

Hyperacuity Sensing for Image Processing

Sumit Basu
sbasu@MIT.EDU

Advisors:
Warren Jackson (Xerox PARC)
David Biegelsen (Xerox PARC)
David Jared (Xerox PARC)
Fernando Corbató (MIT)

Abstract

The human eye's response to high spatial frequency content in images is highly dependent on the nature of the features: the eye integrates periodic intensity variations at high frequencies to form a solid gray, yet is very sensitive to high-contrast steps in intensity (which contain high frequency components), such as edges. This study presents the development and effectiveness of a scanning system that has a similarly varying response to spatial frequency content. The system is implemented by an array of PSD (position sensitive detector) scanning elements, each of which produces four values which can be composed to retrieve intensity and intensity gradient information over the detector's area. It is shown that this information, along with additional information from immediately neighboring pixels, can be used to determine the frequency content and edge locations of the image area covered by the detector. The accuracy of this determination is shown to vary in a way similar to that of the human eye's response. A single-pass algorithm (BBJJ) is developed which determines how to render the area of the image corresponding to each detector as either an edge or a smooth gradient. A version of the algorithm is also presented for standard scanning detectors, which composes the information from multiple detectors to construct values analogous to those available from each PSD element. The algorithms described were implemented in an existing simulation framework. The performance of each of these methods in comparison to that of a standard scanning system is shown (in simulation) for three major classes of images: line art, continuous tone, and halftoned images.

Introduction

The typical approach to improving performance in scanning systems has been to increase the density of the scanning elements. While this undeniably improves the image quality, it is usually not an efficient or economic solution. Consider, for example, an image that has a single, solid black circle in the center of a white background. A typical scanning system will take the average intensity incident on each detector-sized area of this image and represent this value in an n by n pixel block on the screen, where n is the scaling factor. At lower resolutions, this approach results in “jaggies” (jagged lines that did not exist in the original image - see, for example, Figure 20). Sampling at higher and higher resolutions certainly reduces the visibility of these features, but also results in sampling constant-intensity regions at an unnecessarily high resolution.

In this paper, a method is presented for reconstructing image features without going to higher resolutions by using a novel scanning element, the PSD, and an algorithm which composes the output from these devices and their nearest neighbors to reconstruct sub-detector sized features. An algorithm is also presented for performing this reconstruction using the output from standard grayscale detectors. The goal of this work is not to perfectly rerender an input image - clearly any scanning method will output only a subset of the information content of the original, and thus a perfect reconstruction is impossible. Instead, the goal is to faithfully reproduce only the *visually significant* features - those features that are visible to the human eye at a normal viewing distance. In other words, the standard of reconstructed image quality to be used is not simply the squared pixel-by-pixel error between the original and reconstructed images, but one which takes into account the response of the human visual system. Prior research has suggested convincing arguments for adopting such a stance (see Pearlman [11] and Bloomer [1]). By adopting this standard, it is possible to measure the practical performance of these methods without worrying over features that are effectively invisible to the eye.

The Hyperacuity Model

The inspiration for this work came from the operation of the human visual system. Though the eye has a relatively low “detector density” in terms of light receptors, the visual system employs various techniques that produce what appear to be very high resolution images. Though the exact workings of this system are not known, it is possible to gain some understanding of what goes on from observing the visual response to various image features. This response varies greatly with the size and quality of the features in the image. For a pattern of adjacent vertical strips of 0% and 100% intensities (black and white, respectively), the ability of the human visual system to distinguish the pattern from the average gray drops off significantly at around 300 cycles/inch for a standard viewing distance of 12”. For lower levels of contrast between the intensities, the response drops off at even lower spatial frequencies. On the other hand, the visual system is extremely sensitive to high contrast (non-periodic) steps in an image, such as black-white edges. Though edges contain information at frequencies far higher than 300 cycles/inch, the visual system will not blur an edge and is very sensitive to its exact position. A significant amount of prior research has been performed on this “hyperacuity sensing” ability: Nishihara [10] provides a concise description of the effect and cites several works examining the effect in detail (citations 1-9). Fahle [2], [3] and Martinez-

Uriegas [9] provide explanations and models for the effect. Lennie and D’Zmura [7] examine, for various contrast levels, the transition of the visual system’s response from detailed acuity to integrating effects with increasing spatial frequency.

The hyperacuity effect can be easily observed by simply looking at a newspaper page with a halftoned image. When held far enough away, the dots in the image will appear to blend into continuous tones, while the edges in the headline text will remain sharp and clear. An ideal scanner should exhibit this effect as well. It should render features with spatial frequencies beyond some range as continuous gray tones, yet be able to place an edge with high pixel accuracy.


