A Wavelet Approach to Foveating Images

(Extended Abstract)

Ee-Chien Chang and Chee K. Yap*
Courant Institute of Mathematical Sciences
New York University
New York, NY 10012
U.S.A.

May 14, 1997

Abstract

Motivated by applications of foveated images in visualization, we introduce the **foveation transform** of an image. We study the basic properties of these transforms using the multiresolution framework of Mallat. We also consider practical methods of realizing such transforms. In particular, we introduce a new method for foveating images based on wavelets. Preliminary experimental results are shown.

1 Introduction

Conventional images have uniform resolution. **Foveated images** which have non-uniform resolution arise in biological vision. In figure 1(a) and (b) we show a uniform image and a foveated version of the same image. The process of going from (a) to (b) is called "foveating" the image (a).

One of the most interesting forms of foveated images is based on the complex logarithm function. Such **logmap images** were studied by Rojer and Schwartz [18] and others. The complex logmap is a model consistent with empirical data on the mapping from primate retina to the visual cortex [19, 20]. This neuro-physiological model goes back to the pioneering (and the Nobel prize) work of Hubel and Wiesel [6, 7]. Rojer [17] demonstrated many favorable properties of such images. Perhaps the most striking fact is that the data density in such images grows logarithmically with the diameter of the visual field (as opposed to quadratically in the case of uniform images). Such low density images have been exploited in applications such as realtime recognition of car-plates [3] and in video phones [23]. In this paper, we use the term "foveated" to refer to any variable resolution image, not just those with the geometry of the logmap or having a definite "foveal region".

Visualization Applications. Foveated images have been exploited in computer vision (e.g., [4, 21]), especially in the context of **active vision**[1, 2]. But they are also useful in **visualization**, although this aspect is less explored. For visualization, we must compensate for the loss of peripheral resolution in foveated images by rapidly presenting images that are foveated at points of user interest. Robinson [16] describes four principal eye movements in the human visualization. We could use eye-tracking technology [25, 10, 11] to follow the eye and to present the corresponding foveated images. In principal, if the foveation uses the correct parameters and eye-tracking is good, a user should not be able to detect any difference from viewing a uniform image. A system with such properties has been reported by [13]. But in our

*yap@cs.nyu.edu, supported by an National Science Foundation Grant #CCR-9402464.

Figure 1: (a) Original uniform image with 375x469 pixels (b) Foveation at woman's face

visualization applications, we can dispense with eye-tracking — the visualizer may be counted upon to cooperate with the image server. For instance, the visualizer can indicate moving areas of interest in an image via a standard mouse input. Because of the rapid eye activity in normal visualization, we call this approach **active visualization**.

Thinwire Visualization. There are many applications of active visualization techniques. We describe one such scenario, in "thin-wire" models of computation. More precisely, we have a **server** who acts as an image-server and a **client** who wishes visualize images from the server. The images may be very large (e.g., a terrain or map) and perhaps to high resolution. The server has large computational resources (the client need not) and there is only a thin wire (such as the Internet) connecting the server and client. Conventionally, all the bandwidth in such visualization is concentrated in the direction from server to client. But we can greatly reduce this bandwidth by allowing the client to specify areas of interest in the image, and for the server to send the corresponding foveated images. The price we pay is only a modest increase of bandwidth from the client to the server. A basic problem in this scenario is the construction foveated images from uniform images. Our paper focuses on this problem.

Overview of Results. In this paper, we describe a new approach for constructing foveated images using wavelet theory. Previous work on constructing foveated images on computers are based on the idea of averaging over superpixels. Our approach is more general and flexible. We describe a mathematical basis for this approach and obtain some bounds on some practical approaches. A java-based foveated image server demonstrating our approach is accessible from the URL `http://jinai.nyu.edu:9999/visual/demo.html`.

2 Previous Work

We mentioned the work of Eric Schwartz and his collaborators exploiting logmap images for computer vision. In these papers, the foveated image is obtained by dividing the visual field into regions called **superpixels** (or **logpixels** in this setting). See Figure 8 (appendix) for an arrangement of the logpixels. Each superpixel is given the average value of the signal over its region. Wallace et al [24] treats the more general formulation of this approach based on **space-variant geometry**. Their emphasis is the connectivity properties of such geometries, since connectivity is important for vision and image processing. In their applications, this geometry can be pre-computed and stored in a lookup table. But for our visualization applications, this is not flexible enough. Another direction is to produce foveated images from geometric models directly. Franzblau [8] described such a system for rendering foveated images.

