detected cracks are shown in black; all these "missing" pixels need to be inpainted. In the following we will discuss only patch-based inpainting. The idea is to transfer small and well-chosen image parts from undamaged regions of the image to the damaged positions, such that they fit well with the surrounding context. The actual process is explained in Fig. 11, where *target* denotes partially or fully damaged image patch and *source* denotes the patch that will be transferred (as a whole or in part) to the damaged position. Candidates for good source patches are identified by comparing the undamaged part in the target patch with corresponding parts in the possible source patches. When a source patch is identified, the needed part can be extracted to fill in the target, which results in a "composite" patch. There are a lot of possible variations of this method. For example, instead of replacing only the damaged part, a wider area (or in extreme cases the complete source patch) can be transferred in order to avoid visually disturbing discontinuities. For the same reason some local averaging of multiple source patches and the underlying non-damaged regions can be applied. However, this needs to be done with care. In most cases, preserving the authenticity of the undamaged parts should be more important than obtaining a visually pleasing result.

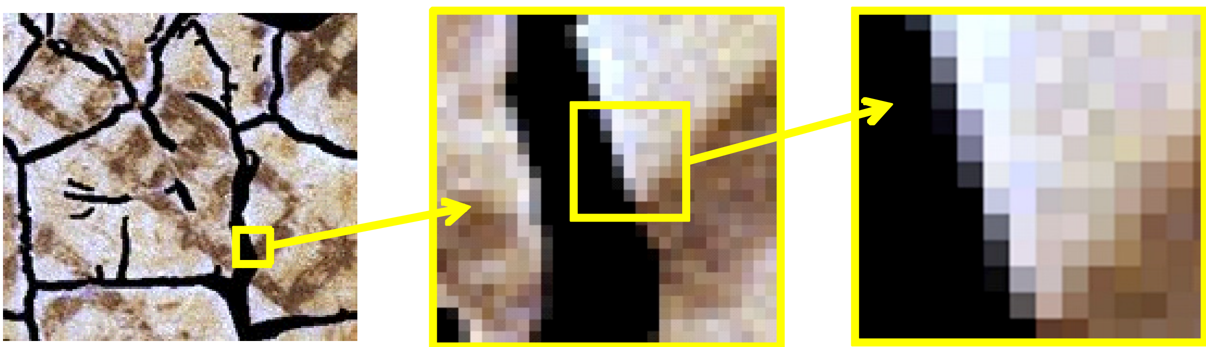

Fig. 10 From left to right: a fragment of the text from Fig. 3 (with detected cracks shown in black) and enlarged details. The inpainting process has to "fill in" all the black pixels.

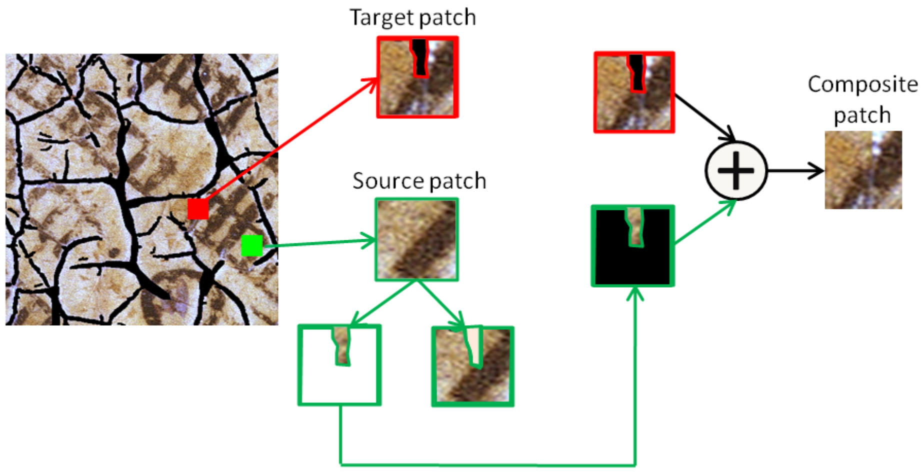

Fig. 11 An illustration of the patch-based inpainting process. The damaged part in the “target” patch is to be replaced by a region of the same shape from an undamaged “source” patch.

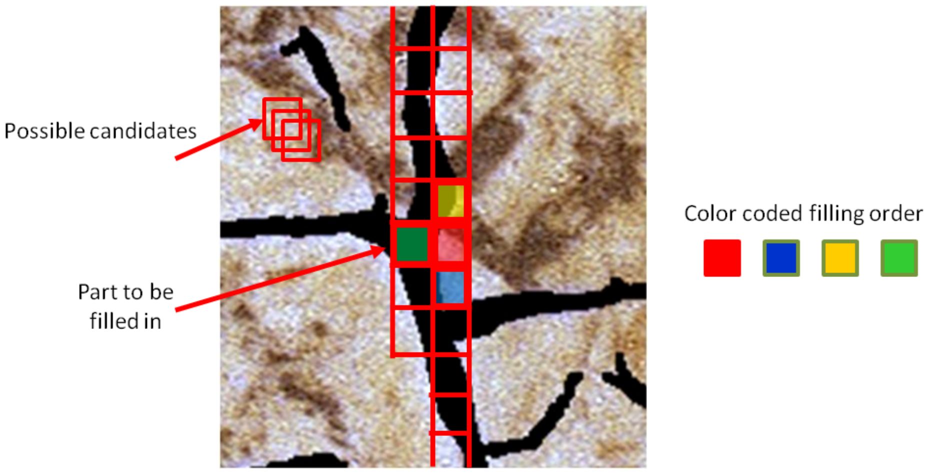

Fig. 12 An illustration of defining filling in crack inpainting (see text).

