

Edge-Based Image Restoration

Andrei Rares, *Member, IEEE*, Marcel J. T. Reinders, and Jan Biemond, *Fellow, IEEE*

Abstract—In this paper, we propose a new image inpainting algorithm that relies on explicit edge information. The edge information is used both for the reconstruction of a skeleton image structure in the missing areas, as well as for guiding the interpolation that follows. The structure reconstruction part exploits different properties of the edges, such as the colors of the objects they separate, an estimate of how well one edge continues into another one, and the spatial order of the edges with respect to each other. In order to preserve both sharp and smooth edges, the areas delimited by the recovered structure are interpolated independently, and the process is guided by the direction of the nearby edges. The novelty of our approach lies primarily in exploiting explicitly the constraint enforced by the numerical interpretation of the sequential order of edges, as well as in the pixel filling method which takes into account the proximity and direction of edges. Extensive experiments are carried out in order to validate and compare the algorithm both quantitatively and qualitatively. They show the advantages of our algorithm and its readily application to real world cases.

Index Terms—Edge structure reconstruction, image restoration, inpainting, sequentiality, T junctions.

I. INTRODUCTION

A N important part of the scientific and cultural heritage of the modern times has been stored in the form of film and photo archives. Unfortunately, the classic storage media for these information sources are bound to gradually decay in time, risking the total loss of the valuable information they are carrying. Fortunately, with the arrival of the digital era, the digitized films and photographs can now be copied easily and virtually without information loss. An equally important aspect is the opportunity of doing restoration in superior ways, never possible in the past. As such, information that disappeared completely from its physical support can now be restored thanks to advanced algorithms developed in the restoration community. Modern technologies have brought along economical benefits, too. The digitized content is now cheaper and easier to store,

Manuscript received November 29, 2003; revised August 3, 2004. This work was supported by the EU's IST research and technological development program. It was carried out within the BRAVA project ("Broadcast Archives Restoration Through Video Analysis") [8]. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Vincent Caselles.

A. Rares was with the Information and Communication Theory Group, Mediamatics Department, Faculty of Electrical Engineering, Mathematics and Computer Science, Delft University of Technology, Delft, The Netherlands. He is now with the Leiden University Medical Center, Department of Radiology, Division of Image Processing, 2300 RC Leiden, The Netherlands (e-mail: a.rares@lumc.nl).

M. J. T. Reinders and J. Biemond are with the Information and Communication Theory Group, Mediamatics Department, Faculty of Electrical Engineering, Mathematics and Computer Science, Delft University of Technology, 2600 GA Delft, The Netherlands (e-mail: m.j.t.reinders@ewi.tudelft.nl; j.biemond@ewi.tudelft.nl).

Digital Object Identifier 10.1109/TIP.2005.854466

search, reuse and distribute. The digitally broadcasted content will take full advantage of these benefits. But today's quality requirements will only be fulfilled when digital restoration is applied to the degraded archived content.

The current restoration algorithms for image sequences exploit both spatial and temporal information, and perform quite well, in general [1], [2]. However, they fail when there is difficult motion in the sequence [3], [4], in particular, for the detection and correction of blotches. Blotches are artefacts typically related to film that are caused by the loss of gelatin, or the presence of dirt particles on the film [5]. Due to the difficult object movements, wrong motion vectors are extracted from the sequence. As a result, the spatiotemporal restoration process that follows may introduce unnecessary errors that are visually more disturbing than the blotches themselves. The extracted temporal information becomes unreliable, and a source of errors itself.

Instead of protecting the blotches from being restored [4], in our view, the detected artefacts should be restored based on spatial information alone [6], [7], discarding the temporal information. In the BRAVA project [8] of the European Union, we have devised a novel restoration algorithm that takes advantage of the available spatial information in order to restore the degraded film frames. This algorithm is not intended to replace the spatiotemporal algorithms, rather, to complement them in places where they fail. Because of its spatial nature, the algorithm can also be applied for the restoration of missing areas in damaged (usually old) photographs, for the automatic interpolation of damaged pixels from the CCD sensors in new digital cameras, or for the concealment of errors in compressed data due to transmission errors. Another area of application is the reconstruction of occluded objects when they are partly covered by other objects. This can be useful for assessing the correctness of a segmentation procedure, as well as for determining the relative depths of objects [9]. The proposed algorithm only deals with the problem of filling in the missing information.

The task of artefact detection represents a separate problem. With some exceptions, the artefact detection and the restoration are usually treated in different algorithms. In this paper, we explicitly assume that the artefact mask is detected by another algorithm and contains no holes.

A. Related Work

Several spatial restoration approaches for missing data have been proposed already in the literature. They address the problem of filling in missing data from different points of view. In the following, a short categorized overview is given that presents the most popular approaches.

Restoration Based on Partial Differential Equations and Variational Methods: A recent category of algorithms centered around the idea of "image inpainting" has shown promising

results on restoring image structure for large, piecewise smooth patches. Masnou and Morel, for example, present in [10] a simple, but effective, approach for filling in missing areas based on the connection of the level lines (i.e., isophotes) that have the same values on the artefact contour. The method was further developed in [11]. In [12]–[16], Ballester *et al.* and Bertalmio *et al.* propose more complex variational approaches for joint interpolation of grey levels and gradient/isophotes directions. In [17]–[19], Chan *et al.* present several inpainting methods based on total variation models, curvature driven diffusions, and Euler’s elastica. In [20] and [21], Bertalmio *et al.* further refine the aforementioned methods by trying to combine in one algorithm different approaches for structure and texture interpolation.

Structure-Based Restoration: Atzori and De Natale propose a spatial restoration method for recovering missing blocks (corresponding to data packets) in video transmission over packet networks in [22]. They use only the information existing in the same frame, by making a “sketch” of the edges around the missing blocks. These edges are connected in a pairwise fashion, if possible, and a smooth interpolation takes place subsequently in the areas delimited by these sketched edges. While this paper uses a spline interpolation to recover the shapes of the edge connections, in [23], they present an alternative based on Bezier curves. In [24], Atzori *et al.* present a spatiotemporal algorithm which first uses a temporal interpolation and then applies a spatial, mesh-based warping to reduce the temporal restoration errors mainly caused by complicated motion. In [7], we present a spatial algorithm for the reconstruction of artefacts based on explicit information about the surrounding edges. The main assumption here is that edges are (locally) straight. Simple edge information is extracted from the image and used to recover the edges inside the artefact. The straight edges reconstructed inside the artefact are then used to guide a smooth interpolation between the edges.

Convolution- and Filter-Based Restoration: With their normalized and differential convolution, Knutsson and Westin [25] defined a general method for interpolating N -dimensional data through convolutions based only on valid data. Their approach is more general and flexible than some restricted convolution, by allowing the association of certainty values to each data point and an applicability operator to the filters to be applied. In [26], Khriji *et al.* presented a restoration technique based on spatial rational filters.

Texture-Based Restoration: In [27], Efros and Leung present a nonparametric texture synthesis algorithm based on Markov random fields. Their approach restores pixels based on the similarity between their local neighborhood and the surrounding neighborhoods. From the candidate neighborhoods, one is randomly selected and the value of its central pixel is pasted at the current location, a process which is able to intelligently imitate the natural randomness of textures. Bornard *et al.* [28] have further developed the aforementioned texture synthesis for image sequences by incorporating temporal information and imposing some local consistency constraints which allow the algorithm to also synthesize structured objects that do not have random appearances. In [29], a method is presented by Criminisi *et al.* that also extends the approach of Efros and

Leung by imposing higher priorities in the restoration order for pixels lying in the neighborhood of edges, thereby preserving better edge sharpness and continuity. In [30], Kokaram presents a parametric texture synthesis algorithm which employs two-dimensional autoregressive models (combined with the Gibbs sampler) in a Bayesian approach. In [31] and [32], he introduces a more general framework for restoring image sequences, based on the Markov chain Monte Carlo methodology. A solution is proposed for jointly detecting and restoring missing data and motion vectors, while also handling occlusion and uncovering situations. In [33], Jia and Tang describe a novel technique based on tensors. Here, edge structure is first reconstructed, followed by texture synthesis. Both steps use adaptive tensor voting. Another way of synthesizing texture is presented in [34] by Acton *et al.* Their approach is based on a diffusion generated by partial differential equations, and a simultaneous reaction based on Gabor filters and AM-FM dominant component analysis. In [35], Hirani and Totsuka combine spatial and frequency information to reconstruct the missing image structure/texture, in a framework of projection onto convex sets.

Connections With Proposed Method: Our approach relates most to the sketch-based method of Atzori and De Natale [22]. It generalizes the algorithm presented in [7] and employs higher level features extracted from the image. Our approach also bears some similarity with the algorithm of Jia and Tang [33] in what concerns the main steps of the algorithm. Each of these steps is, however, differently approached. The novelty of our method consists of the approximation of the incoming edges with circle arcs, the use of the spatial order of edges, and the directional interpolation scheme that restores missing areas parallel to the recovered edges. As opposed to the classic texture-based restoration algorithms, which do not preserve object shapes, we prefer (together with Atzori and De Natale and Jia and Tang) to use explicit edge information to capture the image structure. Our main motivation comes from two observations. On the one hand, edges generally separate areas with different content. Therefore, the interpolation should take place independently on both sides of an edge. On the other hand, edges are more robust against intensity changes such as local shading, thereby being more robust than isophote-based algorithms, for example.

