Pre-attentive considerations for gaze-contingent image processing

Andrew T. Duchowski
Bruce H. McCormick

Department of Computer Science
Texas A&M University
College Station, TX, 77843-3112

ABSTRACT

Gaze-contingent (GC) interfaces must provide adequate peripheral information to the viewer to preserve perceptual fidelity. Specifically, it has been shown that pre-attentive (or preview) benefit must be preserved so that scene perception is not disrupted.[14,8] In this paper we review recent attempts at peripheral degradation of digital imagery. We suggest that to be successful such degradation must preserve potential future visual attractors and, furthermore, not introduce artificial ones. We present a simple multiresolution image processing approach which can be utilized for this purpose.

The feasibility of gaze-contingent processing has recently been questioned.[17] In their paper, Stelmach and Tam processed images by low pass filtering, effectively smoothing extrafoveal regions. The authors then quantized DCT coefficients in the periphery, introducing blocking artifacts. In our paper, we simulate these results and claim that neither of these methods is suitable for GC interfaces. Alternatively, we implement a simple multiple Region Of Interest (ROI) multiresolution scheme in an attempt to degrade the periphery while preserving attentional cues. We evaluate three variants of this approach: a linear degradation function, a nonlinear function, and a function matching Human Visual System (HVS) acuity.

The HVS-matching multiple-ROI algorithm gives good compression and alleviates the degradation of potential visual attractors. Furthermore, MIP mapping offers efficient implementation of the algorithm,[2,5] making it a good candidate for gaze-contingent applications.

Keywords: Gaze-contingent Display, Multiresolution, ROI, Texture Mapping, Image Processing.

1 INTRODUCTION

Interest in gaze-contingent (GC) interface techniques has endured since early implementations of eye-slaved flight simulators. Although often difficult and costly to implement, functional benefits for human-computer interfaces, and technical benefits for data compression, have clearly been recognized.[13]

In most GC implementations (such as flight simulators), emphasis has usually been placed on representing the

foveal Region Of Interest (ROI), [1] while homogeneously degrading the periphery.[11,12] In the Super Cockpit Visual World Subsystem, Kocian considered visual factors including contrast, resolution and color in the design of a GC (head-tracked) display. At the time, no pre-attentional factors were mentioned. In their Simulator Complexity Testbed (SCTB), Longridge et al included an eye-slaved ROI as a major component of the Helmet Mounted Fiber Optic Display (HMFOD). This ROI provided a high resolution inset in a low resolution field which followed the user's point of regard. The precise method of peripheral degradation was not described apart from the criteria of low resolution. However, the authors did point out that a smooth transition between the ROI and background was necessary in order to circumvent the possibility of a perceptually disruptive edge artifact.

Interest in gaze-contingent technology has permeated to various domains including human-computer interfaces, teleoperator environments, and visual communication modalities.[10,16,7] Unlike flight simulators, which deal primarily with graphical data, these latter applications often require processing of digital imagery including video. Recently, the commercial feasibility of GC video processing has been questioned.[17] In their paper, Stelmach and Tam processed images by low pass filtering, effectively smoothing extrafoveal regions. The authors then quantized DCT coefficients in the periphery, introducing blocking artifacts. The authors have concluded that due to the high cost of implementation and only modest perceptual quality benefits, GC processing is not suitable for general purpose image processing.

Although preserving a foveal ROI is central to GC technology, we believe the key to the effectiveness of GC display modalities lies in the proper processing of peripheral information, taking care to preserve pre-attentive preview. Since Posner's seminal work on the spotlight of attention[15] and Treisman's conjunctive search,[18] recent visual search models suggest that significant visual processing is carried out in the periphery.[20] Specifically, it has been shown that pre-attentive (or preview) benefit must be preserved so that scene perception is not disrupted.[14,8] In their study of normal reading, Pollatsek and Rayner showed that parafoveal preview extends up to 3 degrees outside the current foveal window. Although it is difficult to extend results of reading studies to scene perception, Henderson showed that eye movements were not disrupted only when objects about to be fixated next were present in the display. Based on these findings we suggest peripheral degradation must preserve potential future visual attractors and, furthermore, must not introduce artificial ones. To this end we recommend heterogeneous processing of the periphery including smooth degradation of resolution and encoding of multiple ROIs. Ideally, resolution degradation should match visual acuity of the Human Visual System (HVS), while multiple ROIs should be used to preserve peripheral visual attractors thereby providing preview benefit.

In this paper we simulate the blurring operation presented by Stelmach and Tam, by processing still images with either a 3×3 box filter or with a 3×3 Gaussian filter to produce smoothing effects. To simulate the authors' introduction of blocking artifacts, we utilize the Karhunen-Loève Transform (KLT) [2] to imitate the effects of DCT quantization. We claim that neither blurring nor quantization is suitable for GC interfaces since in the former, the presence of potential visual attractors such as edges is diminished, and in the latter, artificial visual attractors are introduced.

To remedy these detrimental effects, we first assess the suitability of substituting a spatially-varying Gaussian filter instead of the usual homogeneous smoothing operation. We also apply a post-KLT smoothing operation to reduce blocking artifacts. Neither of these amendments improves the perceptual quality of peripheral areas.