Kortum and Geisler [13] implemented an efficient real-time system. They use superpixels based on a very simple rectangular geometry (see figure 7 in the appendix). They described psychological experiments in which subjects reported little perceptual difference between foveated and uniform images. They also noted the possibility of using foveated image to reduce bandwidth of transmission. Reeves and Robinson [22] gives a method to foveate MPEG-standard video in a thin-wire environment. MPEG-standard could provide a few levels of resolution. Here, an image is foveated with only two levels. The client/viewer can interactively specify the region of interest to the server/sender. Gross et al [9] discusses foveation on a mesh rather than image. They also use wavelets in a way that looks remarkably similar to ours, but not for foveation purposes.

3 Foveation Theory

There are various practical methods to foveate an image, as noted above. We would like to provide the mathematical formulation of this process. This would allow us to evaluate the various methods. Our approach is based on wavelet theory [5, 15] and in particular the multiresolution approach of Mallat and Meyer [14]. Koornwinder [12] can be a general reference for the background material in this paper.

As usual, one can begin with one-dimensional images or functions. Given an image $I(x)$, we would like to define a foveated image $I_0(x)$ that corresponds to the transformation from (a) to (b) in figure 1. More generally, we define an infinite sequence of foveated images $\{I_n(x) : n \in \mathbb{Z}\}$ at different resolution levels such that, in a suitable sense, $I_n(x) \to I(x)$ as $n \to \infty$. We consider mostly functions f, g in $L^2(\mathbb{R})$, which is viewed as a Hilbert space with inner product

$$\langle f, g\rangle := \int f(x)\overline{w(x)}dx$$

where the integration is over $\mathbb{R}$ unless otherwise specified. Also, $||f|| := \langle f, f\rangle^{-1/2}$.

A **foveation** is determined by a **weight function** and a **scaling function**,

$$w : \mathbb{R} \to \mathbb{R}_{\geq 0}, \quad s : \mathbb{R} \to \mathbb{R}.$$

We require the weight function to be positive except at finitely many values (so $w(x)^{-1} < \infty$ at all but finitely many x). The scaling function is analogous to the "father wavelet" usually denoted $\phi(x)$ in wavelet theory. We define

$$I_0(x) := \int I(t)C^{-1}(x)s\left(\frac{t-x}{w(x)}\right)dt,$$

where

$$C(x) := ||s\left(\frac{\cdot - x}{w(x)}\right)||.$$

The **weighted translation** (or, w-translation) of s by x is defined to be

$$s_x(t) := C^{-1}(x)s\left(\frac{t-x}{w(x)}\right). \tag{1}$$

Now the foveated image could be rewritten as

$$I_0(x) = \langle I, s_x\rangle.$$

Examples. An example of a scaling function is the Haar function

$$h(t) := \begin{cases} 1 & \text{if} \quad t \in [0,1], \\ 0 & \text{otherwise.} \end{cases}$$

The weighted translation of h is

$$h_x(t) = \begin{cases} w(x)^{-1} & \text{if} \quad t \in [x, x + w(x)], \\ 0 & \text{otherwise.} \end{cases}$$

A variant is the centered Haar function $h^*(t)$, defined to be $h^*(t) := h(t + 0.5)$. Figure 2 shows weighted translations of the Haar and centered Haar function under different weight functions.

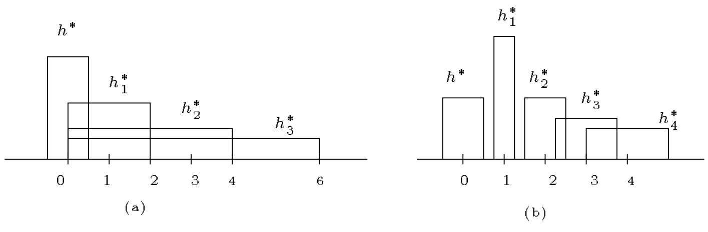

Figure 2: (a) and (b) illustrate centered Haar under weight functions $w(x) = 2|x|$ and $w(x) = |x|$, respectively. (c) illustrates Haar function under weight function $w(x) = 2|x|$.