The PSD (Position Sensitive Detector) has the capability of achieving this hyperacuity standard. The four values that the device outputs can be composed to find the centroid of light intensity over the area of the detector and a first-order approximate of the 2D gradient. As will be shown in later sections, this information can be used to accurately determine sub-detector edge presence and location. Clearly this information could not be used to resolve a sequence of twenty black and white strips if it were incident on the detector - the device would effectively integrate such detail into a solid gray. However, if a single edge were incident on the detector, its position and orientation could be determined exactly (disregarding noise from the device). Thus, in a manner similar to that of the human visual system, the PSD integrates periodic high frequency content but is sensitive to high-contrast steps in intensity. More specifically, the PSD can be used to accurately resolve sub-pixel edges if there is no more than one edge incident on the detector. The development and limitations of this ability will be shown in the following sections.

Characteristics of the PSD

The PSD device itself is basically a simple PIN junction with contacts at each edge held at virtual ground:

FIGURE 1.

The Position Sensitive Detector



When light is incident at a single point on the detector, electron-hole pairs are generated in the intrinsic layer. The charge pairs separate in the vertical field. The electrons in the n-layer then flow to the contacts (all held at ground potential) with currents at each con-

tact linearly (with slope =-1) proportional to the distance between the point of generation (the point of incident light) and the contact. Any general distribution of intensity can be composed from superpositions of this single-point case, resulting in the following expressions for the currents at the contacts:

$$i1 = \frac{1}{2} \int_{-1}^1 \int_{-1}^1 I(x, y)(1 - y) dy dx$$

$$i2 = \frac{1}{2} \int_{-1}^1 \int_{-1}^1 I(x, y)(y + 1) dy dx$$

$$i3 = \frac{1}{2} \int_{-1}^1 \int_{-1}^1 I(x, y)(1 - x) dx dy$$

$$i4 = \frac{1}{2} \int_{-1}^1 \int_{-1}^1 I(x, y)(x + 1) dx dy$$

As will be shown in a later section, this information can be composed to find the centroid of intensity over the area of the detector.

The concept of a PSD is certainly not new: these devices have been used extensively for obtaining position information. In typical applications, the entire sensing area is a single PSD, and the object to be tracked acts as a point source of light; thus the centroids of intensity obtained from the PSD are equivalent to the object's position. The methods presented in this paper, however, make up the first instance of the construction and use of these devices at the scanning element level (150 microns). The goal is to manufacture 100spi-300spi (spot per inch) arrays of PSD's over areas the size of a printed page (8.5" by 11") so that all of the information on a page can be captured in a single scan. At this point, arrays of twenty PSD's have already been built and tested. Preliminary results indicate that the devices are sufficiently accurate to allow use of the algorithms presented in this paper.

The motivation for this study was to determine whether the hyperacuity sensing properties of the PSD devices could be used to reconstruct visually significant sub-detector-sized features in an image. Because no full-sized PSD arrays exist at this point, all of the work was done in simulation. A framework for simulation of the PSD device and standard detectors called "hyper" was built in C by David Jared. This system used pre-scanned images as input and produced the output of a simulated sensor array. To do this for the PSD case, a group of n by n pixels was treated as a single "detector," and the discrete analogues of the expressions for $i1$ through $i4$ above were used to produce the contact currents. For the standard detector case, the system used the average intensity over all of the pixels in the n by n block for the detector current. The algorithms presented in this paper were developed and tested within this simulation framework.

The BBJJ Algorithm

The goal of this study, in terms of using hyperacuity information, was to render edge locations with sub-detector accuracy while using only local information (i.e., hyperacuity values available from the detector and its nearest neighbors). In other words, if a sin-

gle black-white edge passed through a detector, the goal was to rerender this edge (and preserve its location) on the area of the image corresponding to this detector. In contrast, the standard approach in scanning systems is to render the average intensity over the corresponding area of the image (see Hall [5]).

The algorithm developed in this study to achieve this goal has been named the “BBJJ” algorithm. It models the intensity distribution of a given detector in one of two ways: as a linear edge between two constant intensities or as a non-edge (i.e., a continuous variation in intensity). For the case of an edge, it attempts to determine the location of the edge and the intensity value on each side of the edge, using the values available from the detector and its nearest neighbors. For the case of the non-edge, it creates a continuous model of the intensity distribution using information from the detector alone.