In contrast, we implement a simple multiresolution scheme based on the well known MIP texture mapping algorithm[19] in an attempt to degrade the periphery while preserving attentional cues. We evaluate three variants of the approach: a linear degradation function, a nonlinear function, and a function matching Human Visual System (HVS) acuity. Our technique also allows representation of multiple regions of interest. We compare the multiresolution procedure to the above blurring and quantizing methods and demonstrate that it better preserves visual fidelity.

[1] Also known as the Area or Interest (AOI).

[2] Also known as the *eigenvector*, *principal component*, or the *Hotelling* transform.[6]

2 IMAGE PROCESSING

To simulate smoothing effects, box (averaging) and Gaussian filters were used, while a lossy version of the Karhunen-Loève Transform (KLT) was implemented to simulate blocking artifacts. These filters are termed conventional here since they can be found in most image processing textbooks.[6]

The multiresolution algorithm presented in this paper is a straightforward adaptation of the MIP texture mapping algorithm used extensively in computer graphics.[19] Our algorithm preserves the original image resolution within multiple user-selected ROIs and gradually degrades the periphery. In this paper, we cover three degradation functions offered by our algorithm, but details of the implementation are described elsewhere.[3]

2.1 Smoothing

Smoothing is accomplished by two separate methods: application of an averaging box filter and a Gaussian filter. In both cases, a 3×3 filter mask was used to convolve the periphery of the image. In the case of the box filter, the smoothed image, $g(x, y)$, is obtained using the relation

$$g(x,y) = \frac{1}{M} \sum_{(n,m)\in S} f(n,m)$$

for $x, y = 0, 1, \ldots, N-1$ excluding the rectangular ROI. S is the 3×3 box filter mask region and M is 9. In the case of Gaussian smoothing, the Gaussian filter, $h(i, j)$, is used, with $\sigma = 1$:

$$h(i,j) = \frac{1}{2\pi\sigma^2} e^{-\frac{i^2+j^2}{2\sigma^2}}$$

2.2 Blocking artifacts

Blocking artifacts are visible boundaries between reconstructed image subblocks when the image is processed by some blocking transform, such as the Discrete Cosine Transform (DCT). Although the KLT differs from the DCT, if the KLT is applied to image subblocks, using a lossy reconstruction, similar blocking artifacts will appear. The KLT used here processes the image by 4×4 subblocks. Each subblock is arranged by rows into a 16-dimensional vector, x, and a 16×16 covariance matrix, C, is obtained for each vector. After obtaining the corresponding 16 eigenvectors of C, and arranging them as rows of a 16×16 transformation matrix, A, the KLT then consists of multiplying a centralized image vector, $(x - m_x)$, where m_x is the mean vector, by A to obtain the new image vector, $y = A(x - m_x)$. If all eigenvectors of A are used in the transformation, the original image will be reconstructed losslessly. In our implementation we chose to use all 16 eigenvectors within the ROI, and only 4 of the 16 eigenvectors in the periphery. As expected, the resultant image exhibits blocking artifacts in the periphery while the ROI is left intact.

2.3 Spatially-varying smoothing

Spatially-varying Gaussian smoothing was used in an attempt to relax the smoothing effects of the homogeneous Gaussian by gradually diminishing resolution away from the primary ROI. Using the above Gaussian filter, we let σ range linearly with pixel distance from almost 0 at the primary ROI center to 8 at the image boundaries. That is, the pixel distance from the primary ROI center to the image borders is mapped to a linear range of $(0, 8]$.

2.4 Smoothing as a post-processing step

In an attempt to soften the harsh blocking effects of the KLT, a 3×3 box filter was used to smooth the periphery in a post-processing step following the KLT. That is, a pixel averaging operation was performed after KLT reconstruction. The same lossy KLT was used as above, with the primary ROI left unfiltered.

2.5 MIP mapping

In essence, MIP mapping decimates the original image, $f(x, y)$, into k "compressed" subimages, each of different resolution. Each subimage, $g_i(x, y)$, contains a quarter as much resolution information as the previous subimage, $g_{i-1}(x, y)$. The subimage $g_0(x, y)$ is simply a copy of the original image $f(x, y)$, while the subimage $g_k(x, y)$ is the average of $f(x, y)$ compressed into a single pixel. This multiresolution set of subimages comprises a representation of the original image in *resolution space*. Notice that in this resolution space, fine-to-coarse resolution information is distributed nonlinearly (by decreasing powers of 2).

To reconstruct the image, the intensity of each pixel at spatial location (x, y) is calculated as a combination of pixel intensities in resolution space. The reconstructed image is segmented into k resolution bands, matching the number of compressed subimages. At each band boundary, i, pixel information is obtained directly from subimage $g_i(x, y)$. Between band boundaries i and $i+1$, pixel information is linearly interpolated between subimages $g_i(x, y)$ and $g_{i+1}(x, y)$. Reconstruction as described here is carried out in *image space*.

Reconstruction depends on a *mapping function* between image and resolution spaces. The choice of mapping function is crucial since it is responsible for segmentation and distribution of the k resolution bands in image space. In this implementation we chose $k = 8$ and circular representations of ROIs. Degradation of the periphery is achieved by one of three mapping functions: a linear function, a non-linear function, or a function matching the acuity of the human visual system (HVS). The linear and nonlinear functions were chosen as approximate lower and upper bounds, respectively, on the empirical HVS-matching function. Each mapping function segments the image into concentric resolution regions. In all three mapping variants, resolution within the central 5° of each ROI is consistent. Although this is not a restriction imposed by MIP mapping, we chose to maintain consistent resolution within the foveal region so that different peripheral degradation methods could be examined.