Throughout this paper, we compare our proposed algorithm with the related restoration scheme of Atzori and De Natale, both qualitatively as well as quantitatively. We also present a qualitative comparison with the algorithm of Masnou [11], which uses a variational approach applied to the image isophotes.

B. Outline

In Section II, we present the main steps of the algorithm. Section III concentrates on how the structure of the missing areas is recovered. Section IV describes our interpolation method, which takes into account the structure recovered in the previous section. Section V is devoted to presenting and discussing experimental restoration results, as well as comparisons with other algorithms. Finally, Section VI draws conclusions and outlines future work.

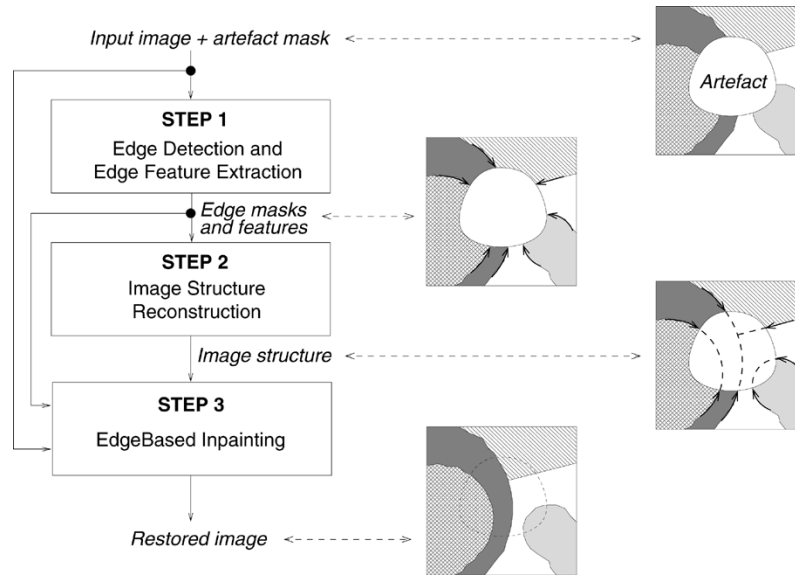


Fig. 1. (Left) General algorithm outline and (right) an illustration of the inputs/outputs for each stage.

II. ALGORITHM OVERVIEW

The spatial restoration algorithm that we propose consists of three main steps, depicted in Fig. 1: 1) edge detection and edge feature extraction; 2) image structure reconstruction; and 3) edge-based inpainting.

The input to our algorithm is an image and an artefact mask. Here, we assume that the artefact mask is detected by another algorithm. For the sake of simplicity, but without loss of generality, in the remainder of this paper we consider that the mask consists of only one artefact, and that the image is grey-valued.

Assuming that the artefact location, size and shape are independent from the image content, then the structure of the original image inside the artefact area is a continuation of the structure outside it. More specifically, the edges inside the artefact are continuations of the outside edges. We, therefore, use the edge information explicitly to guide the restoration process.

In the first step, edges are detected around the artefact, based on the contours of the segments that result from a watershed segmentation. Ideally, these edges separate two objects (or at least two different homogeneous regions), both of which are partially occluded by the artefact. The object edges are extracted in clockwise order, from a point of view lying inside the artefact. Simple edge features are extracted for each edge, such as the luminance values on both sides of the edge, and the local gradient magnitude along the edge. Only relevant edges are then kept for the next steps (e.g., those that have at least a certain gradient magnitude).

In the second step, we try to recover the structure of the image within the artefact area. This problem is ill-posed: Virtually anything could have existed in the area covered by the artefact, before the degradation took place. We have to “invent” content in places where it was lost, based on some assumptions about the usual image properties. In our case, we have modeled the edges as locally circular shapes (equivalent to a second order polynomial). This modeling was subject to several constraints, such as color matching, and noncrossing of the object edges. Our model tries to couple edges that are strongly related to each other, thereby reconnecting the pairs of edges that were part of




Fig. 2. Intensity feature for a group formed by the edges couples $A_1 - A_2$, $B_1 - B_2$, and $C_1 - C_2$.

the same object contour. The matching of the edges is on the one hand based on the similarity of the aforementioned edge features, and on the other hand on *continuity* and *sequentiality* criteria. For an edge couple (e.g., $A_1 - A_2$, $B_1 - B_2$, or $C_1 - C_2$ in Fig. 2), the continuity is measured by fitting a circle to the pair of edges and measuring the goodness of the fit (e.g., the spatial deviation of both edges from the fitted circle). Unlikely edge couples are ignored and the remaining ones are iteratively joined into edge groups (see Fig. 4). An edge group is a (as large as possible) set of consecutive edge couples such that no two couples cross each other. Based on the set of possible groupings, specific configurations that represent potential image structures within the artefact are created. Each configuration is then rated by its “sequentiality,” which is a measure indicating the likeliness of a particular configuration (essentially trying to minimize the number of crossing couples). The score of a configuration is based upon the sequentiality together with the other features that estimate the continuity and similarity of edges. The best configuration is then found by selecting the configuration that minimizes this score. After finding the best configuration, spare edges [e.g., $E_1 - E_2$ in Fig. 5(a) and (b) or $E_2 - E_B$ in Fig. 7(c)], i.e., edges that were not included in any edge couples, are traced one by one into the selected configuration. They are traced up to the point where they meet another edge (or edge

couple), or, alternatively, they gradually vanish toward the opposite side of the artefact. In this way, the structure of the image is recovered inside the artefact area.

Finally, in the third step, the artefact is restored by inpainting, taking into account the recovered image structure. Essentially, the inpainting procedure restores a pixel based on the surrounding recovered edges. The surrounding edges indicate which pixels on the artefact border are used for the interpolation. Then, based on the distance to these border pixels, the pixel inside the artefact is interpolated.

The sketch-based interpolation of Atzori and De Natale [22] follows the three main steps presented in Fig. 1. Since the first step does not concern the restoration directly, we address the differences between our scheme and the one of Atzori and De Natale for the other two steps. In step two, the differences concern the set of features used, and, additionally, the way we combine them in order to characterize the overall acceptability of the reconstructed structure. In the last step, our interpolation method tries to draw strips “parallel” to the nearby edges, resulting in smooth patches. Atzori and De Natale have used a “patch repetition” approach, in which the areas around the artefacts are mirrored across the artefact edge. Many other smaller differences between the two methods exist as well in the above steps (e.g., in step two, the way we normalize the values of different features in order to bring them in the same range).

III. IMAGE STRUCTURE RECONSTRUCTION

The structure reconstruction step is crucial to our proposed restoration scheme, since the explicit image structure that is recovered represents the “skeleton” of the restoration process. The input to this step represents a list of *edges* coming into the artefact, in clockwise order. The output of this step is a list of *edge couples* arranged in *groups* of edges, and a list of *spare edges*.

To build accurate pairwise connections between edges, we make use of *local* features, as well as *global* features. Local features describe how well two edges match each other if they were part of the same edge couple. Global features express the goodness of a complete configuration of edge couples. The local features are 1) the two luminance values on both sides of each edge in the edge couple, 2) the local gradient magnitudes of both edges, and 3) the degree to which the edge couple fits a common circle. The global feature expresses the degree to which edge couples do not cross each other within a configuration. The overall cost c^{cfg} of a particular configuration is given by

$$c^{cfg}(\mathbf{Z}) = \frac{4 \times c^{loc}(\mathbf{Z}) + c^{glb}(\mathbf{Z})}{5} \quad (1)$$

where \mathbf{Z} represents the configuration of groups of edge couples, c^{loc} is the cost related to the four local features, and c^{glb} is the cost associated with the single global feature. All costs have values between 0 and 1, with 0 indicating a perfect match and 1 indicating a complete mismatch. The process of building up the final configuration \mathbf{Z} is presented later in this section.

A. Local Features

Before specifying the couple-related costs, we first give intensity and magnitude representation of an edge. The set of intensities (λ_i) on the clockwise side of the edges is given by:

$\mathbf{\Lambda} = \{\lambda_i = \hat{m}(\mathbf{I}(\mathbf{b}_i)) \mid i = 1, \dots, N^E\}$, where \mathbf{I} represents the intensity image, \mathbf{b}_i is the vector of pixels on the artefact border between edge i and the (clockwise) next edge, \hat{m} is the median operation, and N^E is the number of edges. The set of edge gradient magnitudes (γ_i) is given by: $\mathbf{\Gamma} = \{\gamma_i = \tilde{m}(\mathbf{J}(\mathbf{e}_i), [N^{\text{MAX}}, N^{\text{MAX}} - 1, \dots, 2, 1]) \mid i = 1, \dots, N^E\}$, where \mathbf{J} is the gradient magnitude of \mathbf{I} (obtained after some smoothing, in order to remove noise), \mathbf{e}_i is the ordered vector of edge pixels, with its *head* lying on the artefact border and its *tail* stretching outwards, \tilde{m} is a weighted median operation, $[N^{\text{MAX}}, N^{\text{MAX}} - 1, \dots, 2, 1]$ represents the vector of weights used by \tilde{m} , giving more weight to the edge pixels near the artefact border, and N^{MAX} is the maximum number of pixels in an edge (we used $N^{\text{MAX}} = 12$ in our experiments). $\mathbf{\Gamma}$ and $\mathbf{\Lambda}$, together, are not redundant, since γ_i is not always directly related to $|\lambda_i - \lambda_{i-1}|$. In fact, γ_i indicates the smoothness of edge i .