As opposed to other algorithms which attempt to “segment” an image, i.e., divide it into regions of various image types and apply different heuristics to each type, the BBJJ algorithm was designed to “autosegment” the image on a detector-by-detector basis. Instead of attempting to determine what areas of the image are grayscale, line art, or halftone, it applies a constant set of heuristics to every detector. The area corresponding to each detector is thus rendered only on the basis of the information from the detector and from its nearest neighbors.

The steps of the algorithm are presented in outline form below and then discussed in further detail in later sections. Note that a separate paper describing the specific means of implementing the algorithm will be available shortly.

FIGURE 2.

Outline of the BBJJ Algorithm

1. **Compute I_0 (total intensity) and the x and y moments (\bar{x} and \bar{y}) of the intensity distribution incident on the detector/macrodetector**
2. **If the magnitude of the centroid ($\sqrt{\bar{x}^2 + \bar{y}^2}$) is large, render area as an edge**
 - a. **Find gray value on one side of edge (G_1, G_2)**
 - +If a clear context value exists, use it
 - +Otherwise, use interpolated values
 - b. **Find parametrization of edge from moments using lookup tables**
 - c. **Convert gray-on-gray moments to gray-on-black moments to find actual edge position**
 - d. **Use actual edge position to determine other gray value**
 - e. **Use edge position and parameters to render area as a G_1, G_2 edge**
3. **If the magnitude of the centroid is small, render area as a non-edge**
 - a. **Compute the plane/bilinear model parameters**
 - b. **Render area using plane/bilinear model**


Step two, the heart of the BBJJ algorithm, provides the basis for the hyperacuity abilities of the detector. The theoretical development of this step will thus be presented first.

Afterwards, the particulars of the entire algorithm (all steps) for the PSD (Position Sensitive Detector) and standard detector cases will be discussed.

Several basic conventions will be made to simplify the following discussion of the algorithm. First, the detector's area is on a coordinate system as shown below:

FIGURE 3.

The detector coordinate system



Second, the intensity function $I(x,y)$ is defined over this area, giving the intensity at each point on the detector. The integral of this function over the area of the detector results in the total light intensity incident on the detector, I_0 . Third, a high intensity value (100.0) indicates a white region, while a low intensity (0.0) indicates a black region.

Assumptions

We begin by assuming we know the first moments of intensity (\bar{x} and \bar{y}) over the area of the given detector. The x moment is given by:

$$\bar{x} = \frac{\int_{-1}^1 \int_{-1}^1 xI(x, y) dx dy}{\int_{-1}^1 \int_{-1}^1 I(x, y) dx dy}$$

and the y moment by:

$$\bar{y} = \frac{\int_{-1}^1 \int_{-1}^1 yI(x, y) dx dy}{\int_{-1}^1 \int_{-1}^1 I(x, y) dx dy}$$

We also must know the total intensity incident on the detector, I_0 :

$$I_0 = \int_{-1}^1 \int_{-1}^1 I(x, y) dx dy$$


Lastly, it is assumed that there is no more than one edge incident on a given detector, and that all edges are linear unless otherwise indicated. The implications of this last assumption will be discussed in later sections.

Initial Estimate of Edge Location

If it were known that the detector covered an area containing a uniform gray area and a uniform black area of intensity zero (a grey on black edge), the position of the edge could be determined exactly. This is illustrated in the one-dimensional case below:


FIGURE 4.

Gray on Black edge



The moment would fall in the exact center of the region of intensity G_2 , regardless of G_2 's value, as long as it was non-zero and G_1 was zero. In other words, we have $x_0 = 1 - 2(1 - \bar{x})$. The edge location for a gray-black edge can be similarly found in the two-dimensional case. The edge is parametrized by the quantities ρ and θ . θ is defined as the angle, measured from the x axis, which describes the orientation of the edge (i.e., a line passing through theta will be normal to the edge location). Note that θ is always directed towards the high-intensity region. ρ is defined as the distance to the edge along a line passing through θ . As a result, ρ is not necessarily positive, as illustrated in the cases below.

FIGURE 5. Parametrization of Edge Location with ρ and θ for Two Edge Locations




A unique mapping can be made between (\bar{x}, \bar{y}) space and (ρ, θ) space. An exact formulation for ρ and θ , while possible, is unwieldy and computationally expensive to solve. A lookup table was used in the implementations in this study.

Returning to the one-dimensional case above (Figure 4), it is clear that when the value of G_I is raised to a non-zero value, the moment is shifted towards the G_I region. As a result, the moments will no longer map to the correct edge location. It is thus necessary to find what the moments would have been had the distribution been that of a gray-black edge instead of a gray-gray edge.