Patch-based inpainting methods can be further categorized in terms of the employed optimization process as *local* or *global*. The local optimization takes only one best suited candidate (patch) for each spatial position, while the global methods consider (or combine) multiple candidates in terms of a wider spatial context. In both cases it is important to define some *priorities* for filling the patches. Especially for local methods, the filling order is very important. In the case of a larger missing region, the inpainting process should always start from border areas, i.e. from those target patches that contain undamaged parts and preferably some distinctive structures (like edges) that can guide the search for correct source patches. Once these patches are filled in, the process continues towards the inner regions (see Fig. 12). Of course, this process is prone to error propagation and in this respect global methods are advantageous over the local ones.

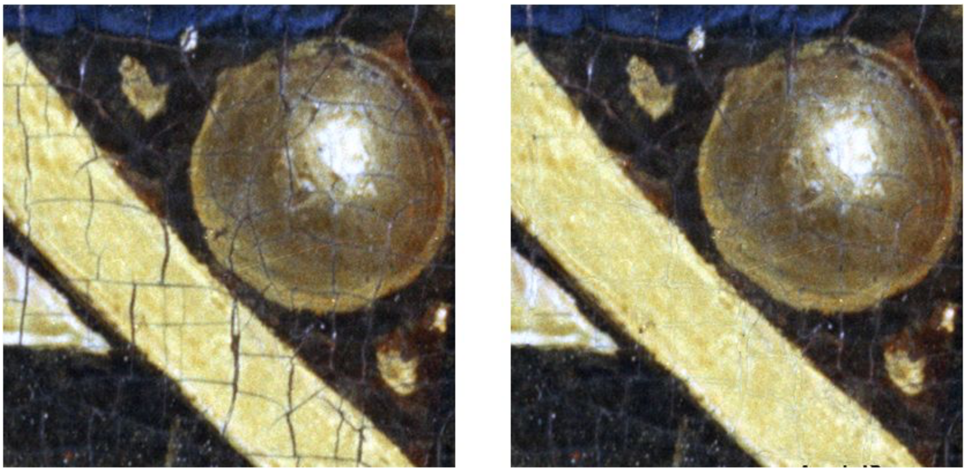

Fig. 13 An example of patch-based inpainting, using the method of (Ružić et al, 2011). The original fragment on the left is an enlarged part of the broach from Fig. 2.

An example of crack inpainting using the global patch-based approach from (Ružić et al, 2011) is shown in Fig. 13 on an enlarged part of the broach from Fig. 2. Notice that this virtual inpainting method is conservative in terms of preserving painted features and the authenticity of the painting. While large cracks have been removed some of the subtle ones remained and the underlying painting composition is kept intact.

Our recent context-aware patch-based inpainting method from (Ružić et al, 2015) improves further inpainting performance by incorporating textural descriptors and statistical modelling of the spatial context. Its application in virtual restoration of paint loss areas was already illustrated in Fig. 8. This method has been evaluated and further optimized for painting restoration in (Pižurica et al, 2015). Here we apply our context-aware approach to virtually restore the book from the panel *Annunciation to Mary*. In some of our previous publications (Ružić et al, 2011; Cornelis et al, 2013; Pižurica et al, 2015) parts of the inpainted book with different methods have been reported. In this Chapter, we report for the first time the complete inpainted text. Also, we used here for the first time as input a crack map produced by a deep learning technique, and in particular with our method from (Sizyakin et al, 2018) described earlier in this Chapter. Fig. 14 shows the original high-resolution image (before the last restoration treatment) and Fig. 15 shows the inpainted result. The inscriptions appear indeed clearer after inpainting. While the deciphering of this text is still under investigation, virtual inpainting proves already useful for enhancing the appearance of the characters. This preprocessing of the image should facilitate its further automatic analysis and visual inspection.

Fig. 14 Book in the panel *Annunciation to Mary*.

Fig. 15 Patch-based inpainting of the book in Fig. 14 with the method of (Ružić et al, 2015) and using as input the crack map produced by a deep learning method (Sizyakin et al, 2018).

Analysis of pearls

Another interesting application of digital image processing is the development of objective metrics that characterize and discriminate between styles of different painters. The *Ghent Altarpiece* is very rich in depictions of precious stones and jewels, which provide a magnificent test set for the experiments. For our impression of the beauty and realism of a painterly representation, as well as for distinguishing between the painter styles, the *brightness distribution* is usually much more important than the color attributes. A mathematical description of the spatial distribution of brightness in a painted pearl was introduced in (Platiša et al, 2011) and further elaborated in (Platiša et al, 2012, Pižurica et al, 2015). This description is related to the so-called *spatiogram* representation (Birchfield et al, 2005; Conaire et al, 2007), which extends the classical *histogram* with some information about the spatial arrangement of the pixel intensities. An image histogram tabulates the relative number of occurrences of a given grey value (brightness), or a range of grey values within that image. Each of these grey value ranges constitutes a *bin* (each bin will have an index b within some set $\{1,\dots, B\}$). By means of *bin counts* (i.e. the number of pixels with a grey value in a particular bin) c_b, a histogram characterizes only *global* distribution of grey values without capturing their spatial relationships. In other words, the histogram tells us which *fraction* of image pixels belongs to a certain brightness range but gives us no clue about their spatial positioning or spreading.