A **standard weight function** is one of the form

$$w_{\text{std}}(t) = w_{\alpha,\beta,\gamma}(t) := \alpha|t - \gamma| + \beta,$$

where α, β, γ are fixed constants and $\alpha, \beta \geq 0$. We call α the **rate**, β the **foveal resolution** and γ the **gaze point**. Intuitively, this weight function defines the point $t = \gamma$ as the fovea with maximum resolution β. The resolution decreases linearly with the distance from the gaze point, at the rate α. Unless otherwise stated, $\gamma = 0$ in the following. When $\beta = 0$, the resolution at the gaze point is that of the original image. Figure 1(b) is a foveated image using a standard weight and the Haar scaling function.

Figure 3 shows the foveated image of the function $I(t) = \sin(t/10) + \sin(t/45)$ using $h(t)$ and $w_{\text{std}}(t)$ with different rates (but $\beta = 0$).

It is easy to conclude that if $I(x)$ is a constant function then $I_0(x) \equiv I(x)$. Next suppose the foveal resolution is $\beta = 1$ and rate $\alpha = 1$. If $I(x) = \sin(x)$ then

$$\begin{aligned} I_0(x) &= \int \sin(t) h_x(t) dt \\ &= (|x| + 1)^{-1}(\cos(x) - \cos(2x + 1)). \end{aligned}$$

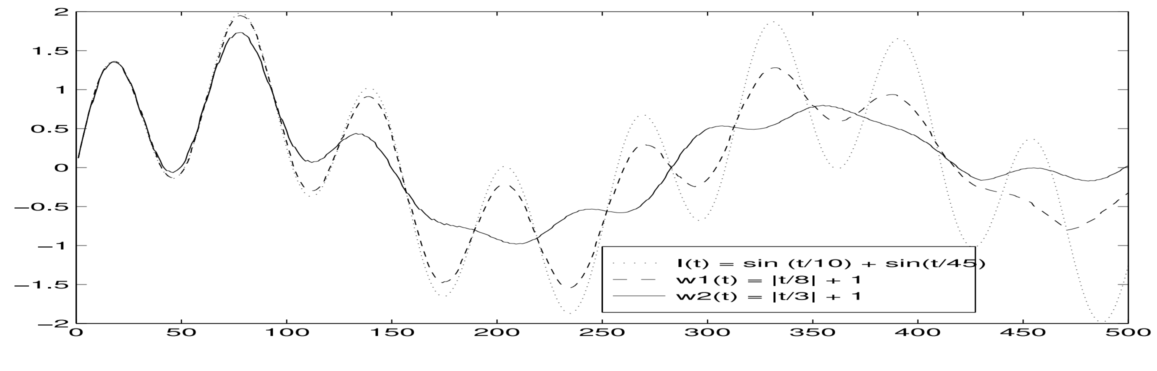

Figure 3: Foveation with standard weight

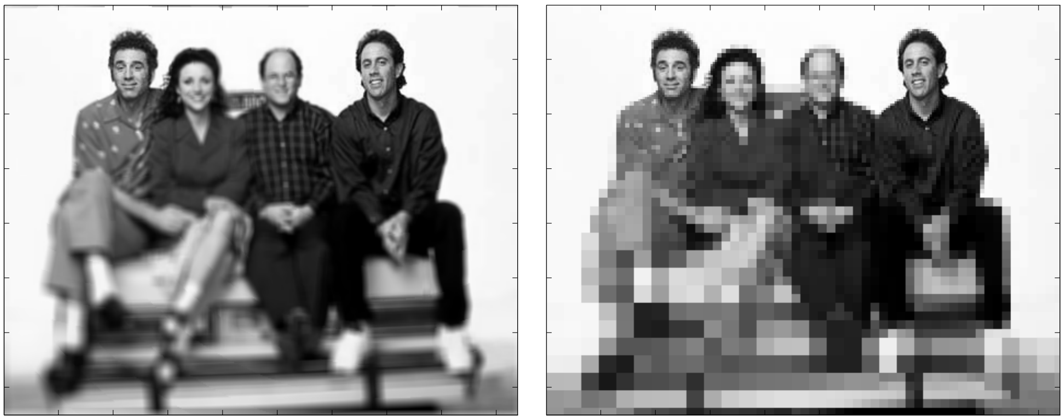

Figure 4: Blended image with 2 gaze points (on the extreme left and extreme right faces with rates $\frac{1}{20}$ and $\frac{1}{30}$ respectively). (a) The left image is obtained using the definition of foveation. (b) The right image is obtained using wavelet foveation.

