A Perceptually Relevant Approach to Ringing Region Detection

Hantao Liu, *Student Member, IEEE*, Nick Klomp, and Ingrid Heynderickx

Abstract—An efficient approach toward a no-reference ringing metric intrinsically exists of two steps: first detecting regions in an image where ringing might occur, and second quantifying the ringing annoyance in these regions. This paper presents a novel approach toward the first step: the automatic detection of *regions* visually impaired by ringing artifacts in compressed images. It is a no-reference approach, taking into account the specific physical structure of ringing artifacts combined with properties of the human visual system (HVS). To maintain low complexity for real-time applications, the proposed approach adopts a perceptually relevant edge detector to capture regions in the image susceptible to ringing, and a simple yet efficient model of visual masking to determine ringing visibility. The approach is validated with the results of a psychovisual experiment, and its performance is compared to existing alternatives in literature for ringing region detection. Experimental results show that our method is promising in terms of both reliability and computational efficiency.

Index Terms—Luminance masking, perceptual edge, ringing artifact, texture masking.

I. INTRODUCTION

IN CURRENT visual communication systems, the most essential task is to fit a large amount of visual information into the narrow bandwidth of transmission channels or into a limited storage space, while maintaining the best possible perceived quality for the viewer [1]. A variety of compression algorithms, for example, such as JPEG and MPEG/H.26x, have been widely adopted in image and video coding trying to achieve high compression efficiency at high quality [2], [3]. These lossy compression techniques, however, inevitably result in various coding artifacts, which by now are known and classified as blockiness, ringing, blur, etc. [4]. The occurrence of the compression induced artifacts depends on the data source, target bitrate, and underlying compression scheme, and their visibility can range from imperceptible to very annoying, thus affecting perceived quality [5]–[7]. During the last decades a lot of research effort is devoted to reduce coding artifacts, so to improve the overall perceived quality of artifact impaired image material [8]–[10]. In the video chain of a current television set, e.g., various video enhancement algorithms, such as deblocking, deringing, and deblur, are typically employed to reduce compression artifacts prior to display. In such a scenario, objective metrics, which determine the quality degradation caused by each individual artifact, and adapt the processing chain for artifact reduction accordingly, are highly needed. In addition, the receiving end of a digital video chain usually has no access to the original image, and in most cases there is even only limited access to the encoding parameters of the bit-stream. Hence, objective metrics used in these types of applications are constrained to a no-reference approach, which means that the impairment assessment relies on the compressed image only.

In the last decades, a considerable amount of research has been devoted to the development of a blockiness metric (e.g., see [11] and [12]), which has been already implemented for the optimization of image quality (e.g., see [13]–[15]). Another common distortion type, namely ringing [4], intrinsically results from loss in the high-frequency component of the video signal due to coarse quantization. In the spatial domain, ringing, which is fundamentally associated with Gibb's phenomenon, manifests itself in the form of ripples or oscillations around high-contrast edges. The occurrence of ringing artifacts spreads out to a finite extent surrounding edges, depending on the underlying properties of the compression scheme. For example, in block-based DCT coding ringing appears as a ripple outward from the edge up to the encompassing block's boundary [4]. As an example, Fig. 1 illustrates ringing artifacts induced by JPEG compression.

Research on the design of a blockiness metric has shown that an efficient no-reference approach intrinsically exists of two steps: 1) the detection of regions in an image where blockiness might occur, and 2) the determination of the blocking annoyance in these regions. We use a similar two-step approach for the design of a no-reference ringing metric. This paper only discusses the first step: the detection of regions in the image, in which visible ringing occurs. A successive paper that discusses the quantification of the perceived annoyance of ringing in these regions is published in [16] and [17].

Unlike blocking, whose spatial location is very regular and thus easily predictable, the location of ringing is edge dependent, and as such also image content dependent. This makes the task of detecting ringing regions much more difficult, especially in a no-reference application. In general, ringing can be considered as a form of signal dependent noise, which only occurs near sharp transitions in image intensity when not visually masked by local image characteristics. As such, the occurrence of ringing can be directly associated with strong edges in an image. Additionally, the visibility of ringing is reduced in the case of very low and very high background intensity (i.e., luminance
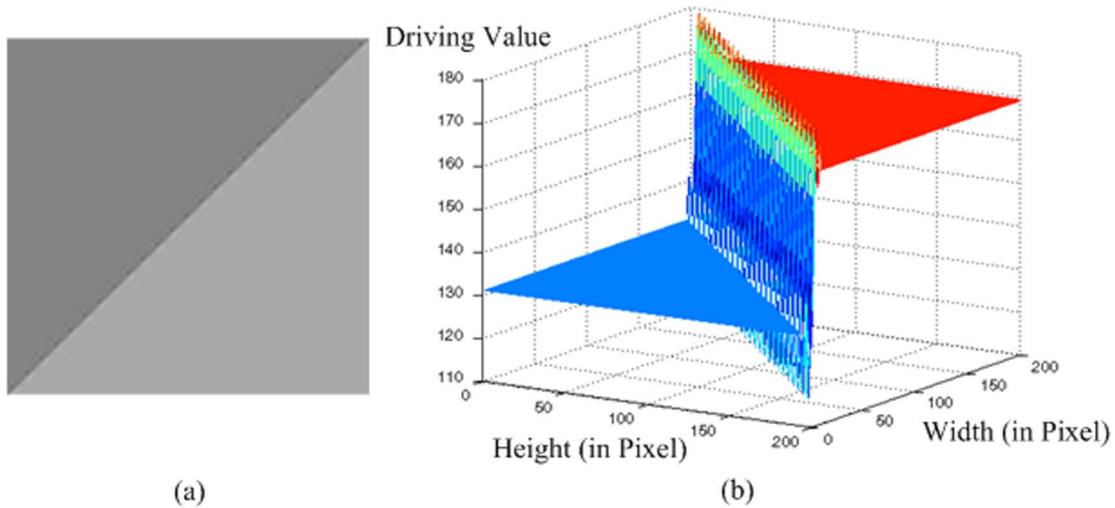
Fig. 1. Illustration of ringing artifacts in an image patch compressed with JPEG (MATLAB's *imwrite* function with $Q = 30$). (a) 2-D image and (b) its spatial intensity distribution (in 8-bits driving values). Ringing can be perceived as intensity fluctuations near the edges, while the image content there should be uniform.
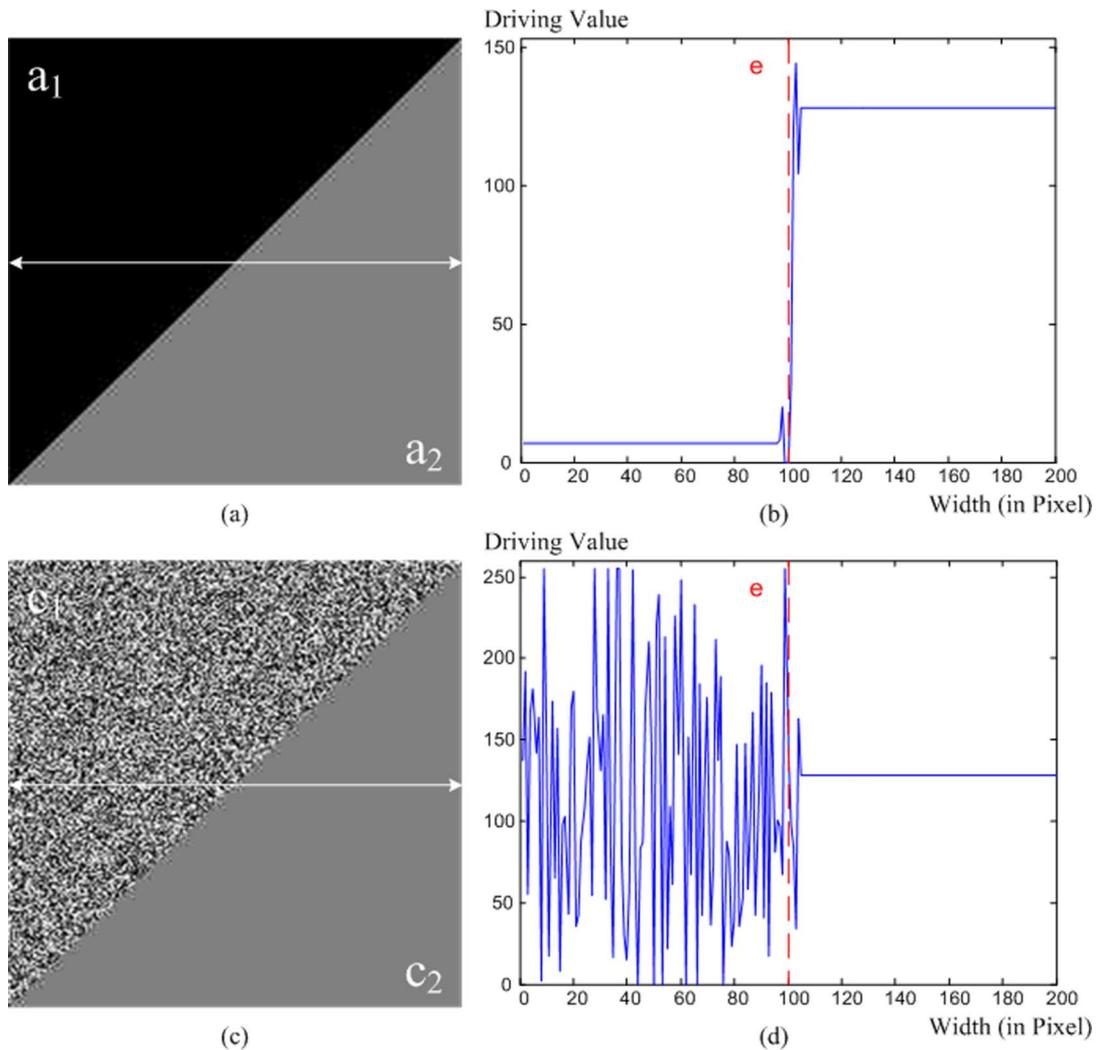
Fig. 2. (a), (b) Illustration of luminance and (c), (d) texture masking on ringing visibility for two image patches compressed with JPEG (MATLAB's *imwrite* function with $Q = 30$). Graphs (a) and (c) show the compressed image patches, whereas graphs (b) and (d) represent the intensity profile (in 8-bits driving values) along the row in the image indicated with the arrow in graphs (a) and (c), respectively. The dashed line "e" in graphs (b) and (d) refers to the position of the edge. Note that although both sides of the edge at "e" exhibit ringing artifacts, the visibility of ringing differs.