The algorithm permits the existence of several ROIs within the image. The purpose of our algorithm is to simply encode the primary ROI subtended by the fovea and allow peripheral ROIs (termed here potential, or peripheral, Regions Of Interest, or pROIs) to preserve potential visual attractors. Note that the algorithm makes no attempt to detect these attractors. Other means, such as pattern matchers, interest operators, or neural networks, should be employed for this purpose.

2.5.1 Linear mapping

Using a linear mapping function, the k resolution bands within the image are spaced equally apart. The linear mapping function is similar to the function used in the spatially-varying Gaussian smoothing approach above. That is, pixel distance is mapped into the linear range $[0, 8]$. The central region of the ROI is considered level 0, where the original image resolution is used to reconstruct the pixels within this region. At level 8, the worst resolution, which is simply the average of the entire original image, is used for reconstruction. To maintain resolution within the 5° foveal region the linear function is defined as $l = \frac{d}{R}$, where d is pixel distance and R is the radius of the highest resolution region. Figure 1 (a,b) shows the relative (percent) resolution plot, and the corresponding linear mapping function, l. Percent resolution is obtained by taking the inverse of the constant 2 raised to the mapping function, corresponding to the manner in which the image is originally decimated.

2.5.2 Nonlinear mapping

Nonlinear mapping segments the image into k resolution levels by the relation $l = A(1 - e^{-\lambda \frac{d}{R}})$, where d is the pixel distance from the ROI center and R is the radius of the central 5° ROI. A is the asymptote approximated at the image boundary. To act as an upper bound for the HVS-matching function A was chosen as 2.35, and R was set to 105 to match the radius of the area of highest-resolution (see below). The plots of the relative resolution and the nonlinear mapping functions are shown in Figure 1 (c,d).

2.5.3 HVS acuity-matching mapping

The HVS acuity-matching mapping was derived from empirical MAR (minimum angle of resolution) data.[4] Since it is known that acuity at 5° eccentricity is roughly 50% of the acuity at the fovea,[9] resolution at the fovea center is assumed to be 100%, that is, 66 dots per inch (maximum resolution). Assuming a 60cm viewing distance, acuity as a function of eccentricity was calculated and is shown in Table 1. Percent resolution was approximated

Eccentricity (deg.)	MAR (at 75cm)	Min. Separability (inches, at 60cm)	Max. Resolution (dots per inch)
5	.14	.06	33
10	.28	.12	17
15	.32	.13	15
20	.38	.16	13
25	.45	.19	11

Table 1: Resolution as function of eccentricity at 60cm viewing distance

by treating maximum resolution from Table 1 as relative resolution, i.e., 33 dpi maximum resolution ≈ 50%, 17 maximum resolution ≈ 26%, etc. Using percent resolution and assuming a screen display resolution of 100 dots per inch, resolution level distances used in the MIP mapping algorithm were calculated and are shown in Table 2. Resolution level distances were used to determine diameters of decreasing resolution areas. The highest

Eccentricity	0-5°	5°	10°	15°	20°	25°
Resolution	100%	50%	26%	23%	20%	17%
Diameter subtended						
(cm)	–	5.2	10.4	15.8	21.2	26.7
(in)	–	2.1	4.1	6.2	8.3	10.5
(pixels, assuming 100 dpi)	–	210	410	620	830	1050

Table 2: Resolution levels (in pixels)

resolution region within each ROI, for example, is a circular region with a diameter of 210 pixels. Since resolution is distributed by decreasing powers of two, desired relative resolution at eccentricities were mapped into resolution space using the function for the band level, $l = -\frac{\ln(\%\ resolution)}{\ln 2}$. Percent resolution between bands was linearly interpolated prior to the log mapping. The relative resolution function is shown in Figure 1 (e). All three resolution functions are depicted in Figure 1 (f).

2.5.4 Multiple ROIs

Finally, the MIP mapping algorithm was extended to allow multiple ROIs within the image. The purpose of this extension is to allow encoding of distant visual attractors (from the primary ROI) at high resolution. This

effectively overrides the default low resolution assigned to these peripheral pixels in the case of a single ROI. Justification for these peripheral ROIs is two-fold: first, peripheral attractors should be preserved in order to maintain preview benefit; second, peripheral ROIs can be used as potential future regions of interest in anticipation of saccades to these areas.

All ROIs were processed consistently, with each ROI of equal size. ROI membership of each image pixel is determined by a simple distance criteria. That is, the resolution at each image pixel is obtained by the mapping function used relative to the spatially closest ROI. Membership of pixels where multiple ROIs are equidistant is arbitrarily assigned to one of the equidistant ROIs.

The relative resolution distribution in image space, for single and multiple ROIs, is shown in Figure 2. Light areas (white $\equiv$ level 0) are reconstructed with 100% original resolution, dark areas (black $\equiv$ level 8) are reconstructed with 0% (worst) resolution. Black rings are level boundaries. In images (a), (c), and (e), the region of interest is located in the bottom left of the picture. In images (b), (d), and (f), a secondary ROI is added in the upper right quadrant.