Blended Weights. We can obtain new weight functions by combining several given weight functions. Rojer [17] studies the problem of combining several foveated images of a common image into a new foveated image. This is called "blending of foveated images". The following definition corresponds to an effective method for blending images: given weight functions w_1 and w_2, their **blended weight function** w is the pointwise minimum of the two functions: $w(x) = \min\{w_1(x), w_2(x)\}$. Of course this definition extends to the blending of any finite number of images. The weight function illustrated in figure 4 is the blending of two standard weights with different gaze points (suitably generalized to 2 dimensional images).

4 Weighted Distance

For a function I', we want to measure how far I' is from the foveation I_0. It is not useful to use the usual $||I' - I_0||$ norm as we want to place less emphasis on $I'(x) - I_0(x)$ when x is far away from the "foveal region".

Let us remark (and should be no surprise) that $I \in L(\mathbb{R})$ does not imply $I_0 \in L(\mathbb{R})$. For instance, consider $\beta = 0$, $\alpha = 2$ in the standard weight function and the centered Haar as the scaling function. Let $I(x) \equiv h(x)$. Then $I_0(x) = w(x)^{-1}$, for $x > 1$. Since $w(x) = 2|x|$, $||I_0|| = \infty$.

Define for a function f,

$$||f||_w = \left(\int (w(t))^{-1} \, |f(t)|^2 \, dt\right)^{1/2}. \tag{2}$$

Then the w-**distance** between f and g is defined to be $d_w(f,g) := ||f - g||_w$. Let $L(t)$ be a function such that $dL(t) = \frac{dt}{w(t)}$. If $L^{-1}(t)$ is well-defined then

$$\begin{aligned} \int |f(t)|^2 (w(t))^{-1} dt &= \int |f(t)|^2 dL(t) \\ &= \int |f(L^{-1}(y))|^2 dy \end{aligned} \tag{3}$$

Let us compute $L(t)$ and $L^{-1}(t)$ (if these exist) for the standard weight functions:

Lemma 1 *Let $w(t) = \alpha|t| + p$.*
(i) Then

$$\int \frac{dt}{w(t)} = L(t) = \mathrm{sg}(t) \cdot \frac{\ln(\alpha|t| + \beta)}{\alpha}$$

where $\mathrm{sg}(t)$ is the sign function.
(ii) $L(t)$ is defined for all t. It is continuous and monotonic increasing at every $t \neq 0$.
(iii) $L(0) = 0$ and $L(0^+) = \ln(\beta)/\alpha = -L(0^-)$. Hence, $L(t)$ is continuous at $t = 0$ iff $\beta = 1$.
(iv) Define

$$E(t) := \mathrm{sg}(t)\frac{e^{\alpha|t|} - \beta}{\alpha}.$$

Then $L(E(t)) = E(L(t))$ whenever both sides are defined. In particular, if $\beta = 1$, then E and L are inverses of each other.

When $\beta = 1$, we give $L(t)$ and $E(t)$ a special notation:

$$\mathrm{Log}_\alpha(t) := \mathrm{sg}(t) \cdot \ln(\alpha|t| + 1), \qquad \mathrm{Exp}_\alpha(t) := \mathrm{sg}(t)(e^{\alpha|t|} - 1).$$

Notice that, unlike the usual logarithm which is defined only for positive numbers, $\mathrm{Log}_\alpha(t)$ is defined and continuous for all real t.

Note that the weighted distance of a function I from its foveation could be unbounded. For example, take $I(t) = \sin(t)$, standard function $w(t) = 2|t|$, and Haar as the scaling function.