[18]), and ringing is more masking visible in homogenous areas than in textured or detailed areas (i.e., texture masking [19]).

The effect of luminance and texture masking on ringing visibility is illustrated in Fig. 2. Hence, to accurately detect regions

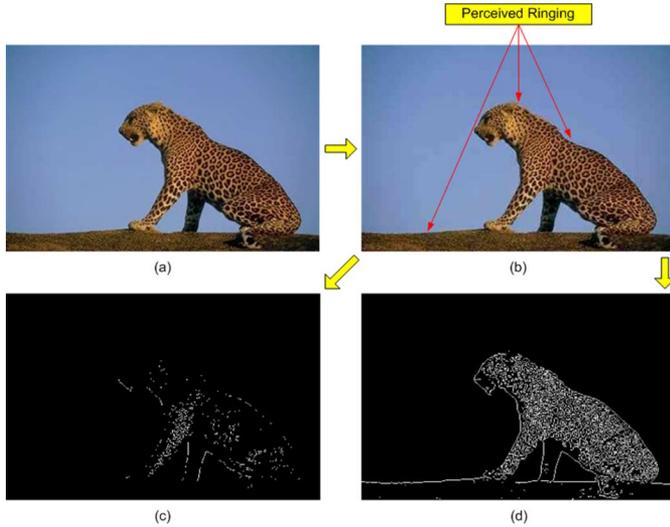Fig. 3. Ordinary edge detector (i.e., Sobel operator) applied for ringing region detection. (a) Original image. (b) JPEG compressed image (MATLAB's *imwrite* function with $Q = 30$). (c) Sobel edge map of (b) using a high threshold (i.e., 23% of I_{max}). (d) Sobel edge map of (b) using a low threshold (i.e., 10% of I_{max}).
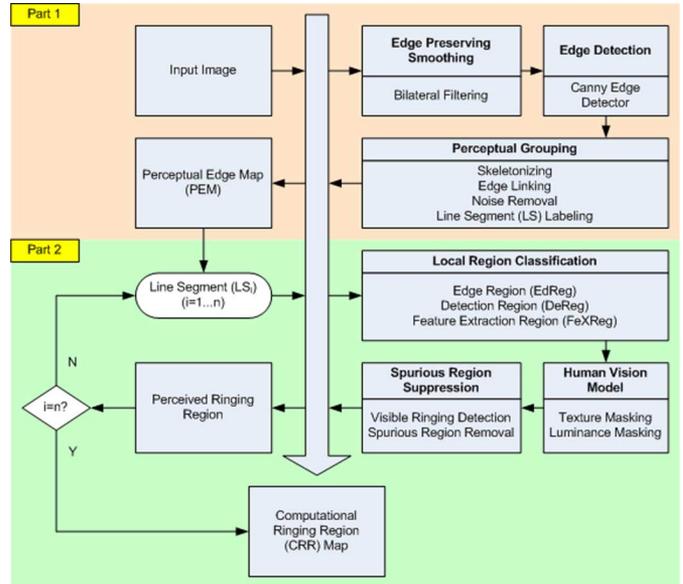

Fig. 4. Schematic overview of the proposed algorithm, with at the top the part to detect edges relevant for ringing, and at the bottom the part to measure visibility of ringing around these edges.

with perceived ringing, two essential aspects need to be explicitly addressed: 1) an (strong) edge detector; and 2) a masking model of the HVS.

A. Review of Related Work

Until recently, only a limited amount of research was devoted to perceived ringing. The methods in [20] and [21] both simply assume that ringing occurs unconditionally in regions surrounding strong edges in an image. This, however, does not always reflect human visual perception of ringing, because of the absence of spatial masking as typically present in the HVS. This issue is taken into account by incorporating properties of the HVS into the detection method, such as, for example, in [22] and [23]. The approach in [22] is based on the global edge map of an image, where binary morphological operators are used to generate a mask to expose regions that are likely to be contaminated with visible ringing artifacts. This procedure involves the identification of regions around all detected edges, and a further evaluation of these regions based on visual masking. In [23], a different way of including HVS masking properties is employed. This method classifies the potential smooth regions (i.e., regions in an image other than edges and their surroundings) into different objects based on their color similarity and texture features. The resulting objects are assigned as background around potential ringing regions. Texture masking is implemented by evaluating the contrast in activity between the potential ringing region and its assigned background (e.g., the higher the contrast in activity, the more visible ringing is assumed to be). Additionally, also luminance masking is implemented to further determine ringing visibility.

There are two main concerns with the methods existing in literature. First of all, the edge detection methods employed in [20]–[23] capture strong edges using an ordinary edge detector, such as a Sobel operator, where a certain threshold is applied to the gradient magnitudes to remove noise and insignificant edges. Depending on the choice of the threshold, these methods run the risk of omitting obvious ringing regions near

nondetected edges (in case of a high threshold) or of increasing the computational power by modeling the HVS near irrelevant edges (in case of a low threshold). Fig. 3 illustrates the effect of the threshold value of a Sobel operator. The edge map in Fig. 3(c), resulting from a high threshold value, largely removes noisy edges while eliminating a number of important edges, at which ringing obviously exists [see Fig. 3(b)]. This may heavily degrade the accuracy of the prediction of perceived ringing. By lowering the threshold [as in Fig. 3(d)], all strong edges are maintained in the edge map, but it also contains more texture edges, which are nonrelevant to ringing detection, and consequently, result in a large number of unnecessary computations for ringing visibility. The second concern with the existing methods is related to the models of the HVS used, for example, in [22] and [23], which are computationally very expensive. The HVS model in [22] involves a parameter estimation mechanism, which requires a number of calculations to achieve an optimal selection. The major cost of the HVS model in [23] is introduced by its clustering scheme embedded, which contains color clustering and texture clustering.

Obviously, the optimal performance in terms of reducing the number of required computations, while maintaining the reliable detection of perceived ringing, can be achieved by optimizing two aspects: 1) the detection accuracy of relevant edges; and 2) the reduction in complexity of the HVS model itself. Hence, what is needed is an edge detector that only extracts edges most closely related to the occurrence of ringing, and a HVS model that is simpler (and thus more applicable for real-time implementation) than the approaches existing in literature. In this paper, both aspects needed to efficiently detect regions with visible ringing are discussed.