The weighted median is used to calculate $\mathbf{\Gamma}$ because, as we get farther from the artefact, it is obvious that the local properties of the edge tend to become less and less related to the missing edge inside the artefact. It may also happen that a third object present in the image lies close to the artefact, without touching it. In this case, the actual edge is partly occluded, and the detected edge bends to follow the border of the third object. As a result, the edge tail is not related to the structure to be recovered in the artefact. Weighing the tail less than the head tries to overcome this situation.

The cost related to the local features of a configuration \mathbf{Z} , c^{loc} , is computed by averaging the costs c_{ij} of every edge couple within that configuration

$$c^{loc}(\mathbf{Z}) = \frac{\sum_{i=1}^{N^G} \sum_{j=1}^{|\mathbf{G}_i|} c_{ij}}{\sum_{i=1}^{N^G} |\mathbf{G}_i|} \quad (2)$$

where \mathbf{G}_i , $i = 1, \dots, N^G$ are the groups of edge couples in configuration \mathbf{Z} , $|\mathbf{G}_i|$ is the number of edge couples in group \mathbf{G}_i , and c_{ij} is the individual cost of edge couple \mathbf{q}_{ij} (the j th couple of group \mathbf{G}_i) [see (3), shown at the bottom of the next page].

The cost c_{ij} of a specific couple \mathbf{q}_{ij} indicates how well the two edges within the couple match each other, i.e., whether they describe the border of the same object. Since they belong to the same object, it seems natural to require that the intensities on both sides of the edges have similar values [first two terms in (3)] and that the strength of the edges match as well [third term in (3)]. Further, we assume that the object edges continue each other smoothly, without abrupt changes of direction (fourth term in the same equation). The cost c_{ij} is then defined by the formula in (3), with λk_{ij}^s representing the intensity on side s of edge k ($k \in \{1, 2\}$) in the couple, as shown in Fig. 2. The *side* index s indicates whether the intensity belongs to the side lying in *clockwise* ($s = C$), or in *trigonometrical* ($s = T$) direction. γk_{ij} represents the gradient magnitude along edge k of couple \mathbf{q}_{ij} . The intensity and gradient subscript notations are different here in order to reflect the affiliation of the edge to couple \mathbf{q}_{ij} from group \mathbf{G}_i . $\beta_{ij}^s \in \{0, 1\}$ are flags indicating whether the next edge on side s of edge “1” from edge couple \mathbf{q}_{ij} belongs to




Fig. 3. Behavior of the circle-fitting related measures. (a) Both the spatial deviation and the angular consistency indicate a good edge couple. (b) The spatial deviation indicates a good edge couple, whereas the angular consistency indicates a bad one. (c) Both the angular consistency and the aperture quality indicate a good edge couple. (d) The angular consistency indicates a good couple, whereas the aperture quality indicates a bad one.

a couple in the same group ($\beta_{ij}^s = 1$) or not ($\beta_{ij}^s = 0$). These binary flags effectively switch off cost contributions of the respective luminances and gradients in places where they are rendered irrelevant by spare edges (e.g., spare edge E_1 between edge couples $B_1 - B_2$ and $C_1 - C_2$ in Fig. 2 prevents the comparison of $\lambda 1_B^C$ with $\lambda 2_B^B$), or by edge couples from other groups. ω_{ij} is the cost of fitting a circle to couple \mathbf{q}_{ij} .

Let us discuss in more detail why the fourth term in (3), ω_{ij} , is essential. First, when an edge couple has spare edges on both sides, none of the first three features is of any help. Therefore, we need a supplementary feature in order to be able to do the matching. Second, when there are more objects with similar appearance that are occluded by the artefact (e.g., the fingers of a hand), the three intensity-based features alone are not sufficient to discriminate between them. Third, exploiting the continuity of the edges within a couple can help in selecting the right couples. For example, the shape of the potential edge couple $e_1 - e_8$ in Fig. 4(b) is less *natural* than the shape of the couple $e_1 - e_{12}$. Obviously, the reconstruction of object shapes is an ill-posed problem which we need to avoid. We do it by putting constraints on the edge reconstruction. Namely, we use smoothness and convexity constraints that we implement by means of a model fitting (to ensure reliable parameters).

The *naturalness* of a couple is a psychological term, rather than a physical measurement. It describes the way humans perceive the edge continuation, and not the deviation from a theoretically objective ground truth (which does not exist in practice). Naturalness is discussed in the Gestalt theory on perceptual grouping, the grounds of which were laid as early as 1923 by psychologist Max Wertheimer [36]. This theory has shown that some visual cues, such as proximity, similarity, good continuity, closure, etc., allow us to group parts of an image into objects (or groups of related objects). For example, in Fig. 4(b), the naturalness of edge couple $e_1 - e_{12}$ is expressed by a combination of properties such as similar local direction (i.e., tangent) and constant curvature.

These observations led us to define the naturalness of a couple by how well they fit a circle.¹ The cost ω_{ij} of fitting a circle to the edge couple is defined by

$$\omega_{ij} = 1 - \delta_{ij} \times \varphi_{ij} \times \theta_{ij} \quad (4)$$

where δ_{ij} is the *spatial deviation* of the couple from the fitted circle, φ_{ij} is the *angular consistency* factor, and θ_{ij} is the *aperture quality* factor. ω_{ij} returns values between 0 (ideal case) and 1 (worst case). For the δ_{ij} , φ_{ij} and θ_{ij} parameters, the significance of these values is reversed (0 represents the worst case, while 1 represents the ideal case). This enables us to propagate a “worst case” value identified with either δ_{ij} , φ_{ij} or θ_{ij} , to the circle fitness measure ω_{ij} .

The *spatial deviation*, δ_{ij} , indicates how far on the average the edge pixels lie with respect to the fitted circle. First, a distance δ_{ij}^* that represents the median of the distances from the edge pixels in couple \mathbf{q}_{ij} to their closest points on the fitting circle is defined

$$\delta_{ij}^* = \hat{m} (|d([e_{ij}^1, e_{ij}^2], \mathbf{x}_{ij}^{\odot}) - r_{ij}^{\odot}|) \quad (5)$$

where e_{ij}^1 and e_{ij}^2 are the two edges of couple \mathbf{q}_{ij} , concatenated here in a single vector for the median operation, r_{ij}^{\odot} and \mathbf{x}_{ij}^{\odot} are the radius and the center of the fitting circle, respectively, and $d(\cdot)$ represents the euclidean distance. In order to bring the value of δ_{ij}^* between 0 and 1, we use the following normalization:

$$\delta_{ij} = \frac{K^\delta}{\delta_{ij}^* + K^\delta} \quad (6)$$

where $K^\delta = 2.16$ is a constant chosen to calibrate δ_{ij} in such a way that if it is above a predefined threshold, it indicates a valid edge couple.

¹To avoid numerical problems for straight edge couples (i.e., a large radius of the fitted circle), all radii above a certain threshold (10^4) were limited to that threshold.

$$c_{ij} = \sqrt{\frac{(\beta_{ij}^C (\lambda 1_{ij}^C - \lambda 2_{ij}^T)^2 + \beta_{ij}^T (\lambda 1_{ij}^T - \lambda 2_{ij}^C)^2 + \beta_{ij}^C \beta_{ij}^T (\gamma 1_{ij} - \gamma 2_{ij})^2 + \omega_{ij}^2)}{\beta_{ij}^C + \beta_{ij}^T + \beta_{ij}^C \beta_{ij}^T + 1}} \quad (3)$$

The spatial deviation determines how well the couple fits a circle, but does not take into account the “direction” of the edges. From Fig. 3(a), one can observe that the normal edge couple lies on the fitted circle in the following clockwise order: *tail 1 – head 1 – head 2 – tail 2*, while an erroneous edge couple lies in the order *tail 1 – head 1 – tail 2 – head 2* [see Fig. 3(b)]. In both cases, the spatial deviation is small. To penalize these incorrect continuations, we introduce the *angular consistency*

$$\varphi_{ij} = \frac{|\Delta^* \alpha_i - \Delta^* \alpha_j|}{|\Delta^* \alpha_i| + |\Delta^* \alpha_j| + \varepsilon^\varphi} \quad (7)$$

where $\Delta^* \alpha = \alpha^t \setminus \alpha^h$, with α^h and α^t being the angles measured from the center of the fitted circle to the *head* and *tail* of the edge, respectively, and $\varepsilon^\varphi = 10^{-6}$ a small value to avoid a potential division by zero. The operator “ \setminus ” defines the smallest angle (in absolute value) between two angles α_1 and α_2 (both between 0 and 2π) as follows:

$$\alpha_1 \setminus \alpha_2 = \begin{cases} \Delta \alpha_{1,2} = \alpha_1 - \alpha_2, & \text{if } |\Delta \alpha_{1,2}| \leq \pi \\ -\text{sign}(\Delta \alpha_{1,2}) \times (2\pi - |\Delta \alpha_{1,2}|), & \text{otherwise} \end{cases} \quad (8)$$

Finally, we also want to penalize very wide angles between the heads of the two edges in a couple, since such configurations are very unlikely. For example, the edge couple in Fig. 3(d) is much less common than edge couple in Fig. 3(c). This is measured by the *aperture quality*

$$\theta_{ij} = \begin{cases} \sqrt{1 - \frac{|\tilde{\Delta} \alpha_{ij}^h|}{2\pi}}, & \text{if } |\tilde{\Delta} \alpha_{ij}^h| < |\tilde{\Delta} \alpha_{ij}^t| \\ \sqrt{\frac{|\tilde{\Delta} \alpha_{ij}^h|}{2\pi}}, & \text{otherwise} \end{cases} \quad (9)$$

where $\tilde{\Delta} \alpha_{ij}^h = \alpha_i^h \setminus \alpha_j^h$, $\tilde{\Delta} \alpha_{ij}^t = \alpha_i^t \setminus \alpha_j^t$, and $\theta_{ij} \in [0, 1]$.