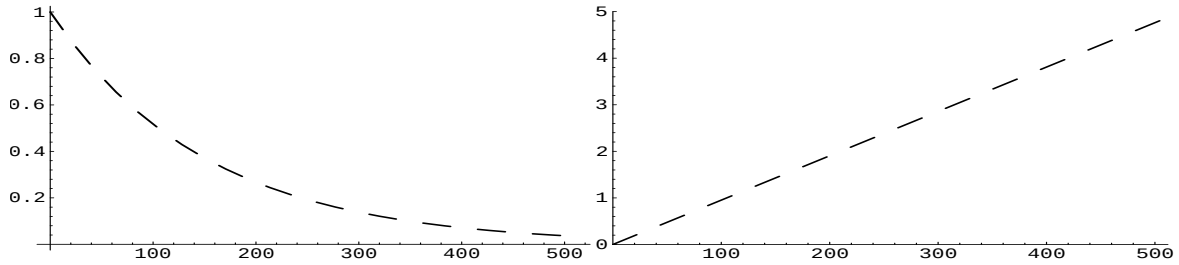

(a) Relative resolution: *resolution* vs. *pixel distance*

(b) Linear mapping: band level l vs. *pixel distance*

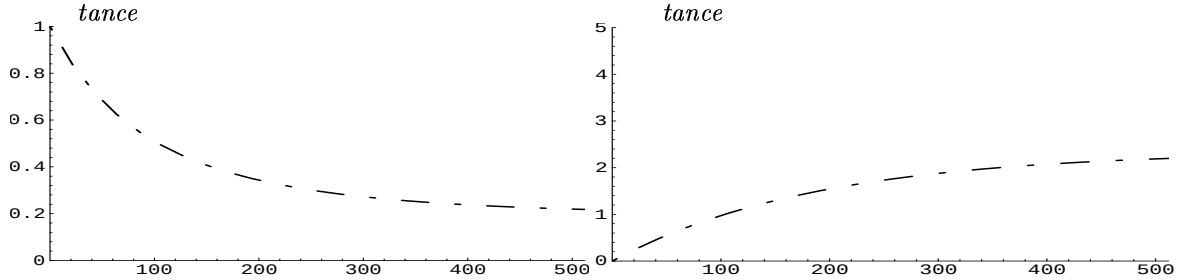

(c) Relative resolution: *resolution* vs. *pixel distance*

(d) Nonlinear mapping: band level l vs. *pixel distance*

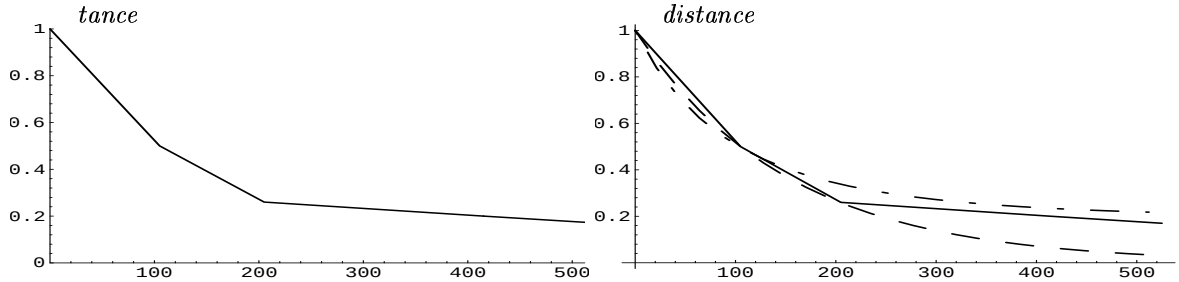

(e) Relative resolution: *resolution* vs. *pixel distance*

(f) Relative resolution: *resolution* vs. *pixel distance*

Figure 1: MIP mapping functions

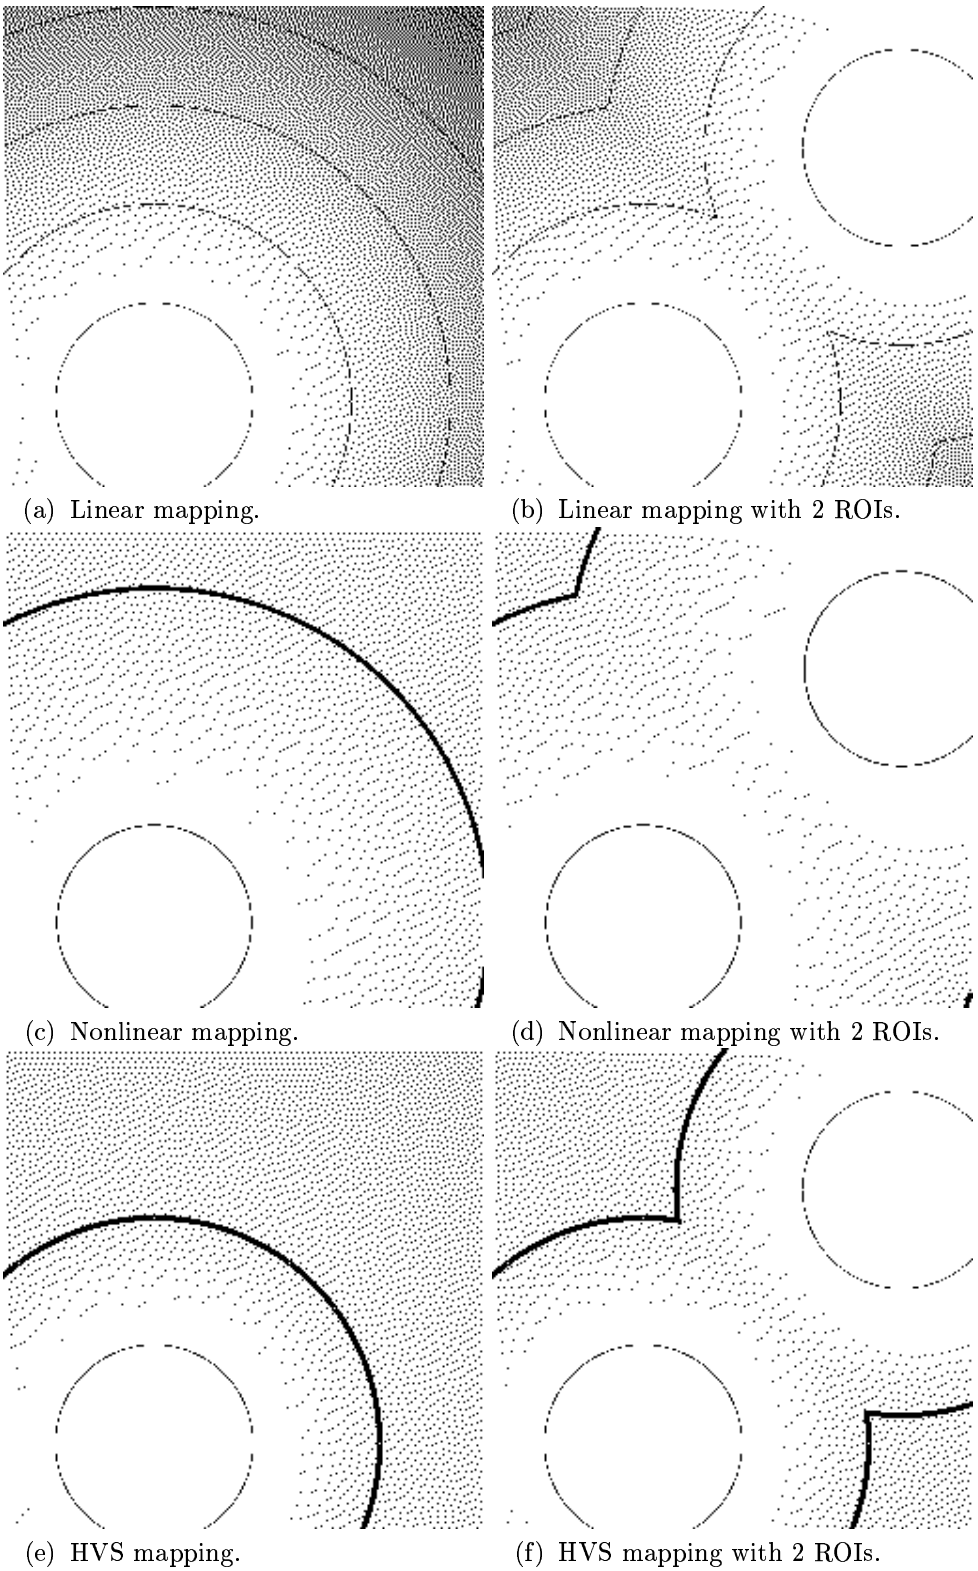

(a) Linear mapping.

(b) Linear mapping with 2 ROIs.

(c) Nonlinear mapping.

(d) Nonlinear mapping with 2 ROIs.

(e) HVS mapping.

(f) HVS mapping with 2 ROIs.

Figure 2: Relative (percent) resolution distribution.

3 RESULTS

Combinations of conventional filters were applied to the periphery of a 512×512 still, monochrome image leaving a 95×120 rectangular ROI unfiltered. The primary ROI in these images is centered on a large flower at the bottom left of the image. This ROI was chosen for the flower's size. A secondary ROI is centered on an elongated, tufty white flower in the upper right of the image. The tufty white flower is visually significant since it is slightly elongated with respect to the other plants making it somewhat of an oddity. Certain odd objects have been shown to attract visual attention.[1] Furthermore, it is in a relatively uncluttered dark image area, making it conspicuous. Lastly, it is slightly more illuminated in comparison to the rest of the flowers. For these reasons, it was chosen as a potential visual attractor and encoded as a secondary ROI. The secondary ROI is preserved only by the multiple-ROI multiresolution scheme.

The original image and results of processing by conventional approaches are shown in Figure 3. Notice that all the smoothing algorithms tend to eliminate the fine detail of the smaller white flowers in the middle right portion of the image. As a result, the flowers tend to appear as white balls. Also, the white tufty flower in the top right quadrant is deemphasized. The spatially-varying Gaussian approach shows no visible improvement over the homogeneous Gaussian.