II. PROPOSED ALGORITHM

The schematic overview of the proposed algorithm is illustrated in Fig. 4. It mainly consists of two parts: 1) extraction
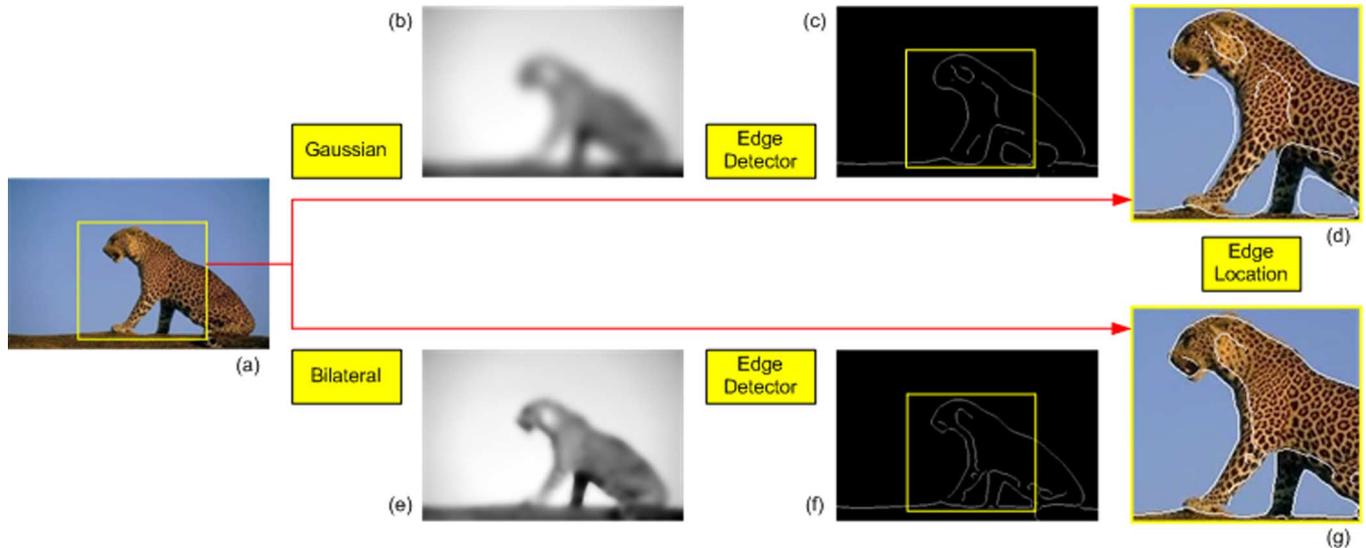
Fig. 5. Bilateral filtering and Gaussian filtering for the detection of perceptually strong edges. (a) Original image. (b) Gaussian filtered image ($\sigma_d = 15$). (c) Edge map of (b). (d) Superposition of (c) on (a). (e) Bilateral filtered image ($\sigma_d = 3, \sigma_r = 100$). (f) Edge map of (e). (g) Superposition of (f) on (a).

of edges relevant for ringing, and 2) detection of visibility of ringing in the edge regions. In the first part, an advanced edge detector is adopted, attempting to select the edges most relevant for ringing (i.e., contours of objects) in combination with the avoidance of the irrelevant edges (i.e., in textured areas). This results in a perceptual edge map (PEM), existing of a set of so-called line segments (LS). In the second part, each LS of the PEM is examined individually on the occurrence of visible ringing in its direct neighborhood, taking into account masking by the HVS. All regions with visible ringing are accumulated in a single binary map, which we refer to as the computational ringing region (CRR) map. Remind that the CRR map is used as input to the second step of the objective metric, in which the ringing annoyance is *quantified*, as published in [16] and [17]. Each part of the ringing region detection algorithm is further detailed in the following sections. The parameters used in the algorithm are specified and discussed in Section IV-B. Note that the entire metric is only based on the luminance channel of the images in order to further reduce the computational load.

A. Perceptual Edge Extraction

As explained above, the detection of visible ringing heavily relies on the accurate and efficient detection of object edges. To achieve this, we propose the application of a Canny edge detector [24] to an image, which first is nonlinearly smoothened. After some additional post-processing, this results in the PEM.

Edge Preserving Smoothing and Canny Edge Detection: When interpreting the surrounding world, humans tend to respond to differences between homogeneous regions rather than to structure within these homogeneous regions [25]. Hence, finding perceptually strong edges mainly implies that texture existing in homogenous regions can be neglected as if viewed from a long distance. This can be implemented by smoothing the image progressively until textual details are significantly reduced, and then applying an edge detector.

Traditional low-pass linear filtering (e.g., Gaussian filtering) smoothens out noise and texture, but also blurs edges, and con-

sequently, changes their spatial location. Since ringing detection intrinsically requires accurate spatial localization of the edges, edge-preserving smoothing is needed. Bilateral filtering was introduced in [26] as a simple and fast scheme for edge-preserving smoothing. It is a nonlinear operation that combines nearby image values based on both their geometric closeness and their photometric similarity, and prefers near values to distant values in both spatial domain and intensity range. In the Gaussian case, it can be expressed as

$$\vec{F}(\vec{x}) = \frac{\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \vec{I}(\vec{\xi})\omega(\vec{\xi}, \vec{x})d\vec{\xi}}{\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \omega(\vec{\xi}, \vec{x})d\vec{\xi}} \quad (1)$$

where

$$\omega(\vec{\xi}, \vec{x}) = \exp\left(\frac{-(\vec{\xi} - \vec{x})^2}{2\sigma_d^2}\right) \exp\left(\frac{-(I(\vec{\xi}) - I(\vec{x}))^2}{2\sigma_r^2}\right) \quad (2)$$

I and F denote the input and output images, χ and ξ are space variables, and the standard deviations σ_d and σ_r characterize the domain and range filtering, respectively. The advantage of using bilateral filtering instead of Gaussian filtering for the localization specific detection of perceptually strong edges is illustrated in Fig. 5.

Subsequently, a Canny edge detector is applied to the bilaterally filtered image to obtain the perceptually more meaningful edges. Since the input image is already filtered, the subsequent Canny algorithm is implemented without its inherent smoothing step, while keeping the other processing steps unchanged. The Canny edge detector uses two thresholds to detect strong and weak edges, and includes the weak edges in the output only if they are connected to strong edges. Their values is automatically set, depending on the image content.

Perceptual Edge Map Formation: Since the HVS does not perceive luminance variations at pixel level, the detected edge pixels are necessarily combined into perceptually salient elements, facilitating further analysis and processing [25]. These perceptual elements, which we refer to as line segments (LS),
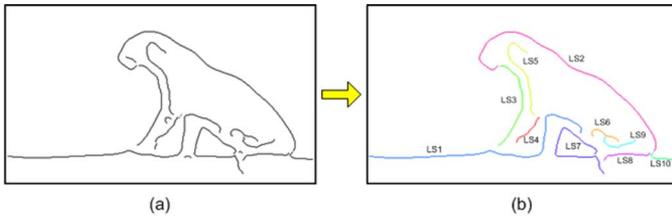
Fig. 6. Construction of the perceptual edge map (PEM). (a) Canny edge map. (b) Related PEM with labeled line segments.

are constructed over the Canny edge map and will be used as the basis for ringing region detection. The following processing steps are implemented to define the LS in the PEM.

1) Skeletonizing: To guarantee that an edge is only one-pixel thick, a kernel of 4×4 pixels is slid over all pixels, and those pixel configurations that have a structure of [1 1; 0 1] or [1 0; 1 1] are replaced by [1 0; 0 1], and those with a structure of [1 1; 1 0] or [0 1; 1 1] are replaced by [0 1; 1 0].

2) Edge Linking: The algorithm links all the edge pixels into a set of elements; each element either contains two endpoints or is a closed loop. If an edge junction is encountered, the tracing procedure breaks, and a separate element is generated for each of the branches.

3) Noise Removal: The elements with the number of connected edge pixels below a certain threshold are discarded. This is done with the ringing detection accuracy and speed in mind.

4) Line Segment Labeling: The resulting elements of connected edge pixels are referred to as line segments (LS), and labeled.

Once this process is complete, we have the PEM. Fig. 6 illustrates the labeling of the LS in the PEM.

B. Ringing Region Detection

Each LS of the PEM is examined individually on the occurrence of visible ringing artifacts in their direct neighborhood, taking into account luminance and texture masking. The regions with visible ringing are then combined in a computational ringing region (CRR) map.

Local Region Classification: In order to characterize the visibility of ringing around a LS, its surrounding is classified into three different zones (see Fig. 7(a) for an example of a single step edge): 1) Edge Region (EdReg): the original edge including the compression induced blur; 2) Detection Region (DeReg): the direct neighborhood of the EdReg, which potentially contains ringing artifacts; and 3) Feature Extraction Region (FeXReg): a region representative for the original local background, which is located outward from the corresponding DeReg. These regions are defined by thickening the LS with a different size for the structuring element of a dilation operation. Fig. 7(b) gives an example, in which for one LS [i.e., LS3 of Fig. 6(b)] the EdReg, DeReg, and FeXReg obtained with a square structuring element of 2, 9, and 17 pixels width, respectively, is shown.

Human Vision Model: Whether ringing is actually visible in the DeReg strongly depends (because of masking in the HVS) on the content of the original background, here represented
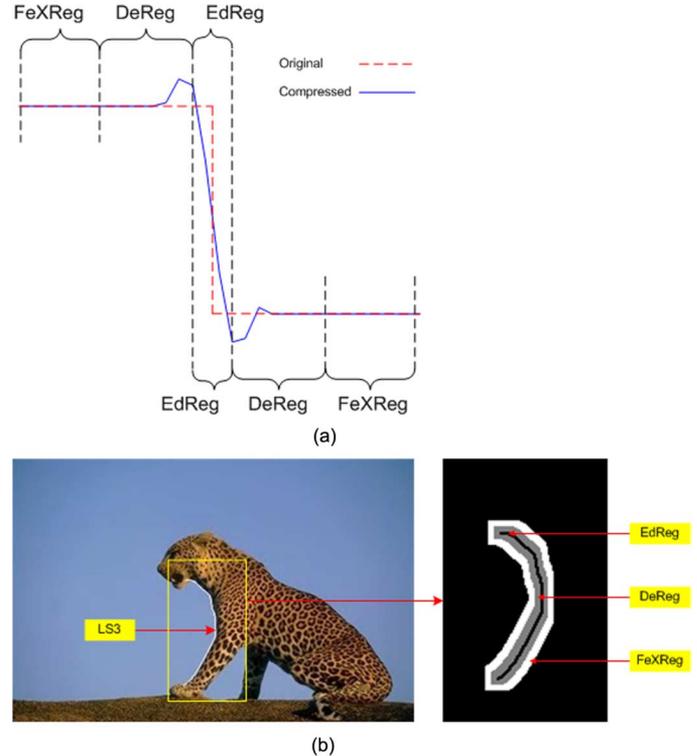

Fig. 7. Illustration of local region classification. (a) Illustration of the three zones for a schematic step edge. (b) Illustration of how the zones are defined around an actual line segment as part of a natural image. In (b) the black line indicates the EdReg, the gray area defines the DeReg, and the white area refers to the FeXReg.

by the FeXReg. Hence, the visibility of ringing is evaluated for each LS by applying a model for texture and luminance masking, using the texture and luminance characteristics of the FeXReg. As a result, DeReg regions, in which ringing is visually masked are eliminated, and only the perceptually prominent DeReg ringing regions remain.