The square roots in (9) are meant to approximately calibrate the values returned by θ_{ij} . Note that the *aperture quality* measure is (more or less) equivalent to the *proximity* property stated by Wertheimer [36]. Moreover, the *aperture quality* is scale independent, which is a desirable property of any extracted feature.

B. Global Feature and Prediction of the Final Configuration

Besides looking at how well edges match within a couple, we also take into account the global configuration that is created, in order to exclude false edge couples. Here, we measure the edge order, or *sequentiality* of the edges, which validates the configuration. Suppose we are dealing with an artefact that splits a number of horizontal objects in two, i.e., they appear (once) on both the left and right sides of the artefact. As an example see Fig. 4. If we inspect the artefact border in clockwise order, the object edges on one side [such as edges $\{e_1, \dots, e_6\}$ in Fig. 4(b)] appear in exactly opposite order compared to the ones on the other side (edges $\{e_7, \dots, e_{12}\}$). This is a very useful property of the edges around artefacts, since it is extremely robust against noisy data. For example, Fig. 4(a) shows how edges can be connected in a wrong way when only the local features are accounted for. Here, the presence of noise resulted in slightly tilted edges (which affected the circle fitting cost), as well as distorted grey levels and gradient magnitudes (which affected the other costs). When the sequentiality

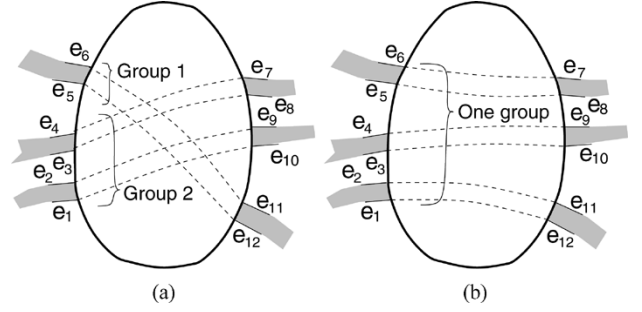


Fig. 4. Contribution of the *sequentiality* parameter. (a) Configuration penalized by the sequentiality parameter. (b) Configuration given preference by the sequentiality parameter.




Fig. 5. Reconstruction examples for fading spare edges.

of the edge couples is also taken into account, the right configuration can be better predicted [Fig. 4(b)]. Edge displacement and changed grey levels do not change edge order so they do not influence the sequentiality feature. The only way in which noise can affect this feature is by hampering the edge detection process, introducing false edges, or missing existing ones. However, the other edges still lie in the same consecutive order, which contributes to the stability of the cost. Most probably, an erroneously introduced edge, or the remaining pair of a missed edge will be treated as spare edges, and, thus, the impact on the sequentiality cost is reduced (since this cost is computed over pairs of edges only).

The sequentiality represents a natural property of most object edges. If edges are not sequential, then they should change their order in the image very often, i.e., they should cross each other, as in interwoven patterns. While interwoven patterns are not unusual, they are certainly not encountered very often.

It is worth pointing out that the sequentiality parameter does not forbid a configuration containing crossing groups—rather, it penalizes it. If the evidence coming from the local features strongly indicates a crossing, separate groups are formed accordingly [resulting in a configuration such as the one in Fig. 4(a)].

Sequential configurations usually have smooth edge couples. This does not mean that the features based on sequentiality and circle fitting are the same: smooth edge couples are not necessarily sequential. Besides, in practice, the detected edges are sometimes displaced or tilted, which affects the smoothness feature. The sequentiality comes to correct for such cases.

Three problems arise when determining the sequentiality of a configuration. First, we must find a way to express it as a number. Second, despite the fact that it is used to calculate the

- 1) Calculate the feature distance d_{kl} between all possible edge pairs (e_k, e_l) with Eq.3 by setting all β flags to 1 (at this point we do not know which edges are spare ones). d_{kk} is set to 1 for all k in order to avoid edges to be connected to themselves.
- 2) Calculate the potential of two edges to make a couple ($K^E = \sqrt{0.1}$ is introduced here in order to avoid that two spare edges are fitted to each other):

$$\pi_{kl} = \pi_{lk} = \begin{cases} TRUE & , \text{ if } d_{kl} < \min(2 \times \min_k d_{kl}, 2 \times \min_l d_{kl}, K^E) \\ FALSE & , \text{ otherwise} \end{cases}$$

- 3) FOR every possible (not necessarily maximal) set of couples

$$Q_i = \{q_j = (e_k, e_l) | \pi_{kl} = TRUE\}$$

- a) All edges which do not belong to a couple from Q_i are considered spare edges and are taken into account only in step [b.vi].
- b) FOR every couple q_j from Q_i
 - i) Initialize group G_1 with couple $q_j = (e_k, e_l)$ and remove the couple from Q_i .
 - ii) Find next edges: e_k^{next} clockwise from e_k and e_l^{next} anti-clockwise from e_l , provided that: $\exists u, v$ such that $e_k^{next} \in q_u, q_u \in Q_i$, and $e_l^{next} \in q_v, q_v \in Q_i$.
 - iii) IF $u = v$ AND q_u does NOT intersect any couple from the current group THEN add couple q_u to the current group, remove it from Q_i and GO TO step ii.
 - iv) Repeat steps ii-iii, this time searching on the other side of the first couple from the current group: anti-clockwise from e_k and clockwise from e_l .
 - v) If $Q_i \neq \emptyset$ and no more couples can be added to the current group, initialize another group with a couple from Q_i and repeat steps i-iv for the new group.
 - vi) Calculate the cost c^{cfg} of the current configuration of groups, Z_h according to Eq.1 (the β flags are now enabled according to the existing spare edges).

4. Choose the configuration with the minimal cost: $Z = Z_m, m = \arg \min_h c^{cfg}(Z_h)$

Fig. 6. Pseudocode for the grouping procedure.

configuration cost, we can measure it only *after* the configuration of edge couples has been formed, based on some cost that does not depend on sequentiality. And third, the sequentiality does not represent a measurement of each edge couple alone. Rather, it is a measurement of the complete configuration, which is an ensemble of edge couples. The latter problem gave rise actually to the formula in (1).

For the moment, let us assume that the configuration of edge couples has already been found. The groups of edge couples in the current configuration are denoted by $G_i, i = 1, \dots, N^G$. Equation (10) then defines the cost related to the global property of sequentiality

$$c^{gfb}(Z) = \begin{cases} 1 - \frac{\sum_{i=1}^{N^G} (\|G_i\| - 1)}{\lfloor \frac{N^E}{2} \rfloor - 1}, & \text{if } \lfloor \frac{N^E}{2} \rfloor \geq 2 \text{ and } \sum_{i=1}^{N^G} \|G_i\| \geq 1 \\ 0, & \text{if } \lfloor \frac{N^E}{2} \rfloor = 1 \text{ and } \sum_{i=1}^{N^G} \|G_i\| = 1 \\ 1, & \text{otherwise} \end{cases} \quad (10)$$

Operator $\lfloor \cdot \rfloor$ rounds to the nearest smaller integer. Thus, the term $\lfloor N^E/2 \rfloor$ represents the maximum number of edge couples

that can be achieved out of the N^E edges. To exemplify this measure, the six couples in Fig. 4(a), have a sequentiality cost of $1 - ((4-1) + (2-1))/(6-1) = 0.2$, while the six couples in Fig. 4(b), have a sequentiality cost of $1 - (6-1)/(6-1) = 0$. It is clear now that the global feature favors fewer but larger groups of edge couples [e.g., Fig. 4(b)], and penalizes more, but smaller groups [e.g., Fig. 4(a)]. As a result, it imposes a (desired) natural constraint on the configurations (in most of the images, edges do not cross each other locally).

The main steps for building up the final configuration are summarized into pseudocode in Fig. 6.