5 Multiresolution and Discrete Foveation

We can now describe a multiresolution framework that is analogous to Mallat and Meyer's analysis. For any weight function w and $n \in \mathbb{Z}$, let $w_n(t) := 2^{-n}w(t)$. (An alternative definition is $w_n(t) := w(2^{-n}t)$). Define I_n to be the foveation using the weight function w_n. This notation is consistent with our previous definition of I_0. Then $\{I_n\}$ is now a sequence of increasingly better foveation as n increases. We can define the set $V_n := \{I_n : I \in L^2(\mathbb{R})\}$ and obtain a tower of spaces

$$\cdots \subseteq V_{-1} \subseteq V_0 \subseteq V_1 \subseteq \cdots.$$

We have so far been describing the continuous foveation transform. In practice, we need a discrete version. Let $B_0 = \{b_m : m \in \mathbb{Z}\} \subseteq L^2(\mathbb{R})$ be a countable set. We treat B_0 as a basis for a subspace $V \subseteq L^2(\mathbb{R})$. Given any foveated function $I(x)$, we can define its projection into V via

$$I^B(x) := \sum_{m \in \mathbb{Z}} \langle I, b_m \rangle b_m(x).$$

Of course, we need some conditions on B. We want V to be an approximation to V_0. A "frame-like" restriction [5] is as follows: for all $I \in V_0$,

$$A\|I\|_w \leq \sum_{m \in B} |\langle I, s_m \rangle|^2 \leq C\|I\|_w. \tag{4}$$

for some constants A, C. Also, as in Mallat, we can require V to be translation invariant in a suitable sense.

6 Practical Foveation Methods

Computing I_0 from its definition is computationally inefficient. We consider practical methods to approximate I_0. There are two general approaches. The first approach is based on **sampling**. Previous papers on foveation can be put under this approach. This is computationally efficient (usually implemented with lookup tables), but lack flexibility when different weight functions are required. In this paper we introduce a second approach. It exploits any multiresolution representation of an image of I: to create a foveated image from I, we 'cut' the necessary pieces from each level of the multiresolution representation and 'glue' them together. There is considerable freedom in choosing the multiresolution representation (e.g., different kinds of wavelet transforms) and in how we cut and paste.

6.1 Foveation via Sampling

We omit details in this extended abstract.

6.2 Foveation via Wavelets

We give an approximation of foveation using wavelets. This is the demonstration at our website URL `http://jinai.nyu.edu:9999/visual/demo.html`. Let ψ be the mother wavelet and ϕ be the corresponding father wavelet. Let $\psi_{m,n}$ be the translated and dilated version of ψ. Specifically,

$$\psi_{m,n}(t) = \psi(2^{-m}(t-n)).$$

Similarly, define $\phi_{m,n}$ in the same way. We choose the scaling function s to be a father wavelet. We also assume that $\{\psi_{m,n}\}$ form an orthonormal basis for $L_2(\mathbb{R})$.

The approximation is as follows.

$$I^{\text{wav}} = \sum_{(m,n)\in B(w)} \langle \psi_{m,n}, I\rangle \, \psi_{m,n}(x),$$

where $B(w)$ is a discrete set depending on the weight function w,

$$B(w) = \{(m,n) | 2^m \geq w(n2^m)\}\,.$$

For the standard weight $w = \alpha|t| + \beta$, the set $B(w)$ is $\{(m,n)|2^m \geq \alpha|n|2^m + \beta\}$, which is $\{(m,n)|2^{-m}c_1 + |n| \leq r\}$ where $c_1 = \alpha^{-1}\beta$ and $r = \alpha^{-1}$. This means that we just have to take at most r coefficients from each level. If the foveal resolution β is greater than 2^m, then no coefficients from level m is in $B(w)$ (m could be a negative integer). This simple structure is useful for fast reconstruction, which are crucial in our application.

If w is the blending of w_1 and w_2, I_w^{wav} and $I_{w_1}^{\text{wav}}$ are the foveation with weight w and w_1 respectively, then

$$I_w^{\text{wav}}(x) = I_{w_1}^{\text{wav}}(x) + \sum_{(m,n)\in B(w_1,w_2)} \langle \psi_{m,n}, I\rangle \, \psi_{m,n}(x), \tag{5}$$

where $B(w_1, w_2)$ is the set $B(w_2) - B(w_1)$. This property suggests application in progressive transmission. It implies that the new foveation with weight w could be computed using the previous foveation with w_1 by adding their difference, which could be conveniently represented by coefficients in $B(w_1, w_2)$.