Unlike smoothing, the KLT introduces blocking artifacts in the image. This is particularly noticeable in the round, white flowers in the middle of the image. These flowers tend to take on a square appearance under the KLT. The squareness of these flowers introduces edges which may compete for visual attention with the tufty white flower in the upper right. Smoothing after the KLT does not offer much help. The resultant image appears worse than either of the images produced by the Gaussian or the KLT alone. Too much original information has already been lost, and too many artificial artifacts have been introduced, for any hope of correction by post-process smoothing.

Figure 4 shows the results obtained by the multiresolution scheme with pre-attentive consideration and multiple ROIs. Notice that both linear implementations of the MIP mapping algorithm extremely degrade the periphery. The degradation by a linear mapping of fine-to-coarse resolution is too severe for perceptual image processing.

The nonlinear and HVS-matching mappings offer some hope in terms of both preserving visual fidelity and offering good compression potential. Nonlinear mapping produces the best visual results, but gives the worst compression potential. The HVS-matching function distributes resolution levels according to empirical data of human visual acuity. Resolution degradation becomes visible closer to ROI centers than with nonlinear mapping. This suggests a higher compression potential, although objects further away from the ROI may be degraded too severely to offer pre-attentional preview benefits. This is especially dangerous for larger images. For this reason, the obvious approach of allowing multiple ROIs was taken in the current MIP mapping implementation.

Images (d) and (f) in Figure 4 contain 2 ROIs, with the periphery around each ROI degraded by nonlinear and HVS-matching degradation functions, respectively. Since allowing an extra ROI into the image encodes more of the total image area at a higher resolution, the multiple-ROI approach reduces the total compression potential. The nonlinear multiple-ROI scheme shows good visual results, but little compression potential is gained. The HVS-matching multiple-ROI approach gives better compression potential, and alleviates the problem of degraded potential visual attractors.

(a) Original image.

(b) Periphery smoothed by 3×3 box filter.

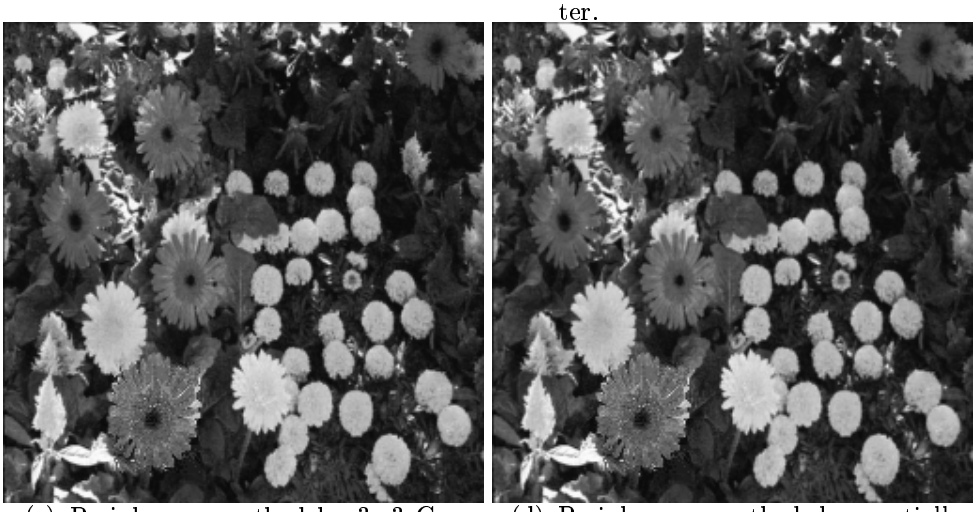

(c) Periphery smoothed by 3×3 Gaussian.

(d) Periphery smoothed by spatially-varying 3×3 Gaussian.

(e) Periphery degraded by lossy KLT.

(f) Periphery degraded by lossy KLT then smoothed by 3×3 box filter.

Figure 3: *Floral* images, processed by conventional image transforms

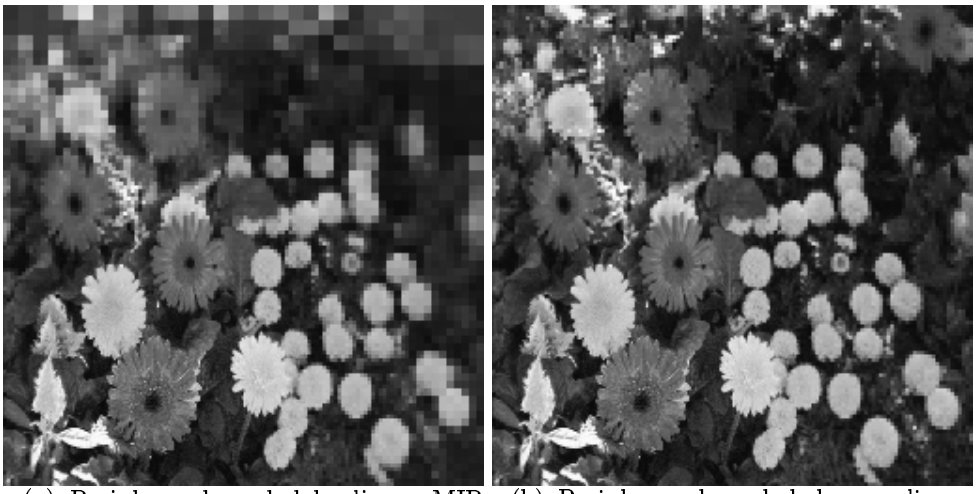

(a) Periphery degraded by linear MIP mapping.