C. Spare Edges Reconstruction

Before one can use the selected configuration Z to restore the artefact, the spare edges must be integrated with the edge couples. Ideally, we should be able to fit circles to spare edges, similarly to what we did with the edge couples, and then calculate where they intersect with the couples. Unfortunately, experiments have shown that fitting circles to spare edges is unreliable and frequently gives unnatural results. This happens mostly because 1) the edges are usually small (remember that when fitting a circle to a couple, the two edges are relatively far apart, making




Fig. 7. (a) Pixel similarity along edges: the value of P is closer to the values of A and B , rather than C or D . (b)-(d) Inpainting of side strips with continuous contour. The side strips are bounded by: (b) an edge couple, (c) an edge couple and a spare edge, and (d) two spare edges.

the fit reliable), and 2) they could be quite noisy (spatially). This motivated us to approximate the spare edges with straight lines (a choice which was validated by experimental results).

To reconstruct the structure of the spare edge inside a strip of the artefact, we iteratively pick the spare edge with the biggest difference between the luminances on its two sides $|\lambda_i^C - \lambda_i^T|$, approximate it with a straight line, and recover it. This is repeated until all spare edges have been traced.

When recovering a spare edge, two situations may occur. 1) The recovered spare edge does not intersect with any other reconstructed edge within the artefact area [e.g., edges $E_1 - E_2$ and $E_3 - E_4$ in Fig. 5(a) and (b)]. 2) The recovered spare edge intersects with another edge that was already recovered inside the artefact [e.g., edge $E_2 - E_B$ in Fig. 7(c)]. In situation 1), we are dealing with a *fading* edge, while, in situation 2), the edge is part of a *T junction*. In Fig. 5(b), the reconstructed spare edge increments the number of *middle strips*² existing inside the artefact. In all other cases, it only adds a new *side strip*,³ even in Fig. 5(a), where the strip $A_1 - A_2 - E_2 - E_1 - A_1$ will be considered a side strip with *fragmented* contour. Fragmented contours occur in places where a reconstructed fading edge intersects the same contour a second time, cutting out a side strip and fragmenting the old contour [e.g., in Fig. 5(a), the *continuous* contours $A_1 - A_2$ and $A_1 - B_1$ become $A_1 - E_1$; $E_2 - A_2$ and $A_1 - E_3$; $E_4 - B_1$, respectively].

IV. EDGE-BASED INPAINTING

If the structure reconstruction step builds the “skeleton” of the missing areas, then we could say that the inpainting step adds the “flesh.” During the inpainting process, the middle strips and the side strips will undergo different types of interpolation. In all cases, however, we rely on the finding that the image structure around an edge is usually “parallel” to that edge.

In the case where we have more edge groups [i.e., crossing edge couples as in Fig. 4(a)], we have to assume that one group lies in the front of the others. Since the information extracted so far provides no guidelines as to which one is in the front

²A middle strip is an area that spans from one side to the other of the artefact and is usually delimited by two consecutive edge couples from the same group [e.g., strips $A_1 - E_1 - E_2 - A_2$ in Fig. 5(b), or $A_1 - B_1 - B_2 - A_2$ and $A_1 - E_2 - E_B - B_2 - A_2$ in Fig. 7(b) and (c), respectively].

³A side strip is an area delimited usually by a single edge couple, or by one or two spare edges [e.g., $E_1 - E_2 - E_1$, $E_3 - E_4 - E_3$ or $B_1 - B_2 - B_1$ in Fig. 5(a), $E_2 - E_B - B_1$ in Fig. 7(c), or $A_1 - E_{A_1} - E_1$ in Fig. 7(d)].

and which one in the background, the choice is made arbitrarily. Only groups consisting of a single edge couple (e.g., a horizon line) are “pushed” to the background, since their reconstruction in the foreground may cover entirely all other groups.

The following subsections describe our interpolation method, starting with the simplest case.

A. Inpainting of a Side Strip With Continuous Contour, Bounded Only by an Edge Couple

This is the simplest case of inpainting. We have a continuous contour and we know that there is an edge at each of the two ends of the contour [e.g., contour $A_1 - A_2$ in Fig. 7(b)]. When the two edges form an edge couple, then a restoration “parallel” to the edges is (broadly speaking) equivalent to drawing circle arcs on both sides of the couple. These arcs are concentric with the couple’s fitted circle, and span from one side of the artefact to the other [e.g., the $P_1 - P - P_2$ arc in Fig. 7(b)]. A pixel along such an arc (e.g., P) is interpolated from the ending pixels of the arc (P_1 and P_2 , in our case), which lie on the artefact border.

To understand the reason why we restore in this way, consider the example in Fig. 7(a). The missing area in region P is likely to be more similar to areas A and B , rather than C or D , although the last two are closer spatially. In fact, D and P are probably very different from each other, since they lie across edge couple $E_1 - E_2$, which means that they belong to two different objects.

Let us denote the circle fitted to edge e_k by c_k , $k \in \{1, 2\}$ [in Fig. 7(b), $c_1 = c_2$ because e_1 and e_2 belong to the same couple]. The circle that passes through P , and is “parallel” to (i.e., concentric with) c_k is denoted by c_k^P . It intersects with the artefact border at two points, $P_1^k = P_1$ and $P_2^k = P_2$. These two pixels are called the *source* pixels from which the intensity λ^P of pixel P is calculated as follows:

$$\lambda^P = \frac{w_1 \lambda^{P_1} + w_2 \lambda^{P_2}}{w_1 + w_2} \quad (11)$$

where w_i , $i \in \{1, 2\}$, is inversely proportional to the distance from P to P_i : $w_i = 1/d(\mathbf{x}^{P_i}, \mathbf{x}^P)$. λ^{P_i} represents the intensity of pixel P_i .

B. Inpainting of a Side Strip With Continuous Contour Bounded by an Edge Couple and a Spare Edge

This is a slightly more complicated case. As an example see Fig. 7(c). Now the side strip is not bounded by a single edge couple, but an edge that belongs to a couple, and one spare edge.

To ease the discussion, the straight line fitted to the spare edge is also denoted as a circle. Like in the previous situation, we strive toward restoring the artefact as much as possible parallel to the edges. This means that, if we are close to one of the edges, for example, then we should interpolate from points on the artefact border close to the circle fitted to that edge. Again this is accomplished by defining two source pixels on the artefact border (one for each bounding edge) from which pixel P is interpolated, but now one of them is related to the spare edge and the other one to the edge couple.

In Fig. 7(c), the first source pixel P_1 is created from the edge e_1 by intersecting c_1^P with the artefact border. The other source pixel is then found by intersecting c_2^P (based on the spare edge e_2) with the artefact border. The intensity of pixel P is now calculated from the source pixels P_1 and P_2 according to

$$\lambda^P = \frac{w_1^a \lambda^{P_1} + w_2^a \lambda^{P_2}}{w_1^a + w_2^a} \quad (12)$$

where w_i^a , $i \in \{1, 2\}$, is inversely proportional to *both* the distance from P to P_i as well as the distance from P to c_i

$$w_i^a = \frac{1}{\max(d(\mathbf{x}^P, \mathbf{x}^{P_i}), \varepsilon^D) \times \max(d^c(\mathbf{x}^P, c_i), \varepsilon^D)} \quad (13)$$

where $\varepsilon^D = 0.01$ is used for protecting against potential divisions by zero, as well as for avoiding unusually large weights due to the proximity of pixel P to either of the circles c_1 and c_2 . $d^c(\mathbf{x}, c)$ represents the euclidean distance between a pixel with coordinates \mathbf{x} and the closest point on circle c . The weights w_i^a place more emphasis on source pixels close to pixel P . Also, in the immediate neighborhood of a reconstructed edge, the source pixel that is close to that edge will dominate, thereby preserving the edge sharpness.

Notice that a side strip with continuous contour can also be formed by two spare edges, for example $E_1 - E_{A_1} - E_{A_2} - E_2$ in Fig. 7(d). Here, P_1 and P_2 are created in a similar way, by intersecting circles c_1^P and c_2^P (corresponding to spare edges e_1 and e_2 , respectively) with the artefact border. The intensity of point P is then again estimated with the formulas in (12) and (13).

C. Inpainting of a Middle Strip With Continuous Contours

The next case of inpainting is a middle strip. In its simplest form, it is only bounded by two edge couples from the same group (see Fig. 8). Again, the interpolation is driven by the structure defined by the bounding edge couples.

Similarly to the side strip case, source pixels from the artefact border are calculated, upon which the interpolation is based. Since we now have two bounding edge couples, two sets of source pixels are created, $\{P_1^k, P_2^k\}$ and $\{P_1^l, P_2^l\}$, each based on one of the two edge couples, \mathbf{q}_k and \mathbf{q}_l , respectively. From the two source pixels that belong to the same part of the contour, $\{P_1^k, P_1^l\}$ and $\{P_2^k, P_2^l\}$, two virtual source pixels are created: P_1 and P_2 . The position of such a virtual source pixel (see also Fig. 8) is defined by

$$\mathbf{x}^{P_i} = \frac{w_1^a \mathbf{x}^{P_1^k} + w_2^a \mathbf{x}^{P_1^l}}{w_1^a + w_2^a}, \quad i = 1, 2 \quad (14)$$

where w_i^a is defined as in (13).