Fortunately, it is possible to transform the moments of a gray-gray edge to those of a gray-black edge if G_I is known. Consider the intensity distribution below. As shown below, this can be seen from two perspectives: a region of intensity G_I and one of intensity G_2 , or a region of intensity 0 and a region of intensity $(G_2 - G_I)$, with a constant intensity G_I added to every point of the detector.

FIGURE 6. Two Representations of a Gray on Gray Edge (offsetting by G_I)



From the second perspective, the intensity distribution of the detector is $I(x) = I'(x) + G_I$, where $I'(x)$ is 0 for $x < x_0$ and $(G_2 - G_I)$ for $x > x_0$. \bar{x}' is then defined as the moment

corresponding to the intensity distribution described by $I'(x)$. The calculation for the first moment of this distribution is shown below for the two-dimensional detector:

$$\bar{x} = \frac{\int_{-1}^1 \int_{-1}^1 xI(x, y) dx dy}{\int_{-1}^1 \int_{-1}^1 I(x, y) dx dy} = \frac{\int_{-1}^1 \int_{-1}^1 x[I(x, y) + G_1] dx dy}{I_0} = \frac{\bar{x}' \int_{-1}^1 \int_{-1}^1 I(x, y) dx dy}{I_0}$$

since $\bar{x}' = \frac{\int_{-1}^1 \int_{-1}^1 xI(x, y) dx dy}{\int_{-1}^1 \int_{-1}^1 I(x, y) dx dy} \Leftrightarrow \int_{-1}^1 \int_{-1}^1 xI(x, y) dx dy = \bar{x}' \int_{-1}^1 \int_{-1}^1 I(x, y) dx dy$

and $\int_{-1}^1 \int_{-1}^1 xG_1 dx dy = G_1 \int_{-1}^1 \int_{-1}^1 x dx dy = 0$

In addition, we can use the relation

$$\int_{-1}^1 \int_{-1}^1 I(x, y) dx dy = \int_{-1}^1 \int_{-1}^1 [I(x, y) + G_1] dx dy = \int_{-1}^1 \int_{-1}^1 I(x, y) dx dy + \int_{-1}^1 \int_{-1}^1 G_1 dx dy$$

to rewrite the expression for \bar{x} as:

$$\bar{x} = \bar{x}' \left(\frac{\int_{-1}^1 \int_{-1}^1 I(x, y) dx dy - \int_{-1}^1 \int_{-1}^1 G_1 dx dy}{I_0} \right) = \bar{x}' \left(\frac{I_0 - \int_{-1}^1 \int_{-1}^1 G_1 dx dy}{I_0} \right)$$

If we now refer to the second term in the numerator as I_{G1} , which corresponds to the total intensity of a detector covered entirely by intensity G_1 , we find the following expression for \bar{x}' , the moment of the gray-black distribution $I'(x, y)$:

$$\bar{x}' = \bar{x} \left(\frac{I_0}{I_0 - I_{G1}} \right)$$


The y moment can be transformed in the same way:

$$\bar{y}' = \bar{y} \left(\frac{I_0}{I_0 - I_{G1}} \right)$$

This corresponds to what the moment would have been if G_1 had been zero (i.e., if the distribution had been that of a gray-black edge). Because we can find the location of the edge precisely when given the moments of a gray-black edge (see the explanation of Figure 4), we can use these transformed moment to precisely determine the location of the gray-gray edge.


In some cases, though, I_{G1} is not available. The transformation of the moments can then be performed using I_{G2} in place of I_{G1} . $I'(x,y)$ is then defined as being (G_1-G_2) (which is negative) for $x < x_0$ and 0 for $x > x_0$, as illustrated below

FIGURE 7. Two Representations of a Gray on Gray Edge (offsetting by G_2)



The expressions for calculating the transformed moments are the same as those presented earlier, with I_{G1} replaced by I_{G2} . These transformed moments correspond to the centroids of the dark region (the region of intensity G_1), as opposed to those of the light region (as in the previous cases, as shown in Figure 5). Similarly, the resulting θ now points in the direction of the dark region, though the ρ is still the distance to the edge. An example of this is shown below.

FIGURE 8. Parametrization of Edge Location After Transforming Moments Using I_{G2}



Looking for Context Gray Values (I_{G1} and I_{G2})

In order to perform the transformations described above, it is necessary to have either I_{G1} or I_{G2} . It is hoped that nearby detectors will provide a “context” for the edge, i.e., that the nearby detectors on each side of the edge will have the same intensity as the respective side of the edge. Ideally, for example, we would find a detector nearby with uniform intensity distribution $I(x,y) = G_1$. Then the total intensity of that detector, I_0 , would be I_{G1} , the gray value needed to transform the moments and render the edge precisely. If the features in the image are widely spaced, such a detector can be easily found by choosing the nearest neighbor in the direction normal to the edge. Because either intensity will suffice, we can use I_{G2} instead of I_{G1} if the detector on the low-intensity side of the edge contains an edge. However, features are often even more closely spaced than this, and methods for dealing with these cases are described in the sections below.