(b) Periphery degraded by nonlinear MIP mapping.

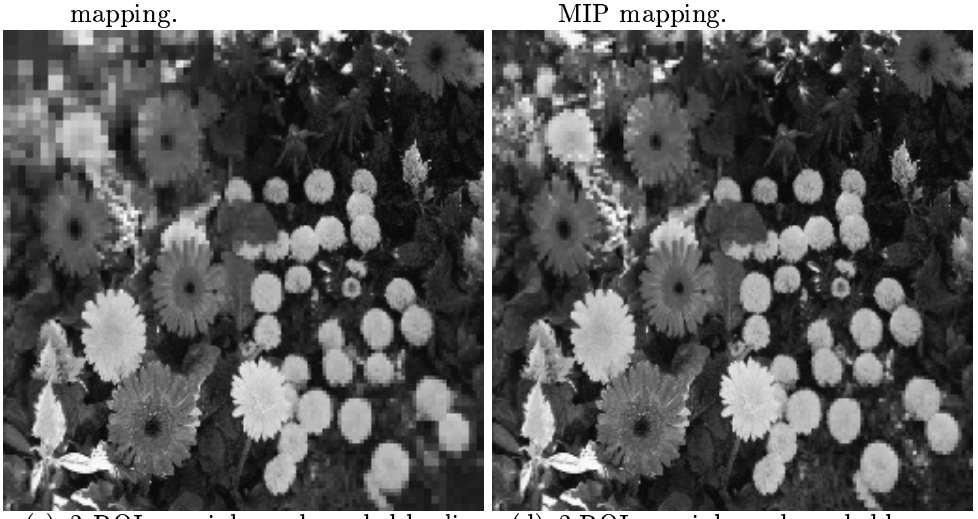

(c) 2 ROIs, periphery degraded by linear MIP mapping.

(d) 2 ROIs, periphery degraded by nonlinear MIP mapping.

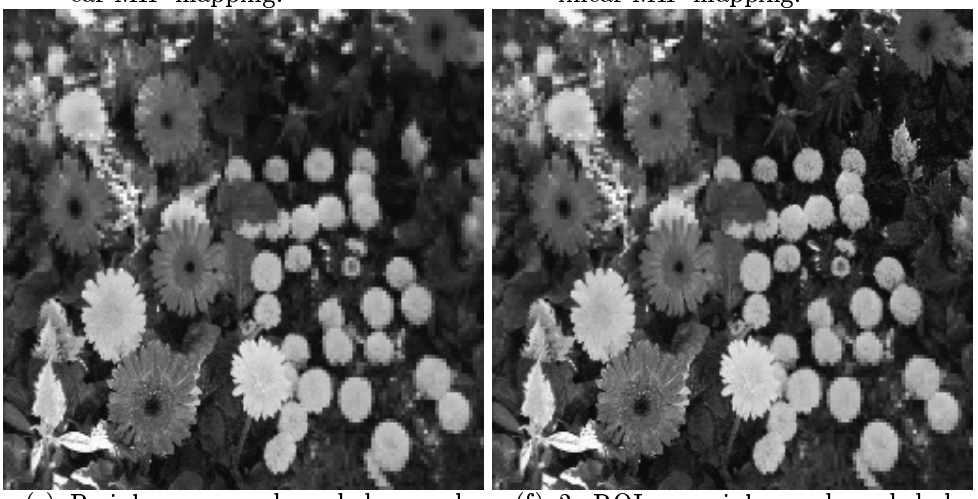

(e) Periphery degraded by HVS-matching MIP mapping.

(f) 2 ROIs, periphery degraded by HVS-matching MIP mapping.

Figure 4: *Floral* images, processed by multiresolution algorithm

4 DISCUSSION

The goal of this research was to explore various methods of peripheral degradation and to find candidates suitable for pre-attentive image processing in gaze-contingent applications. The spatially-varying smoothing filter used here showed no obvious gain over homogeneous smoothing in improving perceptual fidelity. Smoothing filters such as the Gaussian are not suitable for perceptual image degradation. Smoothing filters eliminate too many potential visual attractors. In the case of images with large high frequency content, such as the one used here, too much detail is lost. Homogeneous quantization appears to introduce too many artifacts in the periphery.

The advantage of multiresolution is the ease in which the scheme lends itself to control of peripheral degradation simply by providing the choice of an appropriate mapping function. If greater resolution is required, ROIs can be enlarged, or a different mapping function can be substituted. In this implementation, potential Regions of Interest (pROIs) are chosen with the same spatial dimensions as the primary ROI. An alternative approach may consider form-varying pROIs to match the characteristics of pre-attentive attractors. The form of individual pROIs can be expected to depend on various factors including general properties of the underlying attractor (size, shape, etc.), viewing distance, and the viewer.

Lacking eye-tracking equipment, we were unable to evaluate our results in terms of subjective image quality. Nevertheless, we are confident gaze-contingent techniques will achieve better perceptive quality if the image processing methods used to degrade the periphery are implemented with consideration for parafoveal preview. Particularly, we feel that current monoresolution blurring and quantizing procedures are inappropriate due to their inherent decimation of potential attractors and their homogeneous treatment of the periphery. Encoding potential regions of interest may not only provide preview benefit, but may also improve system response to anticipated eye movements.