This method could be easily generalized to higher dimension. The 2-D standard weight is defined as

$$w_{\alpha,\beta,\gamma}(t) := \alpha\|t-\gamma\| + \beta,$$

where $\|.\|$ could be ∞- or 2-norm in $\mathbb{R}^2$, $\gamma \in \mathbb{R}^2$ and $\beta \in \mathbb{R}$. For simplicity, we use the 2-D Haar function $h(x,y) = h(x)h(y)$. The two dimensional weighted distance is

$$d_w(f,g) := \int\int w(x,y)^{-2}\|f(x,y) - g(x,y)\|^2 dx\ dy$$

.

Figure 4(b) shows a foveation using the Haar wavelet. We use the weighted distance to measure the quality of approximations using different wavelets. Figure 5(a) shows experimental results of this comparison.

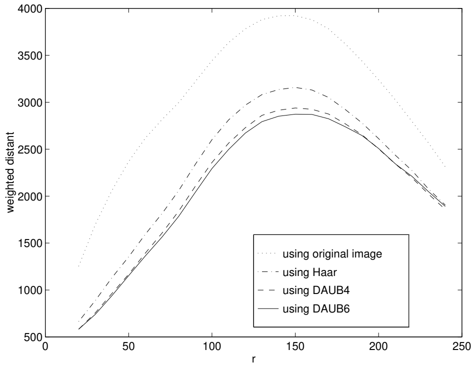

(a)

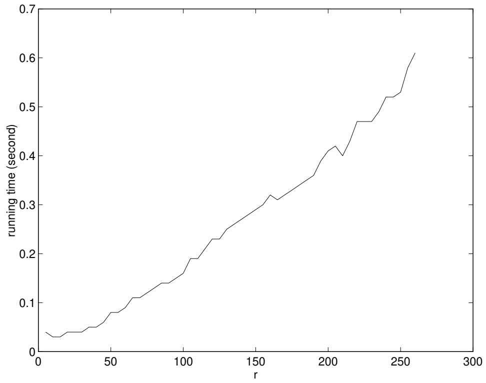

(b)

Figure 5: (a) The weighted distance $d_w(I, I_0)$ and $d_w(I^{\text{wav}}, I_0)$ where I^{wav} is computed using different wavelets. The weight function used is $w(t) = \frac{1}{r}||t||_\infty$. (b) Running time of inverse wavelet transform for weight function $w(x) = \frac{1}{r}|x|$.

6.3 Experimental Results

We now consider two dimensional images. The computation of a foveation via wavelets are computed in three steps.

1. Perform wavelet transform on the image I.
2. Select $C = \{\langle I, \psi_{m,n}\rangle \,|(m, n) \in B(w)\}$.
3. Construct I^{wav} using C by inverse wavelet transform.

In the thin-wire application, step (1) is a precomputation step by the server. Step (2) is a real time computation by the server in response to the client's request. The set C is sent across the thin-wire and step (3) is performed by the client. Steps (1) and (2) are repeated with a new weight function w based on the client's request. In response to the new weight function, the server only need to send the additional coefficients.

In this setting, the performance of the method depends on two measures. The number of bytes send through the thin-wire and somewhat less important, computational efficiency of inverse wavelet transform in stage (3).

The coefficients of wavelet transform could be organized in the standard way [5] as shown in figure 9(a). Furthermore, the set $B(w)$ could be visualized as a union of squares of the same size as shown in figure 9(b). Alternatively, if we change the weight function to $w(t) = \frac{1}{2}||t||_2$ by switching the norm, then $B(w)$ is a union of circles of the same size. It is useful to treat $B(w)$ as a sparse matrix.

Figure 6 shows the number of coefficients required to represent the foveation using standard weight $w(t) = \alpha||t||_\infty$, for various rates α (note that $r = \alpha^{-1}$). Since number of coefficients does not truly represent the number of bytes required, we also show the number of bytes needed to stored the coefficients after lossless compression. In this example, we use GNU 'gzip' for the lossless compression. Before compression, each coefficient is represented by either a two, three, or four-byte integer depending on it's possible range.