Fig. 8. Inpainting of a middle strip with continuous contours.

The intensity of these virtual source pixels P_i are defined as

$$\lambda^{P_i} = \frac{w_1^a \lambda^{P_1^k} + w_2^a \lambda^{P_1^l}}{w_1^a + w_2^a}. \quad (15)$$

Based on the coordinates and intensities of the virtual source pixels P_1 and P_2 , the intensity of pixel P can now be determined as follows:

$$\lambda^P = \frac{w_1^b \lambda^{P_1} + w_2^b \lambda^{P_2}}{w_1^b + w_2^b} \quad (16)$$

where w_i^b is inversely proportional to the distance from point P to the virtual source pixel P_i .

D. Other Cases

When a side strip or a middle strip has fragmented contours [e.g., Fig. 5(a)], they are interpolated similarly to the strips with continuous contours. However, in this case, a virtual source pixel is calculated for each fragment independently [e.g., one for fragment $A_1 - E_3$ and one for fragment $E_4 - B_1$ in Fig. 5(a)] and then the virtual source pixel of the entire fragmented contour ($A_1 - E_3; E_4 - B_1$) is computed as a weighted average of its fragments' source pixels. The rest of the procedure is similar to the previous subsections.

If no edges are detected around the artefact, then the artefact lies probably in a smooth area. In such a case, the intensity of an artefact pixel P is simply the weighted average of the pixels on the artefact border. The weights are inversely proportional to the distance from P to the border pixels.

V. RESULTS

A. Qualitative Evaluation

In this subsection the performance of our proposed algorithm is demonstrated by some visual examples. Fig. 9(a)–(c) shows an artificially degraded version of the “Lena” image, the restored version and a zoom-in on one of the artefact areas in the restored image (for every artefact, the restored structure consisted of a single group of coupled edges). Fig. 9(d)–(f) shows an example of interpolated spare edge (a T junction). Visual inspection of these results shows a good restoration quality. Both sharp and smooth edges are well recovered.

One of the strengths of our restoration scheme comes from its capability of finding and interpolating crossing structures. Fig. 10(e) shows such an example. Here, a group of two edge couples (the margins of the dark grey bar) is crossed by another group of two edge couples (the margins of the light grey bar).




Fig. 9. Restoration results. (a) “Lena” image, degraded with artificial artefacts. (b) Restored “Lena.” (c) Zoom-in on the restored image. (d) “Lena” image, artificially degraded over a T junction. (e) Zoom-in on the original image. (f) Zoom-in on the restored image.




Fig. 10. Comparison on an artificial example with crossing structures. (a) Original image. (b) Degraded image. (c) Restoration by the algorithm of Atzori and De Natale [22]. (d) Structure recovered by the algorithm of Atzori and De Natale. (e) Restoration by our proposed algorithm. (f) Structure recovered by our algorithm.

The restoration shows that the proposed algorithm is capable of reconstructing the correct configuration [Fig. 10(f)].

Obviously, our algorithm works well for objects which fit our assumptions. When edges are neither straight, nor circular (e.g., wiggly edges), the structure reconstruction will not be able to reproduce the initial image content. Also, when the structure becomes complex (e.g., in textured areas), the structure reconstruction step will fail, unless there is a dominant structure orientation (e.g., an image of straws). In these complex cases, the abundance of edges will make the algorithm more prone to errors than in usual cases. Similarly, if many of the edges detected around an artefact are spare edges, the structure reconstruction becomes a very difficult task. In such a case, the luminosity-related costs of most edge couples are cancelled by spare edges, so the final costs may become dependent on circle fitting and sequentiality costs only. Since less features are taken into account, the edge matching gets less reliable than in a normal case, so the probability of mismatches grows. Some edge couples may get treated as two spare edges, or they become coupled with wrong edges, while some spare edges may get erroneously assigned to couples. A thorough analysis of the reconstructed structure can

only be done if a large database with manually segmented images would exist.

B. Quantitative Evaluation

Besides using visual inspection, we have also assessed the performance of the algorithm in a *quantitative* manner. A set of experiments was performed on a set of seven 512×512 images (see the name list in the legends of Fig. 11). These images were chosen because they exhibit some local structure. We have conducted the following series of experiments for each image. Artefacts with random shapes and locations were generated, having sizes of 1, 2, 5, 10, 20, 50, 100, 200, 500, 1000, 2000, 5000, and 10000 pixels. For each size and each image, a single artefact was generated and restored in 100 consecutive experiments (each time with a different, random shape and location). For each restoration, the mean-square error (MSE) was measured between the original and the reconstructed image. The MSE plots are shown in Fig. 11(a), with artefact sizes on a logarithmic scale. For each size and each image, the median MSE for the 100 experiments was plotted (this was chosen in order to avoid the influence of a small percentage of outliers). The MSE values

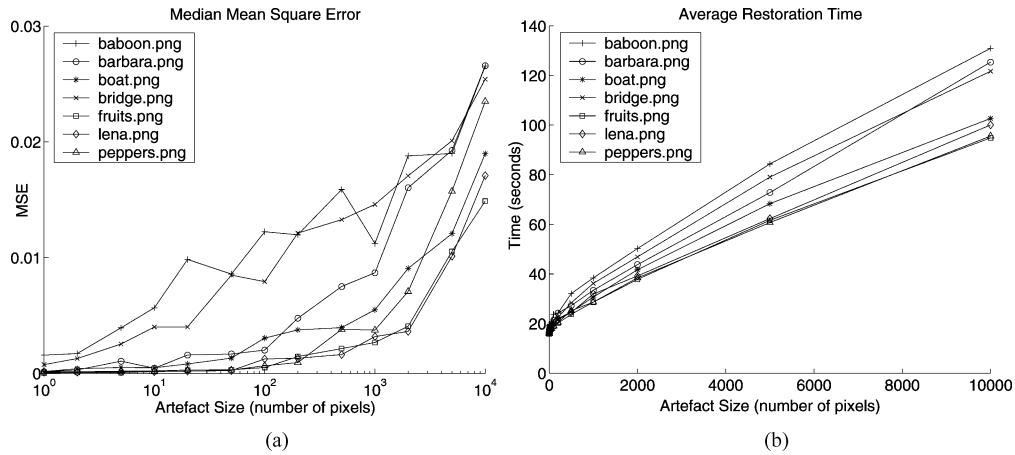


Fig. 11. Plots for the experiments done on the image test set. Each point represents the median result of 100 experiments done on the same image, with random artefacts of the same size. (a) Median MSE, calculated on the grey-value images (grey range: 0 . . . 1). (b) Average restoration time, under Matlab (interpreted code).




Fig. 12. Real case example of film restoration. (a), (e) Original frames, with artefacts of interest surrounded by a white box. (b), (f) Same frames, with main artefacts restored. (c), (g) Zoom-in on the areas of interest in the original frames. (d), (h) Zoom-in on the areas of interest in the restored frames.

stay within acceptable ranges, in general. A growing trend for bigger artefacts is present, as expected (the trend seems to accelerate at larger sizes because of the logarithmic scale used).

Additionally, the associated restoration times are displayed in Fig. 11(b). The artefact sizes are presented here on a linear scale, in order to show the almost linear dependency between the restoration time and the artefact size. The plot also shows a constant overhead, regardless of the artefact size. This overhead is related to the first part of the algorithm, in which object edges are detected, pixels on the artefact borders (together with the list of edges) are arranged in clockwise order, and edge features are computed.

From a perceptual point of view, our algorithm performed satisfactorily for MSE values up to about 0.005. Above this value, the quality of the restoration degraded in a more visible manner. This value is only a rough estimate and should not be taken as an absolute reference, since the MSE is not strictly correlated with the visual quality. Depending on the textural content

and the structural complexity of each image, the restoration errors may start becoming visible at smaller or larger MSE values and/or artefact sizes.

All experiments have been done with the same parameter setting. This showed that the parameter setting was not really sensitive to different images (i.e., different structure configurations), nor to different artefact shapes. Also, adding together costs with different variances did not seem to have a significantly negative impact on the quality of the restoration.

C. A Real Case Experiment

We also demonstrate the algorithm performance on a real case of degraded old film. Each row in Fig. 12 contains, from left to right, an original frame from a degraded film and the same frame in which the main artefacts were subject to restoration using our algorithm (we concentrate only on those artefacts which cover areas containing structure and moving objects). White boxes are

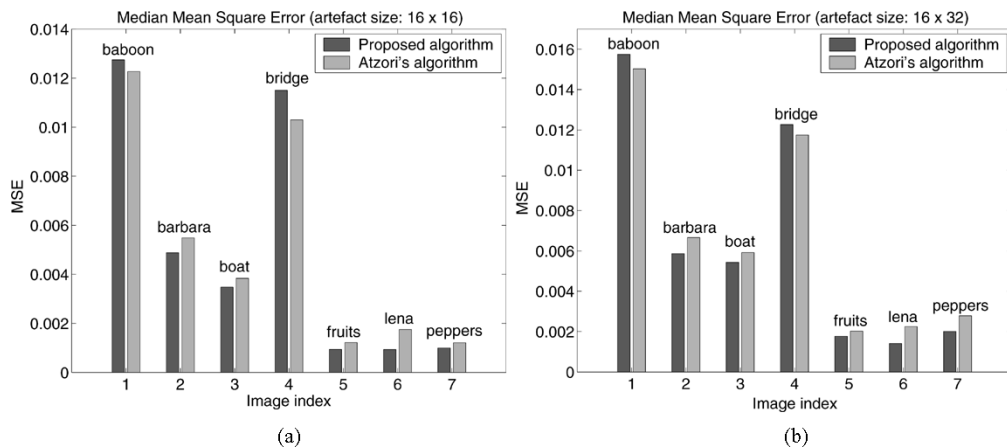


Fig. 13. (Dark bars) Comparison of the median MSE for the proposed algorithm and (light bars) the algorithm of Atzori and De Natale [22], for artefact sizes of (a) 16×16 pixels and (b) 16×32 pixels.