An image spatiogram gives certain information about spatial positioning of the pixel attributes next to their global distribution. Next to the bin count c_b, two additional parameters exist in the spatiogram representation: *spatial center of the bin* $\boldsymbol{\mu}_b$, i.e. the mean value of the spatial positions of all the pixels that have intensities within the given bin, and the amount of *spatial spreading* (covariance of the spatial positions) $\boldsymbol{\Sigma}_b$ of those pixels. Note that the spatial center is a two-dimensional vector, with components along the horizontal direction (say, x-direction) and vertical (y-direction) in the image, or formally: $\boldsymbol{\mu}_b = (\mu_{b,x}, \mu_{b,y})$; $\boldsymbol{\Sigma}_b$ is the covariance matrix (with diagonal elements $\sigma^2_{b,xx}$ and $\sigma^2_{b,yy}$). To enable comparison between regions of different sizes, all spatial coordinates are normalized to the same range. Visualization of these high-dimensional data is not trivial. An effective visualization of the spatiogram features was introduced in (Platiša et al, 2011), with three types of plots, illustrated on Fig. 16:

(S1) connected spatial centers of bins $\boldsymbol{\mu}_b$;

(S_2) $\boldsymbol{\mu}_b$ – positioned bin counts (circle size proportinal to c_b);

(S_3) $\boldsymbol{\mu}_b$ –positioned bin variances (x and y error bars of length $\pm\sigma_{xx}$ and $\pm\sigma_{yy}$, resp.).

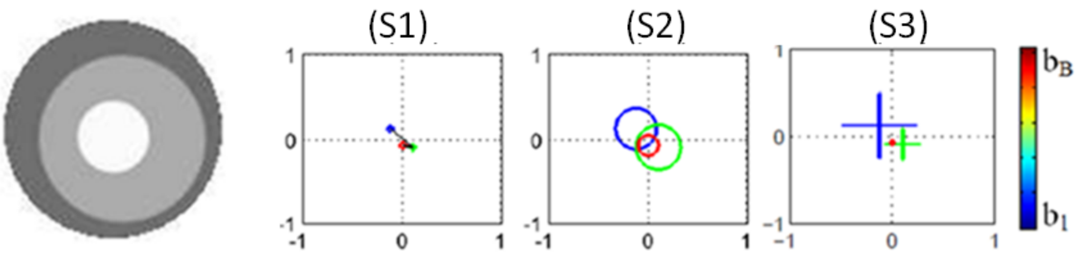

Fig. 16 An illustration of spatiogram plots of (Platiša et al, 2011). Different colors correspond to different bins. (S1): spatial centers of bins. (S2): corresponding bin counts (propotional to circle sizes); (S3): spatial variances within each bin.

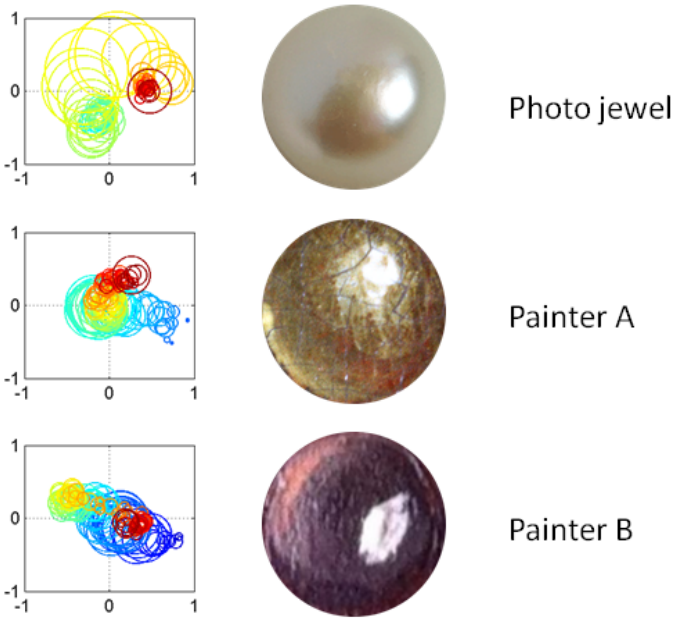

Fig. 17 Spatiogram plots of type S2 are used as digital signatures for pearl images. Different colors represent bins with a particular grey value range, from 0 (blue) to 255 (red). From these plots we can read the spatial centers of different bins and the relative number of pixels within each bin (proportional to the size of the corresponding circle).

The spatiogram-based representation of painted pearls from (Platiša et al, 2011) provides a kind of *digital signature* of the painter (see Fig. 17). Ideally we would like to have these signatures sufficiently consistent for the pearls of one painter and different enough for different painters. Fig. 18 shows a number of pearls extracted from the *Ghent Altarpiece* together with their spatiogram-based digital signatures. When looking at these pearls, we are overwhelmed by their beauty and realism; they are all different, but still so consistent. What is more important from a mathematical analysis point of view is that their digital signatures are pretty consistent as well. In particular, for pearls that are very similar in lighting, we can establish a high similarity between their signatures (see examples in Fig. 19). It is interesting to investigate whether digital signatures for pearls painted by other painters, and especially the copyists of Van Eyck, will be significantly different. Before answering this question, it is useful to look at a somewhat more compact, summarized description that can be extracted from the analyzed digital signatures. Such a compact description was proposed in (Platiša et al, 2011) as an ensemble of four metrics M1 – M4 (corresponding to four numbers for each observed pearl spatiogram) and can be associated with certain visually observable features of the painted pearl, such as the degree of assymmetry or smoothness. These four metrics, given in Table 1, are derived from the centers of spatiogram bins, $\boldsymbol{\mu}_i$, and Euclidean distances between the centers of adjacent bins D_i, where $D_i^2 = \mu_i^2 - \mu_{i+1}^2, \;\; i = 1,\ldots,B-1$. M1 (mean of the distances D_i) captures the (circular) symmetry of the pearl area: the smaller M1 the higher the symmetry (which might be related to the angle between the light source and the pearl surface). M2 (variance of the distances D_i), relates to the uniformity of distances between different bin areas, and thereby to the impression of the surface smoothness. M3 and M4 relate to the ranges of bin centers in the *x*- and *y*-directions, respectively; these metrics tell us about the (dominant) orientation of the asymmetry in the painted pearl.