1) Texture Masking: The visibility of ringing is significantly affected by the spatial activity in its local background, i.e., ringing is visually masked when located in a textured region, while it is perceptually prominent against a smooth background [22], [23], [27] as illustrated in Fig. 2. In this paper, texture masking is modeled classifying the FeXReg of each LS into "smooth" and "textured" objects, depending on the local background characteristics. The DeReg is segmented accordingly, and those DeReg regions of which the corresponding FeXReg is clustered as "textured" are removed. This approach intrinsically avoids explicit modeling of the HVS, and formulates texture masking as a simple yet efficient local pixel clustering procedure. The proposed scheme to implement this is illustrated in Fig. 8(b). It generally involves the following steps.

1) Calculating the local activity of the image content covered by the FeXReg by applying a global threshold to the gradient in pixel intensity to create a local binary map (LBM) of the FeXReg. This yields a profile of local pixel activities, and is formulated as in (3) and (4), shown at the bottom of the next page, where the local activity $LA(i, j)$ at location (i, j) is approximated by the gradient of the image intensity using a gradient operator (e.g., a Sobel operator).

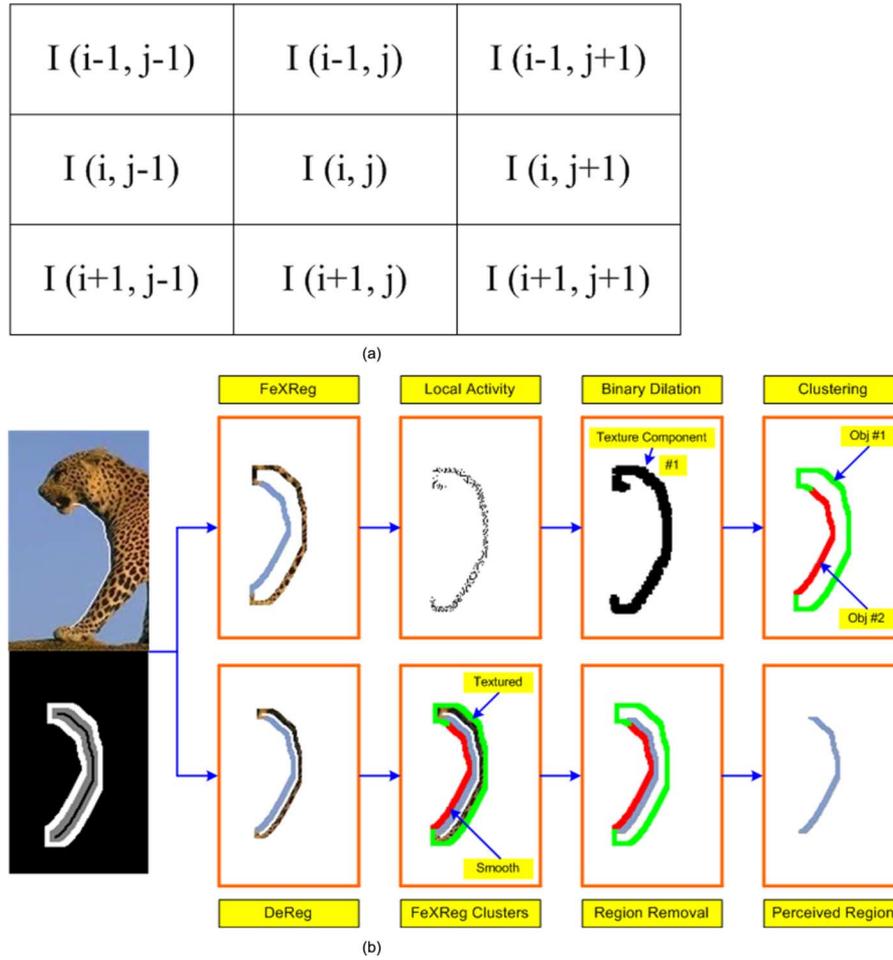| I (i-1, j-1)	I (i-1, j)	I (i-1, j+1)
I (i, j-1)	I (i, j)	I (i, j+1)
I (i+1, j-1)	I (i+1, j)	I (i+1, j+1)

(a)

(b)

Fig. 8. Implementation of texture masking. (a) Pseudoconvolution template used to calculate approximate gradient magnitude. (b) Illustration of the algorithm.

The 3×3 pseudoconvolution template used to calculate the gradient magnitude of a pixel at location (i, j) is shown in Fig. 8(a) ($I(i, j)$ corresponds to the pixel intensity at location (i, j)). The threshold *Thr_txt* is related to the magnitude histogram of the gradient image, and thus, image content dependent.

2) Dilating the LBM using a morphological operator, and labeling (e.g., by 8-connectivity) them into a set of connected components, which are referred to as texture components. This step intrinsically transfers pixel activities to a higher level structure of region activities, motivated by the fact that the human eye is not sensitive to variations at pixel level.

3) Classifying all FeXReg covered by texture components into "texture objects," and the remaining FeXReg into "smooth objects."

4) Removing the regions of DeReg that belong to the "texture objects" of FeXReg, since in these regions ringing is supposed to be masked by texture, and discarding the resulting regions of DeReg with their size under a certain threshold. The maintained regions of DeReg are considered as perceived ringing regions.

2) Luminance Masking: The visibility of variations in luminance depends on the local mean luminance [18], [19], [27]–[29]. As a result, the visibility of ringing is largely reduced in extremely dark or bright surroundings, as illustrated in Fig. 2. The implementation of luminance masking is the same as for texture masking, but to guarantee efficiency, it is only applied

$$\text{LBM}(i,j) = \begin{cases} 0 & \text{LA}(i,j) < \text{Thr\_txt} \\ 1 & \text{otherwise} \end{cases}, \qquad i,j \in \text{FeXReg} \tag{3}$$

$$\begin{aligned} \text{LA}(i,j) = &\left| [I(i-1,j-1) + 2 \times I(i-1,j) + I(i-1,j+1)] - [I(i+1,j-1) + 2 \times I(i+1,j) + I(i+1,j+1)] \right| \\ &+ \left| [I(i-1,j+1) + 2 \times I(i,j+1) + I(i+1,j+1)] - [I(i-1,j-1) + 2 \times I(i,j-1) + I(i+1,j-1)] \right| \end{aligned} \tag{4}$$
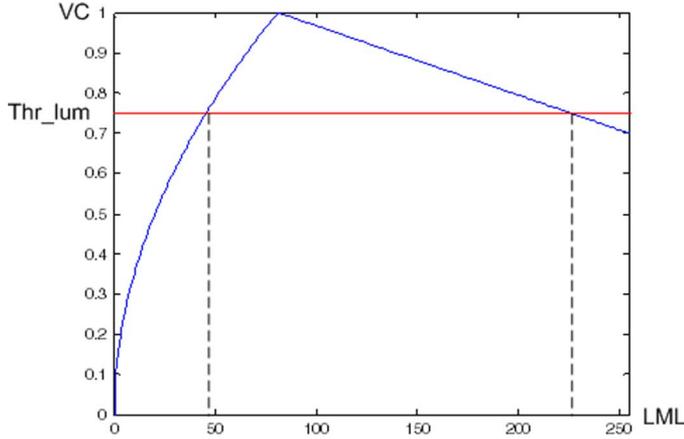
Fig. 9. Implementation of luminance masking via the relation between the local mean luminance (LML) and the artifact visibility coefficient (VC). Thr_lum refers to the threshold used in the implementation.

to those regions of the DeReg remaining after the application of texture masking. The procedure for luminance masking is similarly formulated as a local pixel clustering model, and it mainly contains the following steps.

1) Calculating the local averaged luminance, over a 3×3 template, centered on each pixel that is part of a "smooth object" of the FeXReg

$$\text{LML}(i,j) = \frac{1}{9} \sum_{k=i-1}^{i+1} \sum_{l=j-1}^{j+1} I(k,l) \qquad (5)$$

where $I(i,j)$ denotes the pixel intensity at location (i,j), and $\text{LML}(i,j)$ denotes the local mean luminance. The visibility of ringing due to luminance masking is determined according to the functional behavior shown in Fig. 9 [12], and a local binary map (LBM) is generated by applying a predefined threshold to the visibility coefficient (VC)

$$\text{LBM}(i,j) = \begin{cases} 0 & \text{VC}(i,j) > \text{Thr\_lum} \\ 1 & \text{otherwise} \end{cases} \qquad (6)$$

where $\text{LBM}(i,j) = 0$ indicates a visible pixel location, and $\text{LBM}(i,j) = 1$ indicates a nonvisible pixel location. This generates a profile of local visibility due to luminance masking.