Fig. 14. Comparison with the algorithm of Atzori and De Natale [22]. (a) Original image (zoomed in). (b) Restoration by the algorithm of Atzori and De Natale. (c) Restoration by our proposed algorithm. (d) Full degraded image. (e) Structure recovered by the algorithm of Atzori and De Natale. (f) Structure recovered by our algorithm.

used in the original frames to mark artefacts of interest for our algorithm. These areas of interest are enlarged and displayed next to the full-size frames. The examples from Fig. 12 show that the algorithm performs equally well in real cases of degraded films.

D. Comparisons With Other Algorithms

We have performed a comparison of our algorithm and the sketch-based interpolation of Atzori and De Natale [22]. For each of the seven images from our test set, we have generated artefacts with different sizes and random locations (1000 iterations for each size). For reasons of compatibility with the code we received from Atzori and De Natale, the artefacts were chosen to be only rectangular, having 16×16 or 16×32 pixels. In order to allow a proper comparison of both algorithms, the code of Atzori has been modified such that the input edges for both algorithms are the same, namely, the edges extracted in the first step of our algorithm.

The median MSE of all experiments for each image was measured for both algorithms. The comparison graph is displayed in Fig. 13. For both artefact sizes, our algorithm scored better in five out of the seven images. The fact that both algorithms show larger MSE values for the highly textured images is an indication that the edge detection step had lower performances.

Visually, the restoration quality was not strikingly different for the two algorithms. This is not surprising, given the fact that the algorithms share some similarities. There are, however, more situations in which our algorithm outperforms the other one. Fig. 14 shows an example taken from our quantitative experiments. The circle fitting used in our algorithm enforced a more natural continuation of the edges, by connecting the upper-right edge with the lower-left one. Fig. 10 shows an artificial example of two bars crossing each other and an artefact covering their intersection. Our algorithm was able to detect and reconstruct the right image structure, while the algorithm of Atzori and De Natale failed. The fact that the input edge




Fig. 15. Comparison with the algorithm of Masnou [11]. (a) Original image. (b) Degraded image. (c) Restoration with the algorithm of Masnou. (d) Restoration with the proposed algorithm.

mask is not 100% the same comes from the fact that Atzori's algorithm considers the edges to start right from the artefact border, while we look at edges starting one pixel away from the artefact (thus, the two algorithms never have *exactly* the same edge input). At times, our algorithm benefited from the use of the sequentiality. This global feature has contributed decisively in cases where several edge connections were equally possible.

Due to the type of interpolation used in the last step, our algorithm may sometimes produce smoother than normal areas. However, the patch repetition used by Atzori and De Natale (in this case a mirroring across the artefact border) may introduce its own type of defects, for example when another object lies close to the artefact. In this case, the patches that would be pasted would repeat the object (or parts of it) inside the artefact, although that object does not even touch the artefact. Patch repetition may go wrong in other cases, too. If a strip that presents a constant change of intensity is "interrupted" by an artefact, the mirroring process reverses the gradient direction in the artefact area, introducing a sudden change of intensity in the middle of the artefact. From the bar graphs presented in Fig. 13, it becomes clear that our algorithm performs better for piecewise smooth images, or moderately textured ones. For highly textured images the algorithm of Atzori and De Natale performs better, mainly due to their interpolation scheme based on patch repetition.

We have also performed a comparison with the algorithm of Masnou [11], shown in Fig. 15. The comparison was performed on the example presented in [11]. Both algorithms give good results, as expected. While the two methods may perform similarly in many cases, for overlapped structures or T junctions (as defined in this paper) our algorithm would outperform Masnou's algorithm, which cannot handle them properly.

VI. CONCLUSIONS AND FUTURE WORK

We have presented here an algorithm for the spatial restoration of images. Our goal is to restore frames from image sequences that exhibit "difficult" object motion, making the temporal restoration ineffective. The algorithm uses edge information extracted from the area surrounding the artefact. Based on this information, the missing structure inside the artefact (in the form of object borders) is reconstructed and then the areas between the reconstructed borders are filled by a smooth continuation of the surrounding image data.

The algorithm performs best with piecewise smooth images. In these cases, the restoration results are very good (both visually and numerically), as long as there is enough information around the artefact that is strongly related to the missing data. For highly textured images, the restoration is less effective because the image does not possess a certain "structure"—rather, it is a pattern with some degree of irregularity. In these cases, a texture interpolation method should be employed. This, however, would guarantee only a visually pleasing result, and not a lower error.

One of the main advantages of our method, is that it makes use of both local and global features of the edges in the image. The use of a global feature that validates the edge couples with respect to each other within the recovered structure is a new approach to image restoration. To our knowledge, this is the first algorithm which explicitly takes into account such a global feature, i.e., the sequentiality. The way the interpolation is done, along the reconstructed structures, is also new.

The validity of our structural model was demonstrated by evaluating the algorithm both visually and numerically on various images and across several artefact sizes. Moreover, the same set of parameters was used for all experiments, which demonstrates the robustness of our approach.

By reconstructing overlapped structures, our algorithm actually steps into the three-dimensional area, bringing one structure in front and pushing the others in the background. At this stage, these abilities are rather rudimentary. A superior analysis may certainly be added in the future to ensure the correct depth order of the structures. In any case, since the edge groups that cross each other may give us some depth information, applying the proposed grouping scheme could reveal object occlusions in undegraded images, provided that one can achieve a satisfactory segmentation of the image.

One of the implicit assumptions made in this paper is that the artefact masks do not have holes. Indeed, the overwhelming majority of artefacts from old films does not have holes. When they do have them, a few solutions could be applied. The simplest one is to simply consider that the artefact does not have holes, restore it in the way presented in our paper, and then paste the original content of the artefact holes back into the image (thus, overriding a part of the restoration result). This, of course, neglects the structure that may be present inside the artefact holes,

which might help guiding the structure reconstruction process. In some cases, this information may even be used to decide which group gets painted in in the foreground. Another solution would be to split the artefact mask conveniently such that no resulting sub-mask contains any holes, and then proceed with the normal restoration algorithm.

There are several ways to improve the performances of our algorithm. First, it should be noted that the algorithm presented here uses only a one-pixel-wide layer of pixels around the artefact. By increasing the amount of pixels taken into account, we expect to get more reliable edge features and useful neighborhood information, which will improve the results in situations where the present algorithm has limited effectiveness.

Since the proposed algorithm works well with piecewise smooth images, rather than textured ones, whereas texture-based restoration shows opposite behavior in general, we expect that the combination of the two approaches would improve the spatial restoration of images [20], [21]. Clearly, one needs to be able to decide which scheme to use depending on the surrounding area of the artefact. A special analysis module should be employed for this purpose.

Finally, a more sophisticated approach can be developed for the treatment of the available temporal information, along with the spatial information. Useful information can be extracted about the type of motion that causes the failure of motion estimation [3], [6], and then used to further enhance the results of the current algorithm. These subjects will constitute the focus of our future research.

ACKNOWLEDGMENT

The authors would like to thank L. Atzori for making available to us the code of his restoration algorithm. The sequence used in the real case restoration (Fig. 12) is courtesy of RTP (Radiotevisão Portuguesa).

REFERENCES

- [1] A. C. Kokaram, *Motion Picture Restoration: Digital Algorithms for Artifact Suppression in Degraded Motion Picture Film and Video*. New York: Springer Verlag, 1998.
- [2] P. M. B. van Roosmalen, "Restoration of archived film and video," Ph.D. dissertation, ICT Group, EEMCS Faculty, Delft Univ. Technology, Delft, The Netherlands, 1999.
- [3] A. Rareş, M. J. T. Reinders, and J. Biemond, "Statistical analysis of pathological motion areas," presented at the IEE Seminar on Digital Restoration of Film and Video Archives, London, U.K., Jan. 16.
- [4] P. M. B. van Roosmalen, "High-level analysis of image sequences," Tech. Rep., INA (Institut National de l'Audiotvisuel), the EU Aurora Project, Paris, France, 1999.
- [5] [Online]. Available: http://brava.ina.fr/brava_public_impairments_list.en.html
- [6] A. Rareş, M. J. T. Reinders, and J. Biemond, "Complex event classification in degraded image sequences," presented at the IEEE Int. Conf. Image Processing, Thessaloniki, Greece, Oct. 2001.
- [7] —, "Image sequence restoration in the presence of pathological motion and severe artifacts," presented at the IEEE ICASSP, Orlando, FL, USA, May 2002.
- [8] [Online]. Available: <http://brava.ina.fr>
- [9] N. Nitzberg, D. Mumford, and T. Shiota, *Filtering, Segmentation and Depth*. New York: Springer-Verlag, 1993.
- [10] S. Masnou and J.-M. Morel, "Level-lines based disocclusion," presented at the IEEE Int. Conf. Image Processing, Chicago, IL, 1998.