Table 1. Spatiogram metrics for pearls images from (Platiša et al, 2011).

Symbol	Name	Definition
M1	Mean (D)	$\frac{1}{N}\sum_i D_i$ $\frac{1}{B-1}\sum_i D_i, \quad i = 1,\ldots,B-1$
M2	Var (D)	$\frac{1}{B-1}\sum_i (D_i - M1)^2$
M3	R_x	$\max_i \mu_{x,i} - \min_i \mu_{x,i}$ $\max_i \mu_{x,i} - \min_i \mu_{x,i}$
M4	R_y	$\max_i \mu_{y,i} - \min_i \mu_{y,i}$

Table 2. Spatiogram metrics in identifying artists (Platiša et al, 2012).

	Van Eyck		Van der Veken	
M1	0.210	±± 0.029	0.136	± 0.013
M2	0.055	± 0.019	0.009	± 0.002
M3	0.767	± 0.229	0.733	± 0.028
M4	0.899	± 0.210	0.457	± 0.139

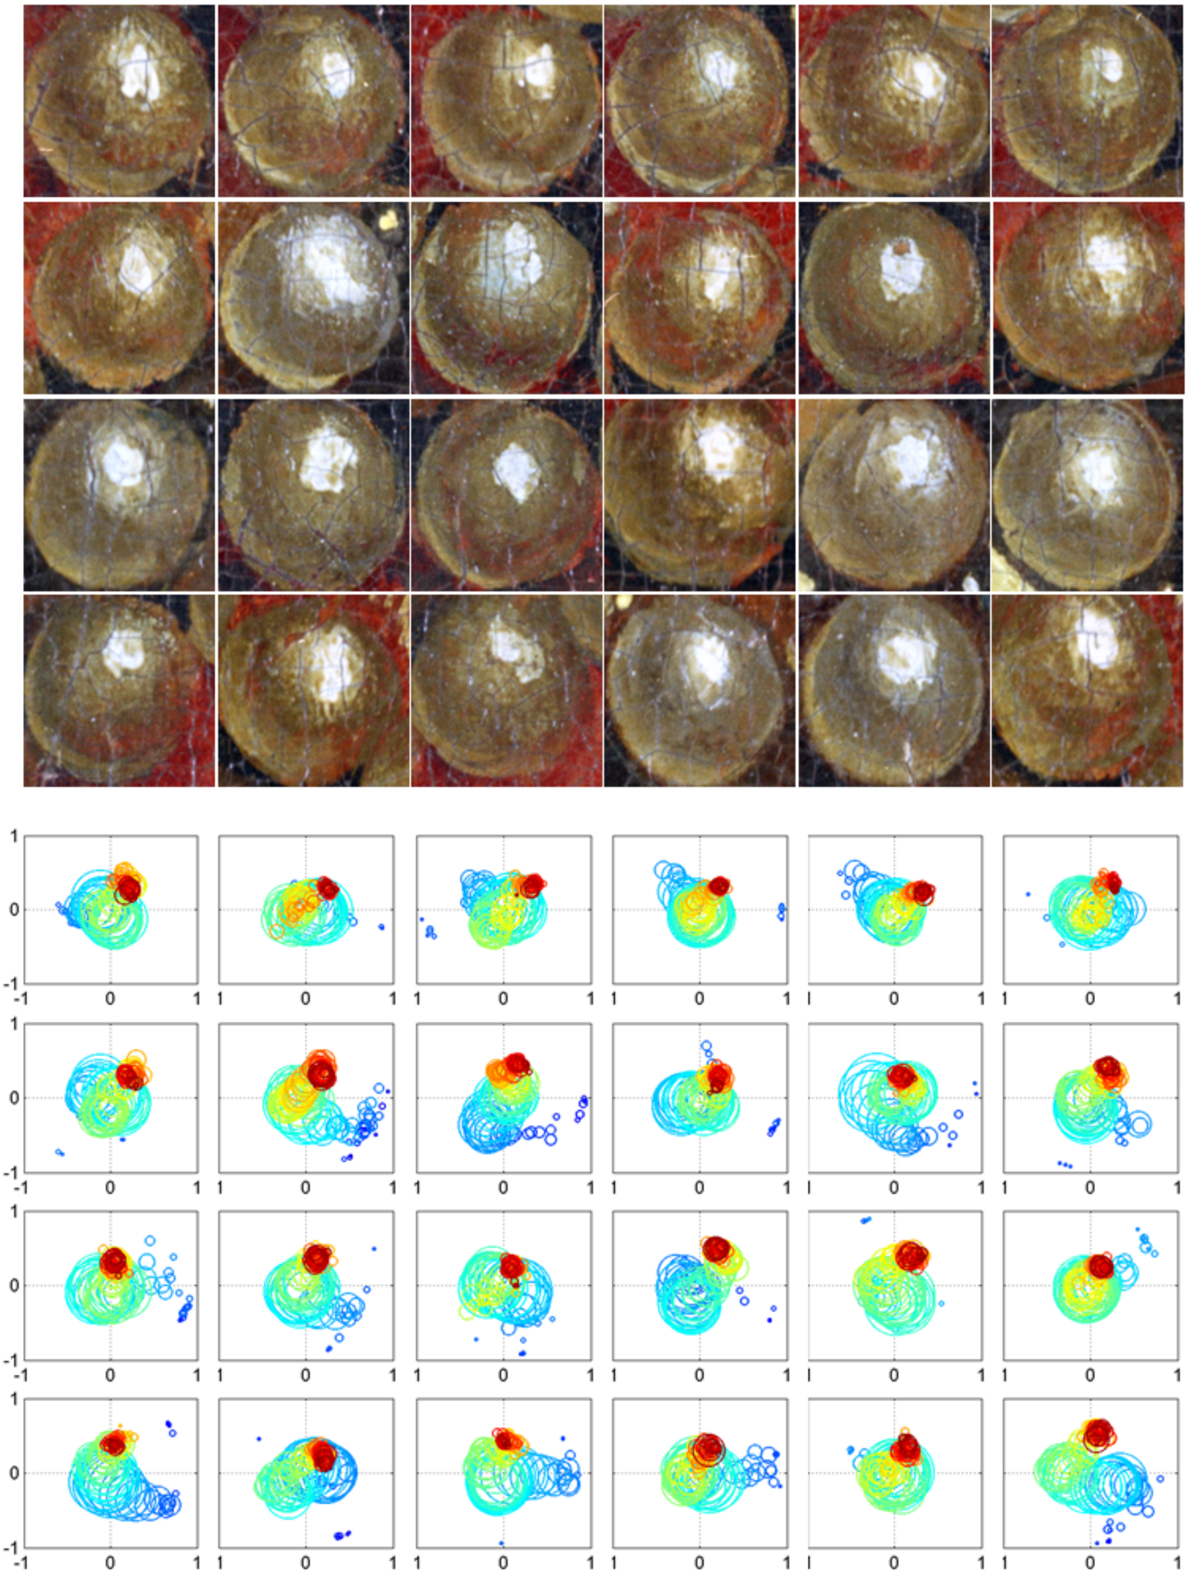

Fig. 18 A selection of pearls from the *Ghent Altarpiece* and their digital signatures presented as spatiogram plots of (Platiša et al, 2012).

The results reported in (Platiša et al, 2012) indicate a great potential of the spatiogram-based digital signatures and the derived M1 - M4 metrics to distinguish between different painter styles. A clear difference was shown between the signatures extracted from the pearls of Van Eyck compared to those of Jef Van der Veken from the *Just Judges* panel (see Table 2 and Fig. 20). Also, clear differences were established between pearls painted by Charlotte Caspers (2010) in the copy of the panel *Angels Playing Music* (2010) and other examined paintings, including pearls from *Maria Maddalena Baroncelli* of Hans Memling (1470). Our recent work, reported in (Pižurica et al, 2015) addressed the consistency of painted pearls in the *Ghent Altarpiece* on the new high-resolution data set. The