2) Dilating the LBM to obtain a set of connected components, which are referred to as invisible components.

3) Classifying the "smooth objects" of FeXReg further into "visible objects" and "invisible objects" depending on the invisible components. This step combined with the one mentioned above intrinsically yields the structures of region visibility.

4) Removing the DeReg that correspond to "invisible objects," i.e., where ringing is not supposed to be visible against a very low or very high intensity background.

Ultimately, only the regions of DeReg that yield visible ringing remain. These regions are combined in the CRR map, of which an example is given in Fig. 10.

C. Spurious Ringing Region Suppression

The ringing region detection method described so far only exposes regions in an image which are likely to be impaired
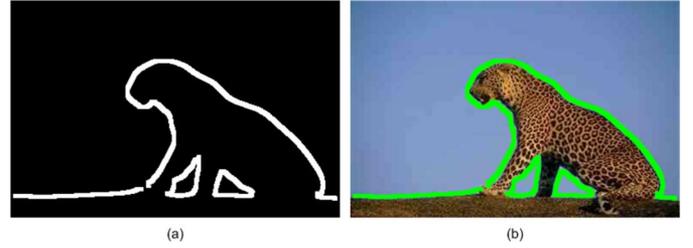

Fig. 10. Example of (a) a computational ringing region (CRR) map corresponding to (b) a JPEG compressed image.

by visible ringing artifacts. The resulting CRR map, however, still includes obvious spurious ringing regions, containing either "unimpaired" or "noisy" pixels misinterpreted as ringing pixels.

"Unimpaired pixels" indicate pixels in the detected regions of the CRR map, which are actually not impaired by ringing. An obvious example of the occurrence of "unimpaired" pixels is in an uncompressed image. The ringing region detection algorithm described so far will find the regions that might be impaired with visible ringing, independent of the compression level. But in an uncompressed image, these regions do not contain visible ringing, and hence, should be removed from the CRR map. Note that without removal of these regions the overall objective ringing metric including the step of quantification of ringing annoyance (see [16] and [17]) would not be less accurate, but less efficient.

"Noisy pixels" are pixels in the detected regions of the CRR map, that actually belong to an edge or texture. They are accidentally misclassified to a ringing region as a consequence of the dilation operation used in the human vision model.

To remove the spurious ringing regions, each detected ringing region (RR) is further examined by calculating its amount of visible ringing pixels. Those RRs with their number of visible ringing pixels below a certain threshold are considered as spurious, and consequently removed from the CRR map. Whether a pixel in a RR is a visible ringing pixel is determined via the local variance (LV) in intensity in its 3×3 neighborhood. The spurious ringing pixels are suppressed by applying two thresholds to the LV, a low threshold (*Thr_v_low*) and a high threshold (*Thr_v_high*). Since unimpaired pixels exhibit no or very small intensity variance in their neighborhood, a pixel with its LV value below or equal to *Thr_v_low* is considered as an unimpaired pixel. In the same way, a pixel with its LV value above or equal to *Thr_v_high* is considered as a "noisy pixel." This can be formulated as

$$\text{VC}_n(i,j) = \begin{cases} 1 & \text{Thr\_v\_low} < \text{LV}(i,j) < \text{Thr\_v\_high} \\ 0 & \text{otherwise} \end{cases} \qquad (7)$$

where

$$\text{LV}(i,j) = \frac{1}{9} \sum_{k=i-1}^{i+1} \sum_{l=j-1}^{j+1} \left[I(k,l) - \frac{1}{9} \sum_{k=i-1}^{i+1} \sum_{l=j-1}^{j+1} I(k,l) \right]^2,$$
$$i,j \in \text{RR}_n \qquad (8)$$

$$\text{Thr\_v\_high} = \alpha \cdot \text{MAX}[\text{LV}(i,j)], \qquad i,j \in \text{LS}_n \qquad (9)$$

where $\text{VC}_n(i,j)$ indicates the visibility of a ringing pixel at the nth ringing region (i.e., RR_n) with its associated line seg-

Fig. 11. Source images.

ment (i.e., LS_n), and $LV(i,j)$ indicates the local variance computed over a 3×3 template, centered at a pixel intensity $I(i,j)$. The value of *Thr_v_low* is chosen to be zero, and the value of *Thr_v_high* is chosen to scale with the strength of corresponding edge (see [23]). Thus, the ringing region RR_n is removed if

$$\frac{\text{SUM}(\text{VC}_n)}{\text{SIZE}(\text{RR}_n)} < R \qquad (10)$$

where $\text{SUM}(\text{VC}_n)$ indicates the number of visible ringing pixels, $\text{SIZE}(\text{RR}_n)$ indicates the size of the given RR, and R indicates the predefined ratio of visible ringing pixels over the detected ringing region.

III. PSYCHOVISUAL EXPERIMENT

To validate our algorithm for ringing region detection, a psychovisual experiment, in which participants were requested to indicate regions of visible ringing in compressed natural images, was carried out.[1] The results were transformed into a subjective ringing region (SRR) map, indicating where in an image on average people see ringing.

A. Subjective Experiment Procedure

A set of eight source images, reflecting adequate diversity in image content, were taken from the Kodak Lossless True Color Image Suite [30]. Fig. 11 shows these source images. They were high-resolution and high-quality color images of size 768×512 (width $\times$ height) pixels. These images were JPEG compressed using MATLAB's *imwrite* function at two different compression levels (i.e., $Q = 25$ and 50). This yielded a test database of 16 stimuli. These stimuli were displayed on a 17-inch LCD monitor with a screen resolution of 1024×768 pixels. The experiment was conducted in a standard office environment [31] and the viewing distance was approximately 40 cm.

Twelve students of the Delft University of Technology, being eight males and four females, were recruited for the experiment. Before they started the actual assessment, each of them was shown three examples of synthetic ringing, synthetic blocking and synthetic blur artifacts, followed by three real-life images

in which ringing, blocking and blur were the most annoying artifacts, respectively. When the participant reported to be able to distinguish ringing from other types of compression artifacts, a set of images with the same level of ringing as used in the rest of the experiment was presented. The participant was requested to mark any region in the image where he/she perceived ringing, independent of its annoyance. The images used during this training were different from those used in the actual experiment. After training, all 16 stimuli were shown in a random order to each subject in a separate session.

B. Subjective Data Processing

The recorded edges per image and subject were transformed into a binary image, in which a white pixel indicated perceived ringing and a black pixel referred to absence of visible ringing. This resulted in an individual ringing region (IRR) map per stimulus and subject. These IRR were then averaged over all subjects to a mean ringing region (MRR) map. From the MRR map, the subjective ringing region (SRR) map was derived by simply applying a threshold (i.e., *Thr_srr*) of 0.5, keeping only those edges near which ringing was perceived by half of the subjects. This threshold was introduced to avoid that subjective outliers would strongly affect the performance comparison between various algorithms. Its actual value is further discussed in Section V.

IV. PERFORMANCE EVALUATION

Our proposed ringing region detection method is validated with respect to the results of the psychovisual experiment, and its performance is compared to existing alternatives in literature. For this performance comparison, we implemented three ringing region detection algorithms recently proposed: 1) region clustering based ringing artifact measure (referred to as RCRM) [23]; 2) morphological filtering based ringing artifact measure (referred to as MFRM) [22]; and 3) no-reference ringing artifact measure (referred to as NRRM) [21]. In literature, all three methods are proved to be promising in terms of ringing region detection.

A. Evaluation Criteria

To evaluate the performance of various ringing region detection algorithms we compared the CRR map as calculated for

---

[1]The data collected from this experiment are available to the image quality assessment community at http://mmi.tudelft.nl/~ingrid/ringing.html

each of the ringing region detection algorithms to the SRR map derived from the psychovisual experiment. These two binary images (i.e., the CRR and SRR map) were compared visually and via a quantitative correlation.

For the visual assessment we produced a comparison map M_C), which is an RGB color image generated by

$$M_C = \begin{cases} M_C(:,:,1) = M_{\mathrm{CRR}} \,\&\, [\mathrm{xor}(M_{\mathrm{CRR}}, M_{\mathrm{SRR}})] \\ M_C(:,:,2) = M_{\mathrm{CRR}} \,\&\, M_{\mathrm{SRR}} \\ M_C(:,:,3) = M_{\mathrm{SRR}} \,\&\, [\mathrm{xor}(M_{\mathrm{CRR}}, M_{\mathrm{SRR}})]. \end{cases} \tag{11}$$

The G (green) channel is assigned to the logical operator AND of the two binary maps, and so, represents the correlated ringing regions. The R (red) and B (blue) channels are assigned to edges occurring only in the CRR map and the SRR map, respectively, and so, represent the uncorrelated ringing regions between both maps. Black regions represent the absence of visible ringing on both maps.