As mention before, $B(w)$ could be stored easily in a sparse matrix. Furthermore, Haar function has the desirable property that a coefficient only affects $2 \times 2 = 4$ coefficients from the next multiresolution level. We code an inverse wavelet transform that exploits these two facts. Figure 5(b) gives its running

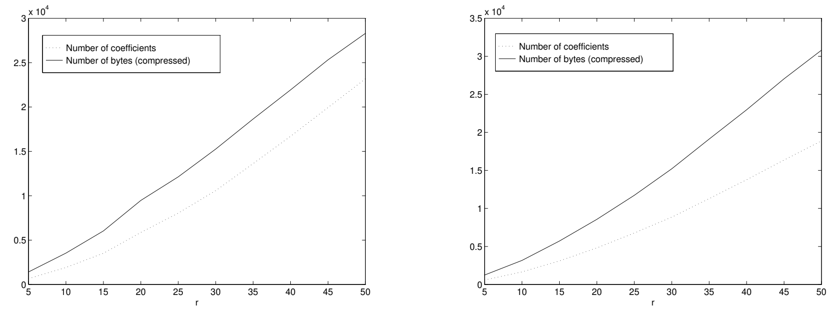

Figure 6: (a) Foveating figure 1(a). Size of original image is 375×469 pixels with 256 gray levels. When $r = 20$, the image is compressed to 0.43 bit per pixel. When $r = 50$, it is compressed to 1.28 bit per pixel. (b) Foveating a picture of mandrill's face. Size of image is 480×500 pixels. When $r = 20$ and $r = 50$, it is compressed to 0.29 and 1.03 bit per pixel respectively.

time on a 512 by 512 image, performed on a Sparc5 station. For example, for $r = 50$, it takes less than 0.1 second. In practice, r is less than 50 for a 512×512 image. A general inverse wavelet transform that does not make use of the sparseness takes about 0.63 second.

7 Open Questions, Further Work

Foveated images have great potential in reducing bandwidth in visualization applications. In this paper, we introduced a new foveation method based on wavelets and formulated the general mathematical framework. Our experimental results show favorable compression rates and timings. E.g., a compression ratio of 20:1 can easily be obtained for our images and typical r values. This ratio will increase for larger images and/or smaller r.

We are currently looking at the mathematical analysis of multiresolution foveation. On the practical side, we will implement continuous foveation techniques in which successive transmissions send only the incremental information. This will further reduce the bandwidth requirement. There are many applications of such techniques: one is the automatic adjustment of resolution to the speed of the gaze point. Another is the problem of "progressive transmission" of large or high resolution images.

Our formulation allows one to compare the various foveation approaches, including the superpixel approach of previous authors. We hope to look into this. Another direction is the foveation of color images. There are new issues here related to the use of color maps.

References

[1] J. Aloimonos, I. Weiss, and A. Bandyopadhyay. Active vision. *Intl. J. Computer Vision*, 2:333–356, 1988.

[2] R. Bajcsy. Active perception. *IEEE Proceedings*, 76(8):996–1005, 1988.

[3] Benjamin B. Bederson. *A miniature space-variant active vision system: Cortex-I.* PhD thesis, Computer Science, Courant Institute, New York University, 1992.

[4] C. Colombo, M. Rucci, and P. Dario. Integrating selective attention and space-variant sensing in machine vision. In Jorge L.C. Sanz, editor, *Image Technology: Advances in Image Processing, Multimedia and Machine Vision*, pages 109–128. Springer, 1996.

[5] Ingrid Daubechies. *Ten Lectures on Wavelets.* SIAM, 1992.

[6] D.H.Hubel and T.N.Wiesel. Receptive fields and functional architecture of monkey striate cortex. *J. Physiol. (Lond.)*, 195:215–243, 1968.

[7] D.H.Hubel and T.N.Wiesel. Sequence regularity and geometry of orientation columns in the monkey striate cortex. *J. Comp.Neurol.*, 158:267–293, 1974.

[8] Mark S. Franzblau. Log mapped focal driven procedural rendering: an application of the complex log map to ray tracing. Master's thesis, Courant Institute, New York University, 1991.