results of that study suggest that the pearls that are larger, more visible and have, a more central place in the panel are indeed painted in a more consistent manner. The study pointed also to some some instances (specific painted objects) that are exceptions to this general rule and that might be of interest to art-historians, and conservators to examine in more detail. The interested reader is referred to our work in (Pižurica et al, 2015) for details.

While digital signatures are in the first place interesting for distinguishing between different painters, they might also find other interesting applications. Think for example of the possibility to modify the digitized scan such that the painting or objects in it become more similar to the style of another painter. We might try to in a way "Van Eyckify" the pearls in the *Just Judges* copy of Van der Veken, so that they appear as if they were painted by Van Eyck. This idea is illustrated in Fig. 21. Imagine that we perturb the pixels in the pearl painted by one painter, and keep perturbing them, until the resulting digital signature becomes similar to that of another painter. The idea obviously can be applied to different objects and different painter styles and give rise to a number of nice applications as well.

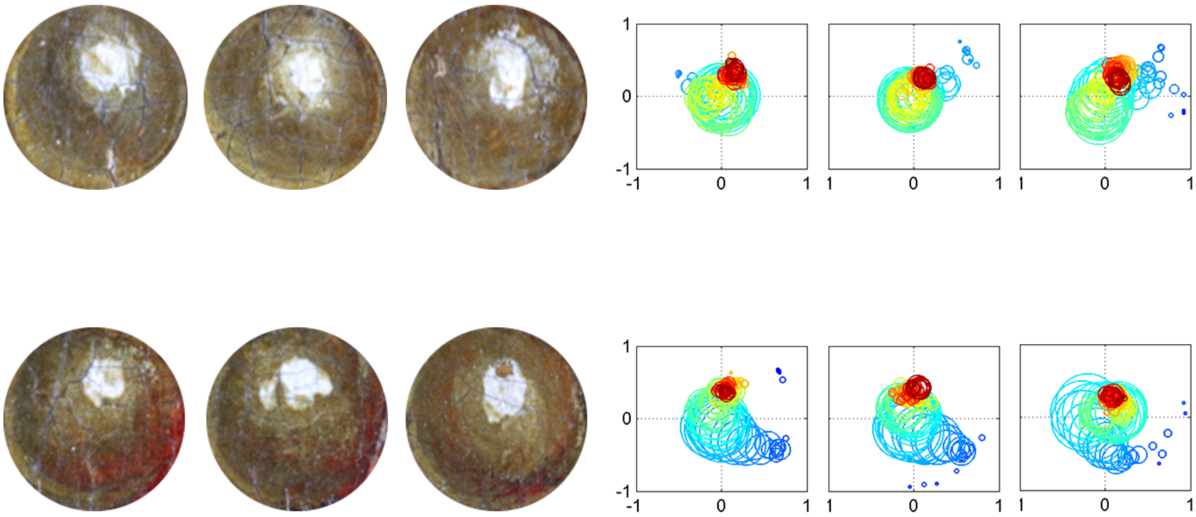

Fig. 19 Examples extracted from Fig. 18 showing two groups of mutually similar pearls with their digital signatures. The example illustrates that the spatiogram-based digital signatures are very similar for pearls with similar lighting.