Interpreting the Transformed Moments

At this point, the magnitude of the transformed moment, $|m|$ ($\sqrt{x^2 + y^2}$), is considered in order to determine the spatial frequency content of the detector. If this value is smaller than some threshold $|m_0|$, the intensity distribution across the detector must be fairly constant, and it is not rendered as an edge. Various techniques for rendering the detector's area in this case are discussed in a later section. Otherwise, the detector area is rendered as a straight edge between two regions of constant intensities I_{G1} and I_{G2} .

Calculating G_2

To render the detector's corresponding area in the image, it is necessary to have the value of G_2 , assuming we are already using G_1 for context. It is possible to determine this value without any additional information. It is known that the total intensity over the detector is

$$I_0 = \int_{-1}^1 \int_{-1}^1 I(x, y) dx dy = \int \int_{A_{G1}} G_1 dx dy + \int \int_{A_{G2}} G_2 dx dy$$

where A_{G1} and A_{G2} are the regions on the two sides of the edge. Because the intensities are constant over this region, and the areas of the two regions must add up to A_d , the total area of the detector, we can simplify this to:

$$I_0 = A_{G1}G_1 + (A_d - A_{G1})G_2$$

A_{G1} can be easily calculated, since we have already determined the position of the edge. We can now solve for G_2 :

$$G_2 = \frac{I_0 - A_{G1}G_1}{A_d - A_{G1}}$$

Rendering the Result

All of the information necessary to render the detector area is now available. There remains the issue, though, of what resolution to rerender the image at. If the detector area is mapped to a single pixel on the output device, there is clearly no advantage to using this algorithm. The subpixel information produced by this algorithm is only beneficial when rerendering the image at a resolution high enough to display that information. This issue will be discussed further in a later section.

Calculating the Moments with Actual Detectors

To apply the BBJJ algorithm, it is necessary to find the information listed in the assumptions above, namely \bar{x} , \bar{y} , and I_0 . This information can be easily derived for the position sensitive detector (PSD) case, as will be shown below. However, for the standard detector case, it is necessary to use twice the resolution of the PSD array to get approximately the same amount of information. This is because the PSD produces three independent values out of $i1$, $i2$, $i3$, and $i4$ (due to the constraint described in the section below), while a standard grayscale detector produces only one (I_0). Two by two arrays of the standard detectors can be composed into “macrodetectors,” from which enough information can be extracted to use a slightly modified form of the BBJJ algorithm.

Position Sensitive Detectors

The expressions for the four currents from the PSD shown earlier can be readily combined to yield the first moments and the total intensity over the area of detector. The total intensity can be calculated in two ways:

$$I_0 = i1 + i2 = i3 + i4$$

This can be easily verified:

$$\begin{aligned} i1 + i2 &= \frac{1}{2} \int_{-1}^1 \int_{-1}^1 I(x, y)(1 - y) dy dx + \frac{1}{2} \int_{-1}^1 \int_{-1}^1 I(x, y)(y + 1) dy dx \\ &= 2 \cdot \frac{1}{2} \int_{-1}^1 \int_{-1}^1 I(x, y) dy dx = I_0 \end{aligned}$$

The development is identical using $i3$ and $i4$. As a result, it is clear that there are only three independent values coming from the PSD, since any three can be used to find the fourth value.

We can now use the total intensity to aid in calculating the moments. The expression for the x moment is given below:

$$\bar{x} = \frac{i4 - i3}{i4 + i3}$$

Similarly, we find for the y dimension that

$$\bar{y} = \frac{i2 - i1}{i2 + i1}$$

The verification for the x moment is shown below:

$$\begin{aligned} \frac{i4 - i3}{i4 + i3} &= \frac{\frac{1}{2} \int_{-1}^1 \int_{-1}^1 I(x, y)(x+1) dx dy - \frac{1}{2} \int_{-1}^1 \int_{-1}^1 I(x, y)(1-x) dx dy}{I_0} \\ &= \frac{2 \cdot \frac{1}{2} \int_{-1}^1 \int_{-1}^1 I(x, y)x dx dy}{I_0} = \frac{\int_{-1}^1 \int_{-1}^1 xI(x, y) dx dy}{\int_{-1}^1 \int_{-1}^1 I(x, y) dx dy} = \bar{x} \end{aligned}$$


The development for the y moment is identical. Thus, for ideal PSD detectors and in the absence of noise, it is possible to exactly determine the first moments and total intensity.