The objective comparison of the CRR map to the SRR map is quantitatively measured by two correlation coefficients, namely $\rho 1$ and $\rho 2$, defined as follows:

$$\rho_1 = \frac{\sum [M_{\mathrm{CRR}} \,\&\, M_{\mathrm{SRR}}]}{\sum M_{\mathrm{SRR}}} \tag{12}$$

$$\rho_2 = \frac{\sum \{M_{\mathrm{CRR}} \,\&\, [\mathrm{xor}(M_{\mathrm{CRR}}, M_{\mathrm{SRR}})]\}}{\sum [\sim M_{\mathrm{SRR}}]}. \tag{13}$$

The numerator of ρ_1 indicates the total number of correlated pixels between the CRR map and SRR map, while the denominator indicates the size of the ringing regions in the SRR map. Thus, ρ_1 quantifies to what extent the subjective ringing regions are detected by the computational models. However, this coefficient by itself is obviously not enough to reflect the detection accuracy of a computational model. A model might be capable of capturing all subjective ringing regions, just by capturing all edges, also those that do not contain visible ringing. These falsely detected ringing regions consequently degrade particularly the efficiency of a subsequent ringing annoyance measurement. The degree of false detections is quantified by ρ_2. Its numerator indicates the size of regions falsely detected by the computational models, and its denominator indicates the size of regions in the SRR map not detected by the human subjects. Evidently, a higher value of ρ_1 combined with a lower value of ρ_2 implies a good detection model.

B. Model Calibration

Our proposed ringing region detection algorithm uses a number of parameters that need to be tuned to optimal, but at the same time robust performance over different image content. For this tuning, we used five new images (not part of the psychovisual experiment). These images were also JPEG compressed with the MATLAB's *imwrite* function at $Q = 25$ and 50. A few experts in the area of compression artifacts (mainly the authors) indicated the regions in the image with visible ringing. The resulting data were used for optimizing the performance of our ringing region detection algorithm. Robustness over content was evaluated by applying these optimized parameters to the new image content of the psychovisual experiment.

Parameters for the Edge Extraction: This set of parameters includes the standard deviations (i.e., σ_d and σ_r) for the bilateral filter to control the extent of the smoothing effect, and the hysteresis thresholding (i.e., *Thr_high* and *Thr_low*) of the Canny edge detector to trace strong edges while preventing breaking of continuous edges. For the bilateral filter the selection of σ_d and σ_r has been intensively discussed for natural images in [26], and they were set accordingly to $\sigma_d = 3$ and $\sigma_r = 100$ in our experiment (see [32] and [33]). For the edge detector Canny sets the *Thr_high* such that a certain percentage (i.e., p) of the total amount of pixels is cumulated in the magnitude histogram of the gradient image, and the *Thr_low* as a fixed fraction (i.e., 0.4) of the *Thr_high* [24]. In our implementation, we used a relatively low value of *Thr_high* (i.e., $p = 85\%$) in order to prevent loosing relevant edges. This may result in irrelevant LSs in the PEM, but these LSs are later discarded by applying the HVS model. In other words, the choice for the thresholds of the Canny edge detector affect the efficiency of the model rather than its accuracy. Finally, the threshold for the noise removal in the PEM formation was set to 20 pixels. Again, this parameter affects the efficiency rather than the accuracy of the model.

Parameters for Region Definition: This set of parameters determines the width of the EdReg, DeReg, and FeXReg regions. The EdReg representing edge blur is chosen to be equal to the one-pixel thick LS. In case this value is too small, blur pixels can easily be detected as spurious pixels in a ringing region (as described in Section II-C). The width of the DeReg is set as a single-sided support dimension of four pixels, which approximates the maximal extent of ringing that spreads out to a region surrounding an edge in JPEG compression [4]. The actual width of the DeReg may vary depending on the underlying properties of the coding technique, but can be adjusted according to [34]. The width of the FeXReg is empirically selected to be the same as for the DeReg. We experienced that the FeXReg may cross an object boundary or reach another edge, which consequently results in spurious pixels in a detected ringing region. The suppression of these pixels has been discussed in Section II-C.

Parameters for the HVS: This set of parameters includes two essential thresholds, i.e., *Thr_txt* for texture masking and *Thr_lum* for luminance masking. The performance of our algorithm is fairly insensitive to variations of these thresholds within the range of [0.6, 0.95] and [0, 0.8] for *Thr_txt* and *Thr_lum*, respectively. Varying these thresholds within their respective range results in a variation of ρ_1 and ρ_2 over [85%, 95%] and [1%, 3%], respectively. For the final performance evaluation of our model, we set Thr_txt $= 0.9$ and Thr_lum $= 0.75$.

Parameters for Spurious Ringing Pixel Detection: This set of parameters contains three threshold values (i.e., *Thr_v_low*, *Thr_v_high* [determined by α as shown in (9)] and R) to further eliminate undesired regions in the CRR map. It should be admitted that this processing step is a fine-tuned optimization to largely remove, for example, the "unimpaired regions" in the CRR map of an uncompressed (or high bitrate compressed) image, thus making the subsequent calculation of ringing annoyance [16], [17] more efficient. The parameters are determined as Thr_v_low $= 0, \alpha = 0.5$ and $R = 0.3$. *Thr_v_low* and α are set according to experiments and observations reported in [23], while R is empirically chosen. R is mainly used to speed

up the algorithm rather than to improve its accuracy. The inclusion of the detection of spurious ringing pixels hardly affects the overall performance of our model: including or omitting the detection of spurious ringing pixels corresponds to a deviation in ρ_1 and ρ_2 over a range of $[-0.5\%, +0.5\%]$. It should, however, be noted that the concept of removing spurious ringing pixels is mainly important for the ringing annoyance estimation, and hence, these parameters might need to be calibrated again for the subjective data of ringing annoyance [16], [17].

Selected Parameters for Methods From the Literature: As mentioned above, we will compare the performance of our algorithm to three alternatives published in the literature. These methods were implemented following the description in the original publications. However, some important parameters were adjusted to ensure a fair comparison. The parameters to control the thresholding of the edge detector were tuned for each algorithm to yield the highest performance possible for the five test images used during calibration. The parameter for determining the extent of ringing artifacts was equal for all algorithms (i.e., a single-sided ringing region support dimension of 4 pixels).

C. Evaluation of Overall Model Performance

The comparison maps for the visual assessment between the SRR map and the (optimized) CRR maps of the various algorithms are given in Fig. 12. The first column shows the test images, the second column presents the SRR maps, and the remaining four columns give the comparison maps of our proposed algorithm, RCRM, MFRM, and NRRM, respectively. In general, most of the ringing regions that were perceived in the psychovisual experiment were also detected by each of the four algorithms. However, our proposed method detects the perceived ringing regions while introducing far less noise (i.e., regions that are not observed subjectively) compared to the other three methods. The correlation coeffcients ρ_1 and ρ_2 between the SRR and each of the CRR maps is given in Fig. 13. These data are summarized into an overall performance, shown in Table I. In terms of detecting perceived ringing regions (i.e., ρ_1), our proposed method outperforms the other three methods by 15% on average. Also in terms of avoiding false detection (i.e., ρ_2) our method is twice as good as the next best one, namely the RCRM. The latter algorithm, however, is lowest in performance based on ρ_1.

V. DISCUSSION

In this paper we present a novel approach to the detection of regions in an image impaired with visible ringing artifacts. The output of the proposed algorithm serves as input for the second step in the objective ringing metric, existing of the quantification of the actual ringing annoyance in each of the detected regions (as published in [16] and [17]). In this respect it is relevant to realize that a good performance of the ringing region detection algorithm mainly contributes to the efficiency of the second step in the objective metric, rather than to the final accuracy of the prediction in ringing annoyance.