5 CONCLUSION

A simple image processing scheme is used to represent multiple regions of interest in still images, while providing gradual degradation of resolution in peripheral areas. Gaze-contingent processing must preserve visual attractors, which smoothing tends to eliminate. Conversely, blocking artifacts introduce false attractors. The approach presented here allows graceful and selective preservation of appropriate peripheral information. The variant of our algorithm matching HVS acuity alleviates both the problems of attenuation of authentic potential visual attractors, as well as the creation of synthetic ones. Furthermore, the algorithm allows good compression and efficient implementation, making it a good candidate for gaze-contingent applications.

6 REFERENCES

[1] Susan J. Boyce and Alexander Pollatsek. An Exploration of the Effects of Scene Context on Object Identification. In Keith Rayner, editor, *Eye Movements and Visual Cognition: Scene Perception and Reading*, pages 227–242. Springer-Verlag, 1992. Springer Series in Neuropsychology.

[2] Franklin C. Crow. Summed-area tables for texture mapping. In *Computer Graphics*, volume 18, pages 207–212. SIGGRAPH, 1984.

[3] Andrew T. Duchowski and Bruce H. McCormick. Simple Multiresolution Approach for Representing Multiple Regions of Interest (ROIs). In *SPIE Conference on Visual Communications and Image Processing*, Taiwan, 1995. SPIE. (Submitted for publication.).

[4] David H. Foster, Salvatore Gravano, and Antonia Tomoszek. Acuity for Fine-Grain Motion and For Two-Dot

Spacing as a Function of Retinal Eccentricity: Differences in Specialization of the Central and Peripheral Retina. *Vision Research*, 29(8):1017–1031, 1989.

[5] Alain Fournier and Eugene Fiume. Constant-time filtering with space-variant kernels. In *Computer Graphics*, volume 22, pages 229–238. SIGGRAPH, 1988.

[6] Rafael C. Gonzalez and Paul Wintz. *Digital Image Processing.* Addison-Wesley Publishing Company, 1987. Second Edition.

[7] Richard Held and Nathaniel Durlach. *Pictorial Communication in Virtual and Real Environments*, chapter 14: Telepresence, time delay and adaptation, pages 232–246. Taylor & Francis, Ltd., London, second edition edition, 1993.

[8] John M. Henderson. Visual Attention and Eye Movement Control During Reading and Picture Viewing. In Keith Rayner, editor, *Eye Movements and Visual Cognition: Scene Perception and Reading*, pages 260–283. Springer-Verlag, 1992. Springer Series in Neuropsychology.

[9] David E. Irwin. Visual Memory Within and Across Fixations. In Keith Rayner, editor, *Eye Movements and Visual Cognition: Scene Perception and Reading*, pages 146–165. Springer-Verlag, 1992. Springer Series in Neuropsychology.

[10] Robert J. Jacob. What You Look at is What You Get: Eye Movement-Based Interaction Techniques. In *CHI '90 Conference Proceedings*, pages 11–18. ACM Press, 1990.

[11] Dean Kocian. Visual world subsystem. In *Super Cockpit Industry Days: Super Cockpit/Virtual Crew Systems*, Air Force Museum, Wright-Patterson AFB, Ohio, 31 March – 1 April 1987. Air Force Systems Command/Human Systems Division/Armstrong Aerospace Medical Research Laboratory.

[12] Thomas Longridge, Mel Thomas, Andrew Fernie, Terry Williams, and Paul Wetzel. Design of an Eye Slaved Area of Interest System for the Simulator Complexity Testbed. In *Interservice/Industry Training Systems Conference*, pages 275–283. National Security Industrial Association, 1989.

[13] Jakob Nielsen. The Next Generation GUIs: Noncommand User Interfaces. *Communications of the ACM*, 36(4):83–99, April 1993.

[14] Alexander Pollatsek and Keith Rayner. What Is Integrated Across Fixations? In Keith Rayner, editor, *Eye Movements and Visual Cognition: Scene Perception and Reading*, pages 166–191. Springer-Verlag, 1992. Springer Series in Neuropsychology.

[15] Michael I. Posner. Orienting of Attention. *Experimental Psychology: Quarterly*, 32:3–25, 1980.

[16] India Starker and Richard A. Bolt. A Gaze-Responsive Self-Disclosing Display. In *CHI '90 Conference Proceedings*, pages 3–9. ACM Press, 1990.

[17] Lew B. Stelmach and Wa James Tam. Processing Image Sequences Based on Eye Movements. In *SPIE Conference on Human Vision, Visual Processing, and Digital Display V*, pages 90–98, San Jose, CA, February 8-10 1994. SPIE.

[18] Anne Treisman. Perceptual Grouping and Attention in Visual Search for Features and for Objects. *Experimental Psychology*, 8(2):194–214, 1982.

[19] Alan Watt and Mark Watt. *Advanced Animation and Rendering Techniques.* Addison-Wesley, Reading, MA, 1992.

[20] Jeremy M. Wolfe. Guided Search 2.0: The Upgrade. In *Proceedings of the Human Factors and Ergonomics Society, 37th Annual Meeting*, pages 1295–1299, October 1993.