Standard Detectors

For a two by two array of standard detectors (a macrodetector), we have four pieces of independent information: the total intensity of each of the detectors, labeled $I1..I4$ as shown below:

FIGURE 9.

Macrodetector Currents and Coordinate System




I_0 is thus simply the sum of the four total intensities:

$$I_0 = I1 + I2 + I3 + I4$$

Because of the quantizing that results from discretizing the space, it is not possible to find the moments of the underlying image exactly. However, we can find approximate moments that provide a reasonably unique mapping to ρ and θ . Each of the detector currents is the total intensity over its area. Note that the detector cannot distinguish between two different intensity distributions that have the same total intensity. We can thus consider the intensity distribution for each detector to be concentrated at a single point at the center of the detector:

FIGURE 10.

Simplified Intensity Distribution Model for a Macrodetector



This distribution would be indistinguishable to the macrodetector from any other distribution with the same total intensity for each detector. The moments for such a distribution would be as follows. Note that \bar{x}_m and \bar{y}_m are defined to be the macrodetector moments.

$$\bar{x}_m = \frac{\left(-\frac{1}{2}\right)I_4 + \left(-\frac{1}{2}\right)I_1 + \left(\frac{1}{2}\right)I_3 + \left(\frac{1}{2}\right)I_2}{I_1 + I_2 + I_3 + I_4} = \frac{1}{2} \left[\frac{(I_3 + I_2) - (I_4 + I_1)}{I_1 + I_2 + I_3 + I_4} \right]$$

$$\bar{y}_m = \frac{\left(-\frac{1}{2}\right)I_1 + \left(-\frac{1}{2}\right)I_2 + \left(\frac{1}{2}\right)I_3 + \left(\frac{1}{2}\right)I_4}{I_1 + I_2 + I_3 + I_4} = \frac{1}{2} \left(\frac{(I_3 + I_4) - (I_1 + I_2)}{I_1 + I_2 + I_3 + I_4} \right)$$

These quantities bear a strong resemblance to the moment expressions for the PSD, in that the expressions for both devices consists of a difference of currents from opposing sides of the device over their sum. In addition, because they are moments in the mathematical sense, they can be transformed using a context gray value as shown in the discussion of the BBJJ algorithm. However, the key difference is that the currents from the opposing sides of the PSD contain information from the entire region of the detector, whereas the current from each standard detector within the macrodetector contains information from only its quarter of the region. This leads to degeneracies when attempting to map these moments to ρ and θ , as will be shown in the next section.


Mapping the Moments to the Edge Parametrization

The relationship of the moments to ρ and θ is analytically complex and expensive to compute. On the other hand, finding the macrodetector moments from ρ and θ is straightforward - the expressions derived above can simply be applied to a distribution that corresponds to a gray on black edge at ρ and θ . A lookup table was thus made mapping \bar{x} and \bar{y} to ρ and θ . This was accomplished by first iterating in a set interval over the (ρ, θ) space and finding \bar{x} and \bar{y} for each point. A value of ρ , $\cos\theta$, and $\sin\theta$ was thus found for a series of moment pairs. As expected, these points were not uniformly distributed over (\bar{x}, \bar{y}) space. The values of each of the three quantities were thus interpolated over the entire moment space (-1 to 1) and resampled at 40 divisions per dimension (e.g., 0.05 per division). It was necessary to use $\cos\theta$ and $\sin\theta$ because they are continuous over the entire (x, y) space whereas θ is not - there is a branch-cut of 2π that must be placed somewhere. If such discontinuous a mapping were interpolated over the plane, the region that crossed over the discontinuity would contain values in the midst of the step.

These lookup tables were generated for both the PSD and the array of standard detectors. Because the exact moment can be found with the PSD and every edge location/orientation maps to a unique moment, the accuracy of this method of attaining the edge location can be made arbitrarily precise (with a fine enough resolution for the lookup table). For the standard detectors, though, this is not the case. The macrodetector moments derived above for the macrodetector do not map to unique edge locations. The basic problem is illustrated in the figure below:

FIGURE 11.