[9] Markus H.Gross, Oliver G. Staadt, and Roger Gatti. Efficient triangular surface approximations using wavelets and quadtree data structures. *IEEE Transactions on Visualization and Computer Graphics*, 2(2), 1996.

[10] Hutchinson et al. Human-computer interaction using eye gaze input. *IEEE Transactions on Systems, Man & Cybernetics*, 19(6):1527–1534, 1989.

[11] R. Jacob. Emerging user-interface media: potentials and challenges – eye tracking. In *Tutorial, SIGGRAPH '89*, 1989.

[12] Tom H. Koornwinder, editor. *Wavelets: an elementary treatment of theory and applications.* World Scientific, Singapore, 1993.

[13] Philip Kortum and Wilson S. Geisler. Implementation of a foveated image coding system for image bandwidth reduction. In *Human Vision and Electronic Imaging, SPIE Proceedings Vol. 2657*, pages 350–360, 1996.

[14] Stephane G. Mallat. A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 11:674–693, 1989.

[15] Yves Meyer. *Wavelets and Operators.* Cambridge University Press, Cambridge, 1992.

[16] David A. Robinson. The oculomotor control system: a review. *Proc.IEEE*, 56(6):1032–1049, 1968.

[17] Alan S. Rojer. *Space-variant computer vision with a complex-logarithmic sensor geometry.* PhD thesis, Computer Science, Courant Institute, New York University, 1989.

[18] Alan S. Rojer and E.L. Schwartz. Design considerations for a space-variant visual sensor with a complex-logarithmic sensor geometry. In *10th ICPR*, volume 2, pages 278–285, 1990.

[19] E.L. Schwartz. The development of specific visual projections in the monkey and the goldfish: outline of a geometric theory of receptotopic structure. *J.Theoret.Biol.*, 69:655–685, 1977.

[20] E.L. Schwartz. Spatial mapping in primate sensory projection: analytic structure and relevance to perception. *Biological Cybernetics*, 25:181–194, 1977.

[21] E.L. Schwartz, D.N. Greve, and G. Bonmassar. Space-variant active vision: Definition, overview and examples. *NeurNet*, 8(7–8):1297–1308, 1995.

[22] T.H.Reeves and J.A.Robinson. Adaptive foveation of mpeg video. *Proceedings, 4th ACM International Multimedia Conference*, 1996.

[23] R.S. Wallace, B.B. Bederson, and E.L. Schwartz. Voice bandwidth visual communication through logmaps: The Telecortex. In *Workshop on Applications of Computer Vision*, pages 4–10, 1992.

[24] R.S. Wallace, P.W. Ong, B.B. Bederson, and E.L. Schwartz. Space-variant image processing. *Intl. J. of Computer Vision*, 13(1):71–90, 1994.

[25] L. R. Young and D. Sheena. Method and design: Survey of eye movement recording methods. *Behavior Research Methods and Instrumentation*, 7(5):397–429, 1975.

Appendix: Superpixel Geometries

Figure 8 and figure 7 illustrate the super-pixel arrangement used in [24] and [13] respectively.

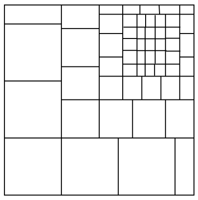

Figure 7: A square superpixel arrangement.

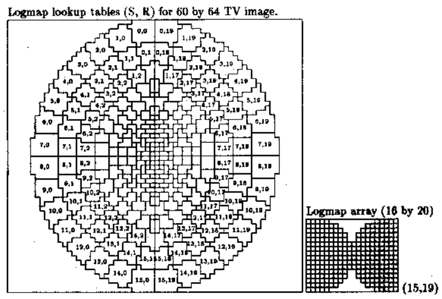

Figure 8: Superpixel geometry based on logmap.

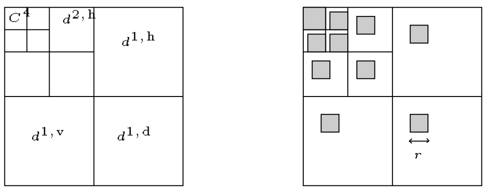

Figure 9: (a) Standard organization of the two-dimensional wavelet transform. (b) The filled area corresponds to the set $B(w)$, where $w(t) = r^{-1}\|t\|_\infty$.