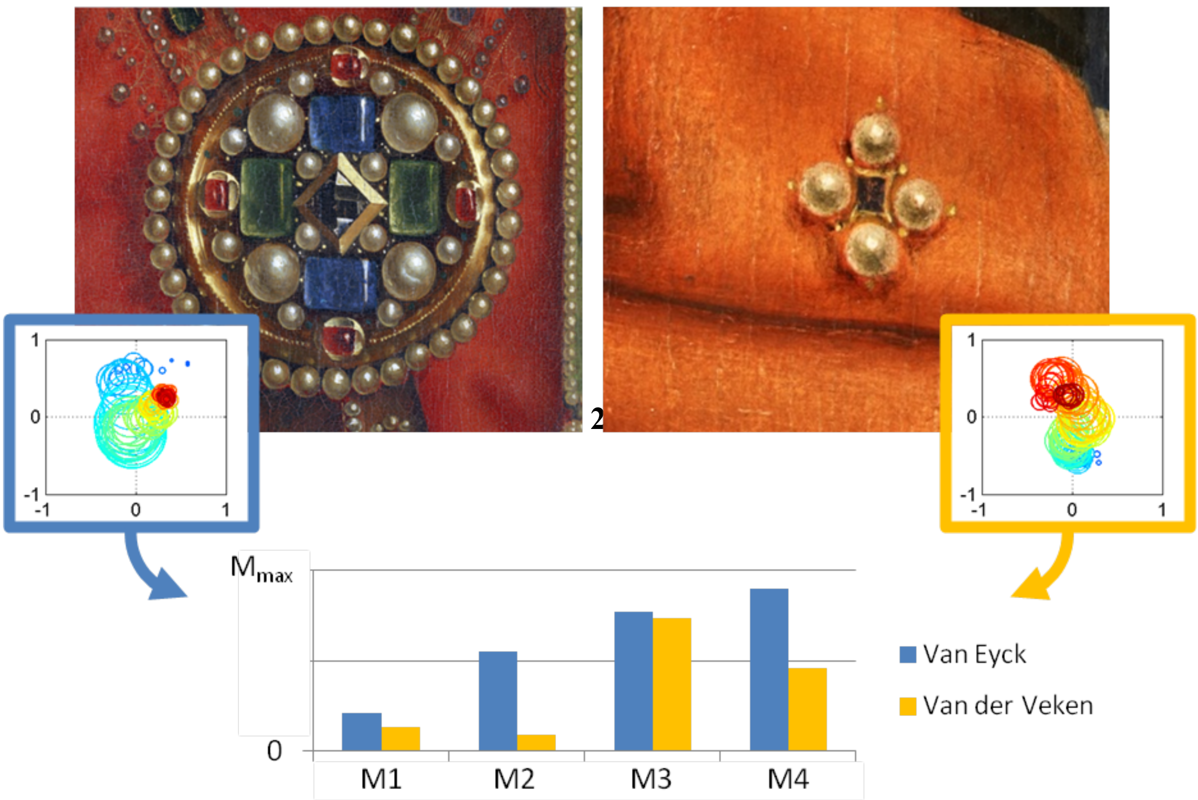

Fig. 20 A comparison of the spatiograms and the M1-M4 metrics derived from pearls of Van Eyck (broach on the panel *God the Father*, left) and the pearls of Van der Veken (detail of the panel *Just Judges*, right). The bar chart with M1 – M4 metrics corresponds to Table 2, where the M2 values were scaled by a factor of 10 for better visualization.

Conclusion

This Chapter showed some novel applications of digital image processing techniques for the virtual restoration and the mathematical analysis of the *Ghent Altarpiece*. All the discussed techniques are applicable to other paintings by old masters as well. The presented results indicate that digital image processing can indeed provide useful support in the conservation/restoration treatments by offering automatic detection of crack patterns and paint loss areas, as well as virtual restoration. A great potential of emerging deep learning tools has been demonstrated, as well as that of sparse coding and patch-based processing. We reported for the first time here digital inpainting of the complete book from the *Annunciation to Mary* panel. After inpainting the inscriptions appear enhanced, which should facilitate further analysis and deciphering of this text. The potentials of mathematical analysis for distinguishing the painter styles were shown and discussed based on a case study with pearls. In this study, spatiogram-based digital signatures enable clear differentiation between different painters. Some other interesting applications of digital image processing and digital signatures are mentioned, such as bringing the painted object closer to the style of another painter. Still, the main interest in extracting digital signatures of painters lies in authentication. In the case of the *Ghent Altarpiece* the development of such techniques may contribute in the future to gaining more insights into the division of hands between Jan and Huber Van Eyck and their respective workshops

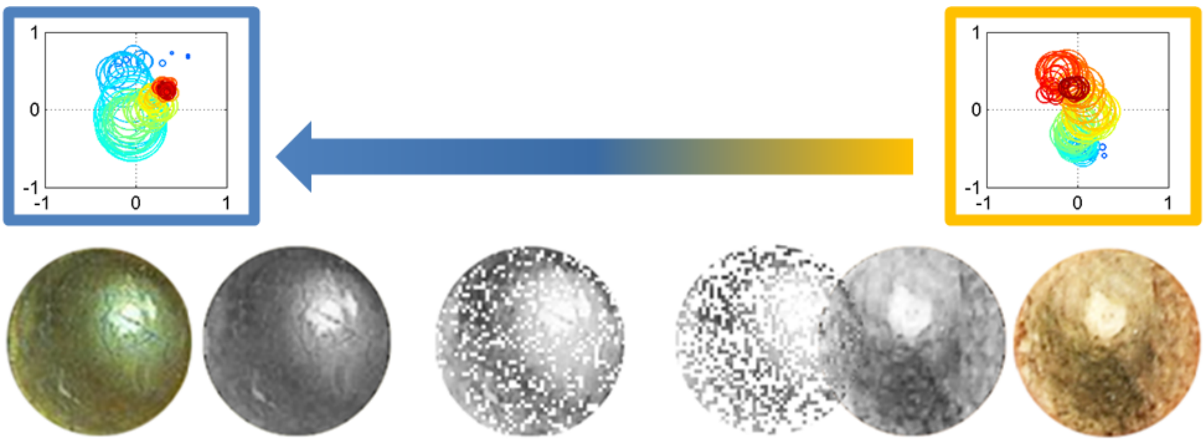

Fig. 21 Another possible use of spatiogram-based digital signatures. Imagine that we perturb the pearl image (e.g., by randomly switching pixels) until the resulting spatiogram becomes similar to the spatiogram of a Van Eyck pearl or to one of another painter such that the object appears as if this other painter painted it.