Degeneracies in the Macrodetector Moment Mapping



If G_1 is zero (black) and G_2 is non-zero, then as long as $0 < x_e < 1$, \bar{x}_m will be 1 (see the expressions for the macrodetector moments above), since the total intensity of the macrodetector (the denominator in the expression for the macrodetector moment) will be the same as the total intensity of the right side ($I_2 + I_3$, the numerator). As a result, many different edge locations will map to the same macrodetector moments. Thus the mapping is degenerate (i.e., not unique) in this region, and the edge location cannot be recovered from the moments. Note, however, that if G_1 is non-zero and G_2 is zero (in the case illustrated above with $0 < x_e < 1$), this problem does not occur, because the numerator of the macrodetector moment expression changes with the edge location. The macrodetector moment thus changes with x_e in this case, and maps uniquely to an edge location. This asymmetry will be later exploited in order to counteract the degeneracy.

Finding Context Gray Values


Since the spatial resolution and characteristics of the PSD and macrodetector differ significantly, the techniques for finding context gray values are quite different.

Context in the PSD Array

As shown earlier, if the exact moments can be determined (i.e., for the PSD), either context gray value is sufficient to transform the moments and exactly determine the edge location. The intensities of interest are those of the neighboring detectors on either side of the edge (along rays of θ and $\theta + \pi$). Because the moments have not yet been transformed by the context gray values at this point, the exact value of θ is not known. However, due to the nature of the detector array's geometry, we can only choose eight unique directions to look for context in (see Figure 12 above). In addition, while the ρ calculated from the transformed moments will differ significantly from the ρ of the original moments, θ stays fairly constant. As a result, it is possible to use the θ from the initial calculation of the moments to choose the detectors for context. We quantize the angle into eight regions of $\pi/4$, beginning at $-\pi/8$, corresponding to the eight detectors surrounding the detector being rendered. This division is shown below (the detector being rendered is at the center):

FIGURE 12.

Quantization of the Theta Space



An accurate estimate must now be found for either I_{G1} or I_{G2} . Thus there are two possible detectors to choose from: the neighbors on either side of the edge. Once the two potential context detectors are chosen, the moments are calculated for each one. The first criterion for choosing which of the two to use is if $\rho < \rho_0$ for either detector. This is to prevent the use of a detector for context that itself contains an edge, since its total intensity value I_0 will not be the desired context value. If both context detectors satisfy this constraint, it is desirable to choose the detector with the smallest moments, since it will have the least variation in intensity.

If both detectors have moment magnitudes greater than ρ_0 , a somewhat less accurate context value can still be calculated by interpolation. If the intensity distribution over

the area of the detector is modeled as a plane of the form $I(x,y) = Ax + By + C$, the values of A, B , and C can be easily determined to be the following:

$$A = \frac{3(i4 - i3)}{4(i4 + i3)}$$


$$B = \frac{3(i2 - i1)}{4(i2 + i1)}$$

$$C = \frac{1}{4}(i1 + i2)$$

This model can then be applied to whichever of the two context detectors has a smaller moment magnitude and evaluated at the corner/edge of interest. When this value is multiplied by the area of the detector, the desired I_0 is attained. For example, if the context detector is in the upper right corner of Figure 12 ($\theta = \pi/4$), the context intensity would be the plane model for that detector evaluated at $(-1,-1)$ (its lower left corner). The accuracy of this method varies with the actual intensity distribution over the detector. In the one-dimensional case shown below, for example, using this method of interpolation yields an accurate value for the edge of the detector:

FIGURE 13.


Interpolation: Best Case



On the other hand, the next case results in a significant error (at $x=1$).

FIGURE 14.

Interpolation: Worst Case



This last case is particularly pathological, because the initially calculated moments within the detector would be fairly small. This would be interpreted as an indication that the intensity distribution over the detector was fairly constant. In turn, this would imply


that it would be safe to use for context (i.e., that the I_0 of this detector could be used for I_{G1}). If x_e was close enough to 1.0, the $\rho < \rho_0$ constraint would be satisfied as well. However, using the I_0 of this detector as context would lead to a very conspicuous error - essentially, the light and dark values for the edge would appear flipped (since the desired dark context intensity would be a light value, and the light intensity would be calculated from this value and the total intensity to yield a dark value). As a result, a final check is made before rendering the edge to make sure that the intensity of the light side of the edge is actually higher than that of the dark side. If this is not the case, there was clearly an error in attempting to find the context gray value, and the pixel is rendered as a non-edge (see the section on rendering non-edges).

Context in the Standard Detector Array

For the standard detectors, the freedom to choose either context intensity no longer exists. It is necessary to find the gray value on the side of the edge with the non-dominant area. This need arises from the macrodetector moment degeneracy presented earlier. In one dimension, the example below clearly demonstrates the problem:

FIGURE 15.