- [11] S. Masnou, "Disocclusion: A variational approach using level lines," *IEEE Trans. Image Process.*, vol. 11, no. 2, pp. 68–76, Feb. 2002.
- [12] C. Ballester *et al.*, "Filling-in by joint interpolation of vector fields and gray levels," *IEEE Trans. Image Process.*, vol. 10, no. 8, pp. 1200–1211, Aug. 2001.
- [13] C. Ballester, M. Bertalmio, V. Caselles, G. Sapiro, and J. Verdera, "A variational model for filling-in gray and color images," presented at the ICCV, Vancouver, BC, Canada, Jul. 2001.
- [14] C. Ballester, V. Caselles, and J. Verdera, "A variational model for disocclusion," presented at the IEEE Int. Conf. Image Processing, 2003.
- [15] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester, "Image inpainting," presented at the SIGGRAPH, 2000.
- [16] M. Bertalmio, A. Bertozzi, and G. Sapiro, "Navier-stokes, fluid-dynamics and image and video inpainting," presented at the IEEE CVPR, 2001.
- [17] T. Chan and J. Shen, "Mathematical models for local nontexture inpainting," *SIAM J. Appl. Math.*, vol. 62, no. 3, pp. 1019–1043, 2001.
- [18] —, "Non-texture inpainting by curvature-driven diffusions (CDD)," *J. Vis. Commun. Image Represen.*, vol. 12, no. 4, pp. 436–449, 2001.
- [19] T. Chan, S. H. Kang, and J. Shen, "Euler's elastica and curvature based inpainting," *SIAM J. Appl. Math.*, vol. 63, no. 2, pp. 564–592, 2002.
- [20] M. Bertalmio, L. Vese, G. Sapiro, and S. Osher, "Simultaneous structure and texture image inpainting," *IEEE Trans. Image Process.*, vol. 12, no. 8, pp. 882–889, Aug. 2003.
- [21] S. Rane, M. Bertalmio, and G. Sapiro, "Structure and texture filling-in of missing image blocks for wireless transmission and compression applications," *IEEE Trans. Image Process.*, vol. 12, no. 3, pp. 296–303, Mar. 2002.
- [22] L. Atzori and F. G. B. De Natale, "Error concealment in video transmission over packet networks by a sketch-based approach," *Signal Process.: Image Commun.*, vol. 15, no. 1-2, Sep. 1999.
- [23] —, "Reconstruction of missing or occluded contour segments using Bezier interpolations," *Signal Process.*, vol. 80, no. 8, pp. 1691–1694, 2000.
- [24] L. Atzori, F. G. B. De Natale, and C. Perra, "A spatio-temporal concealment technique using boundary matching algorithm and mesh-based warping (BMA-MBW)," *IEEE Trans. Multimedia*, vol. 3, no. 3, pp. 326–338, Sep. 2001.
- [25] H. Knutsson and C.-F. Westin, "Normalized and differential convolution: Methods for interpolation and filtering of incomplete and uncertain data," presented at the IEEE CVPR, New York, 1993.
- [26] L. Khriji, M. Gabbouj, G. Ramponi, and E. D. Ferrandiere, "Old movie restoration using rational spatial interpolators," presented at the 6th IEEE Int. Conf. Electronics, Circuits, Systems, Sep. 1999.
- [27] A. A. Efros and T. K. Leung, "Texture synthesis by nonparametric sampling," presented at the ICCV, 1999.
- [28] R. Bornard, E. Lecan, L. Laborelli, and J.-H. Chenot, "Missing data correction in still images and image sequences," presented at the ACM Multimedia, Juan Les Pins, France, Dec. 2002.
- [29] A. Criminisi, P. Pérez, and K. Toyama, "Object removal by exemplar-based inpainting," presented at the IEEE CVPR, 2003.
- [30] A. C. Kokaram, "Parametric texture synthesis using stochastic sampling," in IEEE Int. Conf. Image Processing, New York, Sep. 2002.
- [31] —, "Practical MCMC for missing data treatment in degraded video," presented at the ECCV Workshop on Statistical Methods for Time Varying Image Sequences, Copenhagen, Denmark, 2002.
- [32] A. C. Kokaram and S. Godsill, "MCMC for joint noise reduction and missing data treatment in degraded video," *IEEE Trans. Signal Process.*, vol. 50, no. 2, pp. 189–205, Feb. 2002.
- [33] J. Jia and C.-K. Tang, "Image repairing: Robust image synthesis by adaptive ND tensor voting," presented at the IEEE CVPR, 2003.
- [34] S. T. Acton, D. P. Mukherjee, J. P. Havlicek, and A. C. Bovik, "Oriented texture completion by AM-FM reaction-diffusion," *IEEE Trans. Image Process.*, vol. 10, no. 6, pp. 885–896, Jun. 2001.
- [35] A. N. Hirani and T. Totsuka, "Combining frequency and spatial domain information for fast interactive image noise removal," in *Proc. ACM SIGGRAPH*, 1996, pp. 269–276.
- [36] M. Wertheimer, "Laws of organization in perceptual forms," in *A Source Book on Gestalt Psychology*, W. Ellis, Ed. London, U.K.: Routledge & Kegan Paul, 1938, pp. 71–88.



Andrei Rareș (M'97) was born in Bucharest, Romania. He received the B.Sc. and M.Sc. degrees in computer science from Politehnica University of Bucharest in 1996 and 1997, respectively, and the Ph.D. degree in electrical engineering from the Delft University of Technology, Delft (TU Delft), The Netherlands, in 2004.

From 1996 to 1999, he was a Teaching Assistant with the Faculty of Automatic Control and Computers, Politehnica University of Bucharest. From 1999 to 2003, he was a Ph.D. student in the

Information and Communication Theory Group, Mediamatics Department of the Faculty of Electrical Engineering, Mathematics and Computer Science, TU Delft. Since 2003, he has been a Researcher in the Laboratory for Clinical and Experimental Image Processing, Leiden University Medical Center, Leiden, The Netherlands.

His research interests are in image and video processing, including restoration, object tracking, motion estimation, data compression, and medical image analysis.



Marcel J. T. Reinders received the M.Sc. degree in applied physics and the Ph.D. degree in electrical engineering from the Delft University of Technology, Delft (TU Delft), The Netherlands, in 1990 and 1995, respectively.

Currently, he is a Professor in the Information and Communication Theory Group, Mediamatics Department of the Faculty of Electrical Engineering, Mathematics and Computer Science, TU Delft. He is active in the field of machine learning. Besides studying fundamental issues, he applies machine

learning techniques to the areas of bioinformatics, computer vision, and context-aware recommender systems. His special interest goes toward understanding complex systems (such as biological systems) that are severely undersampled.



Jan Biemond (M'80–SM'87–F'92) was born in De Kaag, The Netherlands. He received the M.S. and Ph.D degrees in electrical engineering from the Delft University of Technology, Delft (TU Delft), The Netherlands, in 1973 and 1982, respectively.

Currently, he is a Professor in the Information and Communication Theory Group and Head of the Mediamatics Department of the Faculty of Electrical Engineering, Mathematics and Computer Science, TU Delft. In 1983, he was a Visiting Professor at Rensselaer Polytechnic Institute, Troy, NY, and at

the Georgia Institute of Technology, Atlanta. His research interests include image and video processing (restoration, compression, content-based retrieval, and motion estimation) with applications in digital TV, 3-D TV, HDTV, multimedia, digital libraries, scan-rate conversion, and computer vision. He has published extensively in these fields and has supervised more than 25 Ph.D. theses covering these fields.

Currently, he is Chairman of the IEEE Benelux Section, a Member of the Educational Activities Subcommittee of Region 8, and a Member of the Nominations and Appointments Committee of the IEEE Signal Processing Society. He served this Society as a Distinguished Lecturer from 1993 to 1994. He is a former member of the Administrative Committee of the European Association for Signal Processing (EURASIP), the IEEE Technical Committee on Image and Multidimensional Signal Processing, and the Board of Governors of the IEEE Signal Processing Society. He served as the General Chairman of the Fifth IEEE-SP/EURASIP Workshop on Multidimensional Signal Processing, Noordwijkerhout, The Netherlands, in September 1987, as the General Chairman of the 1997 Visual Communication and Image Processing Conference (VCIP'97), San Jose, CA, and as the Chairman of the 21st Symposium on Information Theory in the Benelux, Wassenaar, The Netherlands, May 25–26, 2000. Currently, he is General Co-Chair of the IEEE International Conference on Multimedia and Expo (ICME'05), to be held July 2005 in Amsterdam, The Netherlands. He was the recipient of the Dutch Telecom Award "Vederprijs" in 1986 for his contributions in the area of digital image processing, in particular, in image restoration and subband coding.