Acknowledgement

The authors are most grateful to art restorers Bart Devolder and Hélène Dubois working on the *Ghent Altarpiece*, for annotating data in the process of training our methods for crack detection and paint loss detection, and for their insightul comments and inspiring discussions.

Bibliography

[Abry et al, 2015] P. Abry, A. G. Klein, W. A. Sethares, and C. R. Johnson. Signal processing for art investigation. IEEE Signal Processing Magazine, Vol. 32, No. 4, pp. 14–16, July 2015.

[Barni et al, 2000] M. Barni, F. Bartolini, V. Cappellini, Image processing for virtual restoration of artworks, IEEE Multimedia, Vol. 7, pp. 34-37, 2000.

[Bertalmio et al, 2000] M. Bertalmio, G. Sapiro, V. Caselles, C. Ballester, Image inpainting, in SIGGRAPH 2000, New Orleans, USA.

[Berns, 2001] R. S. Berns, The Science of Digitizing Paintings for Color-Accurate Image Archives: A Review, Journal Of Imaging Science And Technology, Vol. 45, No. 4, pp. 305-325, July/August 2001.

[Birchfield et al, 2005] S. T. Birchfield and S. Rangarajan, Spatiograms versus histograms for region-based tracking, in Proceedings IEEE Conf. CVPR, 2005, vol. 2, pp. 1158–1163.

[Bucklow, 1997] S. Bucklow, The description of craquelure patterns, Studies in Conservation Vol. 42, 1997.

[Bucklow, 1998] S. Bucklow, A stylometric analysis of craquelure, Computers and the Humanities, Vol. 31, No. 6, pp. 503-521, 1998.

[Carreras, 2006] 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.

[Cha et al, 2017] Y.-J. Cha, W. Choi, and O. Büyüköztürk, Deep learning-based crack damage detection using convolutional neural networks, Comput.-,Aided Civ. Infrastruct. Eng., vol. 32, pp. 361–378, 2017.

[Chan and Shen, 2001] T. F. Chan, J. Shen, Non-texture inpainting by curvature-driven diffusions (CDD), Journal of Visual Communication and Image Representation, Vol. 12, pp. 436-449, 2001.

[Conaire et al, 2007] C. O. Conaire, N. E. O'Connor, and A. F. Smeaton, "An improved spatiogram similarity measure for robust object localisation," in Proceedings IEEE Conf. ICASSP, 2007, vol. 1, pp. I–1069 –I–1072.

[Cornelis et al, 2013] B. Cornelis, T. Ružić, E. Gezels, A. Dooms, A. Pižurica, Lj. 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, Vol. 93, No. 3, pp. 605-619, March 2013.

[Cornelis et al, 2013a] 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.

[Criminisi et al, 2004] A. Criminisi, P. Perez, K.Toyama, Region filling and object removal by exemplar-based image inpainting, IEEE Transactions on Image Processing, Vol. 13, pp. 1200-1212, 2004.

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

[Elad et al, 2006] M. Elad, M. Aharon, A. Bruckstein, The K-SVD: An algorithm for designing of overcomplete dictionaries for sparse representation, IEEE Trans. on Signal Processing, Vol. 54, No. 11, pp. 4311-4322, 2006.

[Feichtinger and Strohmer, 1998] H. G. Feichtinger and T. Strohmer, *Gabor Analysis and Algorithms*, Birkhäuser, 1998.

[Giakoumis and Pitas, 1998] I. Giakoumis, I. Pitas, Digital restoration of painting cracks, in Proceedings of the IEEE International Symposium on Circuits and Signals ISCAS 98 (1998), pp. 269-272.

[Giakoumis et al, 2006] I. Giakoumis, N. Nikolaidis, I. Pitas, Digital image processing techniques for the detection and removal of cracks in digitized paintings, IEEE Transactions on Image Processing, Vol. 15, pp. 178-188, 2006.

[He et al, 2016] K. He, X. Zhang, S. Ren, and J. Sun, Deep residual learning for image recognition, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770-778.

[Huang et al, 2016] S. Huang, W. Liao, H. Zhang, and A. Pižurica, Paint Loss Detection in Old Paintings by Sparse Representation Classification, in Proc. iTWIST, Aug 24-26, Aalborg, Denmark, 2016.

[Ibarra-Castanedo et al, 2010] C. Ibarra-Castanedo, S. Sfarra, D.Ambrosini, D. Paoletti, A. Bendada and X. Maldague, Infrared vision for the nondestructive assessment of panel paintings, *CINDE Journal,* Vol.31, No.5, Sept/Oct 2010.

[Jacobsen et al, 2017] J.-H. Jacobsen, E. Oyallon, S. Mallat, A. W.M. Smeulders, "Multiscale Hierarchical Convolutional Networks," arXiv:1703.04140v1, Mar 2017.

[Komodakis and Tziritas, 2007] N. Komodakis, G. Tziritas, Image completion using efficient belief propagation via priority scheduling and dynamic pruning, IEEE Transactions on Image Processing, Vol. 16, pp. 2649-2661, 2007.