So far, our algorithm is only tested for JPEG compressed image material. More research is needed to also evaluate its performance for different compression techniques. The algorithm
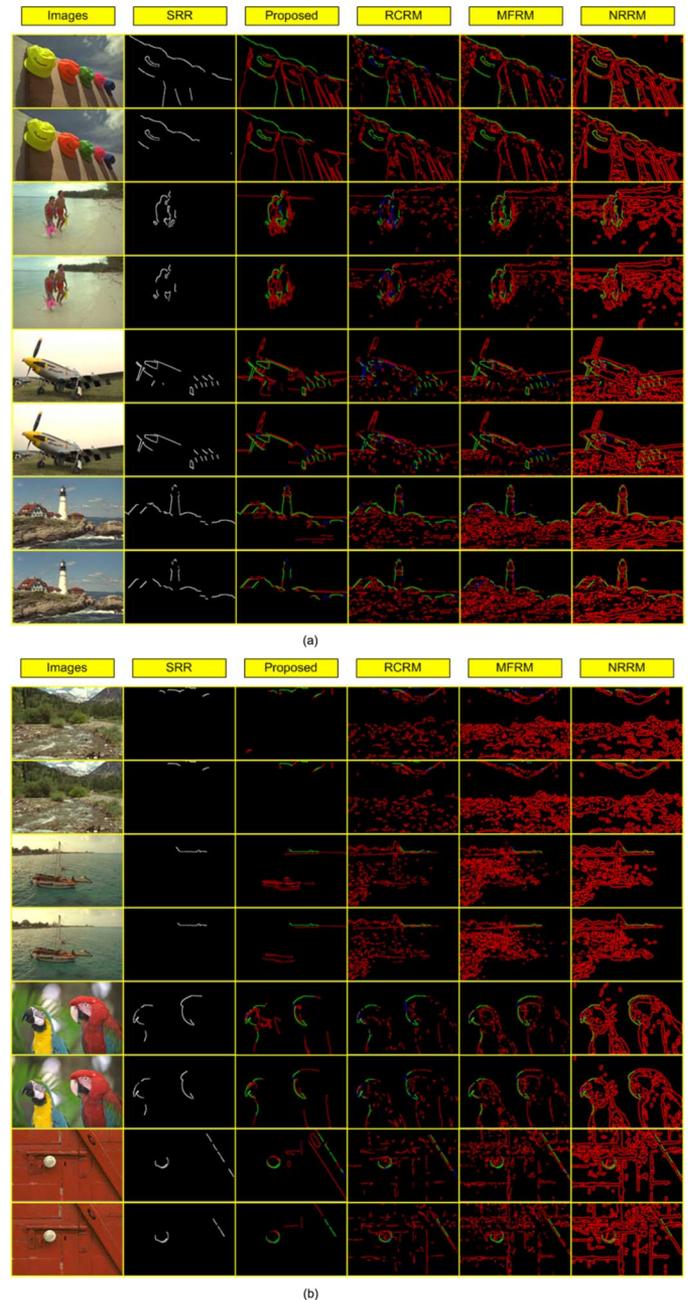

Fig. 12. Experimental results of visual assessment. (a) Images 1-8: Caps (Q25), Caps (Q50), Beach (Q25), Beach (Q50), Plane (Q25), Plane (Q50), Lighthouse (Q25), Lighthouse (Q50). (b) Images 9–16: Stream (Q25), Stream (Q50), Sailing (Q25), Sailing (Q50), Parrots (Q25), Parrots (Q50), Door (Q25), Door (Q50). The second column gives the subjective ringing region (SRR) map, and columns 3–6 give the computational ringing region (CRR) map calculated for our proposed approach, the RCRM [24], the MFRM [23], and the NRRM [22], respectively.

is evaluated for two compression levels, and the corresponding CRR maps are highly comparable. Since in this paper we only measure ringing regions, and not ringing annoyance, this is not surprising. Even for uncompressed images the CRR map will be comparable at first instance, i.e., before removal of spurious ringing regions as discussed in Section II-C.

Our proposed ringing region detection algorithm exists of two essential contributions: an edge detector that only preserves per-
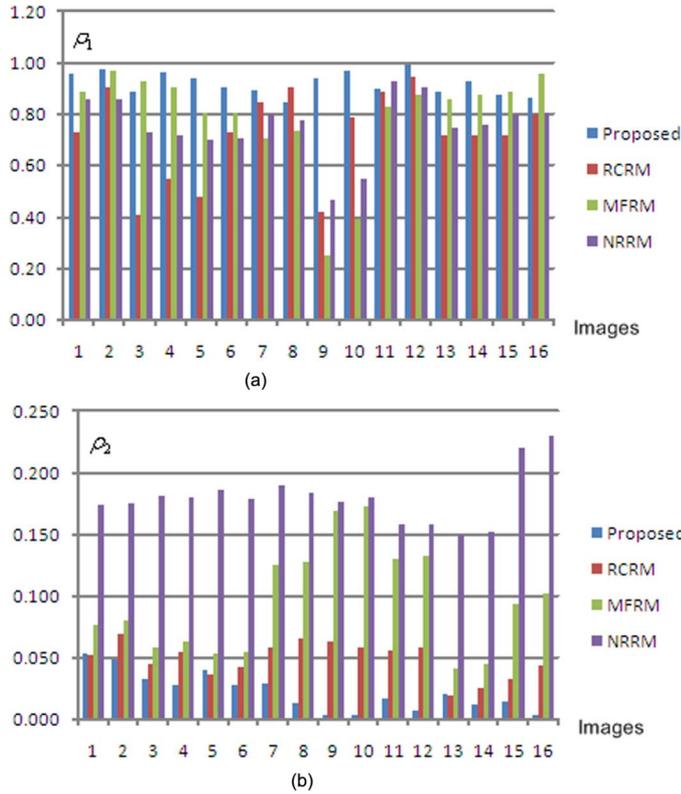
Fig. 13. Quantitative comparison results. (a) Correlation coefficient ρ_1. (b) Correlation coefficient ρ_2.

TABLE I
PERFORMANCE COMPARISON OF THE FOUR RINGING REGION DETECTION METHODS (Thr_STT $= 1/2$ FOR THE SRR MAPS): MEAN AND STANDARD DEVIATION OF THE CORRELATION COEFFICIENTS ρ_1 AND ρ_2

Model	Proposed	RCRM	MFRM	NRRM
$\overline{\rho_1}$	92%	72%	79%	76%
$\sigma(\rho_1)$	0.04	0.17	0.20	0.12
$\overline{\rho_2}$	2.2%	4.9%	9.6%	18%
$\sigma(\rho_2)$	0.02	0.02	0.04	0.02

TABLE II
PERFORMANCE COMPARISON OF THE FOUR RINGING REGION DETECTION METHODS FOR Thr_srr $= 1/3$ OF THE SRR MAPS: MEAN AND STANDARD DEVIATION OF THE CORRELATION COEFFICIENTS ρ_1 AND ρ_2

Model	Proposed	RCRM	MFRM	NRRM
$\overline{\rho_1}$	86%	74%	65%	83%
$\sigma(\rho_1)$	0.09	0.16	0.12	0.12
$\overline{\rho_2}$	4.5%	20%	11%	31%
$\sigma(\rho_2)$	0.03	0.08	0.03	0.04

ceptually relevant edges and a simple, yet efficient HVS. The use of an ordinary edge detector (as in RCRM, MFRM, and NRRM) makes ringing region detection very sensitive to the threshold used; for a high threshold some visually salient edges may not be detected, such that the obvious ringing regions are consequently missed, while for a low threshold many irrelevant edges may be retained, which results in a lot of false ringing regions. Especially for content that is rather insensitive to masking by the HVS (the image "Door" (see Fig. 12) is such an example), the number of detected ringing regions strongly depends on the threshold used for the edge detection. The value of our approach is mainly generated by the bilateral filtering (preserving the perceptually relevant edges) rather than by the edge detection itself. The Canny edge detector could have been replaced by a different edge detector, without expected change in performance.

Table I illustrates the advantage of using texture and luminance masking in ringing region detection (as in our proposed method, and in RCRM and MFRM). It obviously reduces the number of detected false ringing regions (lower ρ_2 value). The NRRM, not including HVS properties, clearly has the highest ρ_2 value. From a practical point of view, this may significantly degrade the efficiency, and to some extent the accuracy of predicting ringing annoyance. Including HVS modeling is especially crucial for highly textured images, such as the image "Stream" (see Fig. 12). This type of content usually masks ringing to a considerable extent, which should be addressed by a robust HVS model. That our HVS model is sufficiently robust against this demanding content is shown by its highest ρ_1 value and its lowest ρ_2 value compared to the other two algorithms including HVS properties (i.e., RCRM and MFRM; see Fig. 13). Additionally, It should be noted that the number of required computations for modeling the HVS is significantly lower for our model than for to the methods RCRM and MFRM. The reduction in complexity is achieved by calculating the HVS only near the perceptually relevant edges and also by simplifying the model of visual masking itself.

The third contribution to our ringing region detection algorithm is a rather ad hoc one: the removal of spurious ringing regions. Due to this spurious ringing region removal, our proposed method captures slightly more visible ringing regions for compression level $Q = 25$ than for compression level $Q = 50$, which is in agreement with the corresponding SRR maps. The impact of compression ratio is less obvious for the other alternative methods. However, this difference in performance is not of major concern, since it can be corrected for in the quantification of actual ringing annoyance, as long as all relevant edges are captured in the ringing region detection.

The performance of the ringing region detection algorithms is evaluated against the results of a psychovisual experiment, represented by SRR maps. From the visual assessment in Fig. 12, it is clear that all algorithms detect ringing regions that do not occur in the SRR maps. This is not surprising, since the SRR maps are derived such that they only maintain ringing regions detected by most of the participants. Hence, it is possible that some perceptible, but not annoying ringing regions are omitted by applying a threshold to the MRR maps (see Section III-B). To evaluate how the selection of this threshold affects the performance of all algorithms, the correlation coefficients ρ_1 and ρ_2 are recalculated for a lower threshold of the SRR map (i.e., Thr_STT $= 1/3$). The results, summarized in Table II, indicate that the actual values of ρ_1 and ρ_2 change for all algorithms, but that the general tendencies are maintained.