[Krizhevsky et al, 2012] A. Krizhevsky, I. Sutskever, and G. Hinton, Imagenet classication with deep convolutional neural networks, Advances in Neural,Information Processing Systems, p. 10971105, 2012.

[Larue et al, 2007] F. Larue, L. Ammann and J. –M. Dischler, A Pipeline for the Digitization and the Realistic Rendering of Paintings, The 8th International Symposium on Virtual Reality, Archaeology and Cultural Heritage VAST (2007), pp. 1–9.

[Le Meur et al, 2011] O. Le Meur, J. Gautier, and C. Guillemot, Examplar-based inpainting based on local geometry, in IEEE Conf. ICIP, 2011, pp. 3462–3465.

[LeCun et al, 2015] Y. LeCun, Y. Bengio, and G. Hinton, "Deep Learning," Nature, vol. 521, no. 7553, pp. 436–444, 2015.

[Lei et al, 2016] Z. Lei, Y. Fan, D. Yimin, and J. Z. Ying, “Road crack detection using deep convolutional neural, IEEE International Conference on Image Processing (ICIP), 2016, pp. 3708–3712.

[Livingstone, 2002] M. Livingstone, *Vision and Art: The Biology of Seeing*, 2002.

[Meeus et al, 2018] L. Meeus, S. Huang and A. Pižurica, Paint loss detection using Fast Deep Convolutional Neural Networks, Technical Report, Ghent University, January 2018.

[Mohen et al, 2006] J. Mohen, M. Menu, B. Mottin, *Mona Lisa: Inside the Painting*, Harry N. Abrams, New York, NY, USA, 2006.

[Pižurica et al, 2015] A. Pižurica, L. Platisa, T. Ružic, B. Cornelis, A. Dooms, M. Martens, H. Dubois, B. Devolder, M. De Mey and I. Daubechies, “Digital Image Processing of the Ghent Altarpiece: Supporting the painting's study and conservation treatment", IEEE Signal Processing Magazine, vol. 32, no. 4, pp. 112-122, July 2015.

[Platiša, et al. 2011] Lj. Platiša, B. Cornelis, T. Ružić, A. Pižurica, A. Dooms, M. Martens, M. De Mey, I. Daubechies, ‘Spatiogram features to characterize pearls in paintings’, In 18th IEEE International Conference on Image Processing (ICIP 2011), 11-14 Sept. 2011, pp. 801-804.

[Platiša, et al. 2012] Lj. Platiša, B. Cornelis, T. Ružić, A. Pižurica, A. Dooms, M. Martens, M. De Mey, I. Daubechies, Spatiogram features to characterize pearls and beads and other small ball-shaped objects in paintings, in *Vision and Material*, M. De Mey, M. Martens, and C. Stroo, Eds, Royal Flemish Academy of Belgium, 2012.

[Poli and Valli, 1997] R. Poli, G. Valli, An algorithm for real-time vessel enhancement and detection, Computer methods and programs in biomedicine, Vol. 52, No. 1, pp. 1-22, 1997.

[Ružić et al, 2011] T. Ružić, B. Cornelis, L. Platiša, A. Pižurica, A. Dooms, W. Philips, M. Martens, M. De Mey, I. Daubechies, Virtual restoration of the Ghent Altarpiece using crack detection and inpainting, in Proceedings of Advanced Concepts for Intelligent Vision Systems ACIVS (2011), Ghent, Belgium, Lecture Notes in Computer Science, 2011, Vol. 6915/2011, pp. 417-428.

[Ružić et al, 2015] T. Ružic and A. Pižurica, Context-aware patch-based image inpainting using Markov random field modelling, IEEE Transactions on Image Processing, vol. 24, no. 1, pp. 444-456, Jan 2015.

[Serra, 1984] J. Serra, *Image Analysis and Mathematical Morphology*, Academic Press, 1984.

[Simonyan and Zisserman, 2014] K. Simonyan and A. Zisserman, Very deep convolutional networks for large-scale image recognition, CoRR, vol. abs/1409.1556, 2014.[Online]. Available: http://arxiv.org/abs/1409.1556.

[Sinisgalli, 2006] R. Sinisgalli, *Alberti’s ”De Pictura”. Il Nuovo De Pictura di Leon Battista Alberti, The New De Pictura of Leon Battista Alberti*, Edizioni Kappa, Roma, 2006.

[Sizyakin et al, 2018] R. Sizyakin, B. Cornelis, L. Meeus, M. Martens, V. Voronin, and A. Pižurica, Crack detection in paintings using convolutional neural networks, Technical Report, Ghent University, January 2018.

[Solanki and Mahajan, 2009] 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), Vol. 1, pp. 97-99, 2009.

[Spagnolo and Somma, 2010] G. S. Spagnolo and F. Somma, Virtual restoration of cracks in digitized image of paintings, Journal of Physics: Conference Series, Vol. 24, No. 1, 2010.

[Sun et al, 2005] J. Sun, L. Yuan, J. Jia, and H.-Y. Shum, Image completion with structure propagation, ACM Transactions on Graphics, Vol. 24, pp. 861-868, 2005.

[Xu and Sun, 2010] Z. Xu and J. Sun, Image inpainting by patch propagation using patch sparsity, IEEE Transactions on Image Processing, Vol. 19, pp. 1153-1165, 2010.

[Walmsley et al, 1994] E. Walmsley, C. Metzger, J.K. Delaney and C. Fletcher, Improved visualization of underdrawings with solid-state detectors operating in the infrared, *Studies in Conservation,* Vol. 39, No. 4, pp. 217-231, Nov 1994.

[Wright et al, 2009] J. Wright, A. Yang, A. Ganesh, S. Sastry, and Y. Ma. Robust face recognition via sparse representation. IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 31, No. 2, pp. 210–227, 2009.