VI. CONCLUSION

In this paper, a novel approach toward the detection of perceived ringing regions in compressed images is presented. The algorithm relies on the compressed image only, which is promising for its applicability in a real-time video chain, e.g., to enhance the quality of artifact impaired video. It adopts a perceptually more meaningful edge detection method for the purpose of ringing region location. This intrinsically avoids the drawback of applying an ordinary edge detector, which has the risk of omitting obvious ringing artifacts near nondetected edges or of increasing the computational cost by measuring ringing visibility near irrelevant edges. The objective detection in agreement with human visual perception of ringing artifacts is ensured by taking into account typical properties of the human visual system, such as texture masking and luminance masking. The human vision model is implemented, based on the local image characteristics around detected edges, to expose only the perceptually prominent ringing regions in an image. The proposed detection method is validated with respect to ringing regions resulting from a psychovisual experiment, and shows to be highly consistent with subjective data. The performance of our approach is compared to existing alternatives in literature, and has been proved to be promising in terms of both reliability and computational efficiency. The proposed ringing region detection method is meanwhile extended with a ringing annoyance metric that can quantify perceived ringing annoyance of compressed images [16], [17].

REFERENCES

[1] Z. Wang and A. C. Bovik, *Modern Image Quality Assessment. Synthesis Lectures on Image. Video & Multimedia Processing*. San Rafael, CA: Morgan and Claypool, 2006.

[2] M. Ghanbari, *Standard Codecs: Image Compression to Advanced Video Coding*. London, U.K.: IEE Press, 2003.

[3] I. Richardson, *H.264 and MPEG-4 Video Compression: Video Coding for Next-generation Multimedia*. Hoboken, NJ: Wiley , 2003.

[4] M. Yuen and H. R. Wu, "A survey of hybrid MC/DPCM/DCT video coding distortions," *Signal Process.*, vol. 70, no. 3, pp. 247–278, 1998.

[5] M. C. Q. Farias, M. S. Moore, J. M. Foley, and S. K. Mitra, "Perceptual contributions of blocking, blurring, and fuzzy impairments to overall annoyance," in *Proc. SPIE, Human Vision and Electron. Imaging IX*, 2004, vol. 5292, pp. 109–120.

[6] C. C. Koh, S. K. Mitra, J. M. Foley, and I. Heynderickx, "Annoyance of individual artifacts in mpeg-2 compressed video and their relation to overall annoyance," in *Proc. SPIE, Human Vision and Electronic Imaging X*, 2005, vol. 5666, pp. 595–606.

[7] J. Xia, Y. Shi, K. Teunissen, and I. Heynderickx, "Perceivable artifacts in compressed video and their relation to video quality," *Signal Process. Image Commun.*, vol. 24, no. 7, pp. 548–556, 2009.

[8] M. Shen and C. J. Kuo, "Review of postprocessing techniques for compression artifact removal," *J. Vis. Commun. Image Rep.*, vol. 9, no. 1, pp. 2–14, 1998.

[9] J. Luo, C. Chen, K. Parker, and T. S. Huang, "Artifact reduction in low bit rate dct-based image compression," *IEEE Trans. Image Processing*, vol. 5, pp. 1363–1368, 1996.

[10] K. Zon and W. Ali, "Automated video chain optimization," *IEEE Trans. Consum. Electron.*, vol. 47, pp. 593–603, 2001.

[11] Z. Wang, A. C. Bovik, and B. L. Evans, "Blind measurement of blocking artifacts in images," in *Proc. IEEE Int. Conf. Image Processing*, 2000, vol. 3, pp. 981–984.

[12] H. Liu and I. Heynderickx, "A perceptually relevant no-reference blockiness metric based on local image characteristics," *EURASIP J. Adv. Signal Process.*, vol. 2009, 2009.

[13] Final Report From the Video Quality Experts Group on the Validation of Objective Models of Video Quality Assessment 2003 [Online]. Available: http://www.vqeg.org

[14] R. Muijs and J. Tegenbosch, "Quality-adaptive sharpness enhancement based on a no-reference blockiness metric," in *Proc. 2nd Int. Workshop on Video Processing and Quality Metrics for Consumer Electronics*, 2006.

[15] I. O. Kirenko, R. Muijs, and L. Shao, "Coding artifact reduction using non-reference block grid visibility measure," in *Proc. IEEE Int. Conf. Multimedia and Expo*, 2006, pp. 469–472.

[16] H. Liu, N. Klomp, and I. Heynderickx, "A no-reference metric for perceived ringing," in *Proc. 4th Int. Workshop on Video Processing and Quality Metrics for Consumer Electronics*, 2009.

[17] H. Liu, N. Klomp, and I. Heynderickx, "A no-reference metric for perceived ringing artifacts in images," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 20, no. 4, Apr. 2010, to be published.

[18] A. B. Watson, *Digital Image and Human Vision*. Cambridge, MA: MIT Press, 1993.

[19] B. A. Wandell, *Foundations of Vision*. Sunderland, MA: Sinauer, 1995.

[20] P. Marziliano, F. Dufax, S. Winkler, and T. Ebrahimi, "Perceptual blur and ringing metrics: Application to jpeg2000," *Signal Processing: Image Commun.*, vol. 19, pp. 163–172, 2004.

[21] R. Barland and A. Saadane, "Reference free quality metric for jpeg-2000 compressed images," in *Proc. Int. Symp. Signal Processing and Its Applications*, 2005, vol. 1, pp. 351–354.

[22] S. H. Oguz, Y. H. Hu, and T. Nguyen, "Image coding ringing artifact reduction using morphological post-filtering," in *Proc. IEEE 2nd Workshop on Multimedia Signal Processing*, 1998, pp. 628–633.

[23] X. Feng and J. Allebach, "Measurement of ringing artifacts in jpeg images," in *Proc. SPIE*, 2006, vol. 6076, pp. 74–83.

[24] J. Canny, "A computational approach to edge detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 8, no. 6, pp. 679–698, 1986.

[25] R. J. Sternberg, *Cognitive Psychology*, 4th ed. Florence, KY: Wadsworth, 2006.

[26] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *Proc. IEEE Int. Conf. Computer Vision*, 1998, pp. 836–846.

[27] H. Liu, N. Klomp, and I. Heynderickx, "Perceptually relevant ringing region detection method," in *Proc. 16th Eur. Signal Processing Conf.*, 2008.

[28] S. Winkler, "Vision models and quality metrics for image processing applications," Ph.D. dissertation, Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland, 2000.

[29] T. N. Pappas and R. J. Safranek, *Perceptual Criteria for Image Quality Evaluation. Handbook of Image and Video Processing*. New York: Academic, 2000.

[30] R. Franzen, Kodak Lossless True Color Image Suite [Online]. Available: http://www.r0k.us/graphics/kodak/ available

[31] I.-R. R. , *BT.500-11, Methodology for the Subjective Assessment of the Quality of Television Pictures*. Geneva, Switzerland: Int. Telecommun. Union, 2002.

[32] C. Liu, R. Szeliski, S. B. Kang, C. L. Zitnick, and W. T. Freeman, "Automatic estimation and removal of noise from a single image," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 30, no. 2, pp. 299–314, Feb. 2008.

[33] S. Paris and F. Durand, "A fast approximation of the bilateral filter using a signal processing approach," *Int. J. Comput. Vis.*, vol. 81, pp. 24–52, 2009.

[34] S. H. Oguz, "Morphological post-filtering of ringing and lost data concealment in generalized lapped orthogonal transform based image and video coding," Ph.D. dissertation, Univ. Wisconsin, Madison, WI, 1999.

Hantao Liu (S'07) was born in Chengdu, China, in 1981. He received the M.Sc. degree in signal processing and communications from the University of Edinburgh, Edinburgh, U.K., in 2005. Since 2006, he has been a Ph.D. student at the Department of Mediamatics, Delft University of Technology, Delft, The Netherlands.

He is currently working on a research project supported by the Philips Research Laboratories, Eindhoven, The Netherlands, developing no-reference objective metrics for perceived artifacts in compressed images. His research interests include image analysis, visual perception, and signal processing.

Nick Klomp received the B.Sc. and M.Sc. degrees from the Delft University of Technology, Delft, The Netherlands, in 2006 and 2008, respectively. From 2007 to 2008, he was working on his Master's thesis at the Department of Mediamatics, Delft University of Technology, developing a no-reference ringing artifact metric for real-time video applications.

Ingrid Heynderickx received the Ph.D. degree in physics from the University of Antwerp, Antwerp, Belgium in December 1986.

In 1987 she joined the Philips Research Laboratories, Eindhoven, The Netherlands, and meanwhile worked in different areas of research: optical design of displays, processing of liquid crystalline polymers, and functionality of personal care devices. Since 1999 she has been the Head of the research activities on Visual Perception of Display and Lighting Systems, and in 2005 she was appointed a Research Fellow in the group Visual Experiences.

Dr. Heynderickx is a member of the Society for Information Displays (SID), and for the SID, she was Chairman of the Applied Vision subcommittee from 2002 until 2007. In 2008, she became a Fellow of the SID and Chairman of the European Committee of the SID. In 2005, she was appointed Guest Research Professor at the Southeast University of Nanjing (China) and Part-Time Full Professor at the University of Technology in Delft (The Netherlands).