One-dimensional Illustration of Macrodetector Degeneracy



If I_{G1} was used for context, we would effectively reproduce the situation shown in Figure 11. On the other hand, if I_{G2} is used, the effective resulting distribution would be as follows:

FIGURE 16.

Figure 11 Intensity Distribution Normalized by I_{G2}



This case is free from degeneracy, since the numerator of the macrodetector moment term is no longer equal to the denominator. As long as there is an edge such that the non-zero intensity ($-G_1$ in this case) falls on more than one detector in each dimension, the macrodetector moment is not degenerate. The two-dimensional constraint for this can be realized by requiring that the non-zero intensity cover more than half of the mac-

rodetector (the dominant area). To achieve this situation, the intensity of the non-dominant area must be used for context.

In order to determine whether the light or dark side of the edge covers the non-dominant area, it is necessary to determine what the light and dark values are. To find this information, we must find an approximation of ρ and θ . Fortunately, though the macrodetector moment calculation is degenerate in terms of its uniqueness and sensitivity, it does not cause errors in θ that cross the quantization boundaries shown in Figure 11.

With the PSD array, we had only one choice for a context detector for any of the eight θ regions. Because we can get a total intensity (I_0) value at twice this resolution for the standard detector array (since the macrodetector of two by two standard detectors is being considered equivalent to one PSD), we have several choices for context. To simplify this discussion, the three groupings of neighboring detectors used for attaining context are shown below for the $\theta = 0$ and $\theta = \pi/4$ cases (corresponding to regions I and II in Figure 12). All other cases are rotationally symmetric to these. Note that when more than one detector is used for context, the I_0 values are averaged to attain the context gray value.



FIGURE 17.

Detector Groupings for Context Gray Values



$$\theta = 0$$

$$\theta = \frac{\pi}{4}$$



INSIDE:



OUTSIDE:



ALONG EDGES:




To find the initial estimates of G_1 (the dark side of the edge) and G_2 (the light side), values are determined for all three groupings shown for both the positive and negative \hat{r} directions. Whichever has the highest intensity value for the positive direction is chosen as G_2 , and whichever has the lowest intensity for the negative direction is chosen as G_1 . Because it is known that $A_1G_1 + A_2G_2 = I_0$ and that $A_1 + A_2 = A_d$, it can be determined which intensity covers the non-dominant area. The value for the corresponding I_0 found earlier is then used to transform the moments and determine the edge location. We have made the assumption earlier on that only one edge can fall within a macrodetector. As a result, the intensity of the dominant area can now be found exactly. This is because if the assumption holds, the dominant area must cover more than 50% of the detector area, and thus must fully cover at least one detector.

While the information of all four detectors is used to construct the edge model, only the central region of the macrodetector is rendered, as shown in the figure below. Instead of iterating by macrodetector steps over the image, the algorithm steps detector by detec-

tor, treating every possible grouping of four adjacent detectors as a macrodetector. This is because the macrodetector is most sensitive to edge position in the central region. In addition, if the entire macrodetector were rendered, it would only be possible to have a single edge running through the entire detector. On the other hand, if only the central region is rendered, the area of one macrodetector is rendered by eight different macrodetector groupings. Edges can thus be followed more closely, reducing the jaggedness of the result, particularly when dealing with curved edges.

FIGURE 18.

Rendered Area of Macrodetector



Rendering Non-Edges

After the recalculation of the moments, it is often the case that the magnitude of the moments is smaller than the threshold required to justify the edge model. In these cases, a variety of alternate models can be applied. The simplest among these is rendering the average gray value over the entire detector. This works well when the image fits the model, i.e. there is no variation in the intensity over the area of the detector. However, often a detector for which the edge model does not apply still contains a significant gradient. The three independent pieces of information for the PSD and the four for the macrodetector are sufficient to take this gradient into account in rendering a non-edge. The particular models used, which make use of all the information available in each case, are described below.

PSDs: the Plane Model

The plane model is used for the PSD array because it requires only three pieces of independent information: i.e., precisely what is available. This model has already been presented in the earlier description of interpolating to find context gray values. As described there, this approach models the entire area of the detector as a single plane and can sometimes overshoot or undershoot the actual values at the edges of the detector. As a result, the edges of the plane models of two adjacent detectors may not be the same intensity, which can produce a false “edge” in the rendered image. Because the spatial extent of such a feature is only a single detector edge, it is usually not significant. If the area in the output image corresponding to a single detector is fairly large, though, the effect can be noticeable.

Standard Detectors: the Bilinear Model

The bilinear model requires the four pieces of independent information that are available from a macrodetector. Green [4] develops this method in detail. It matches the average intensity of each detector at the exact center of the detector. The values between centers are linearly interpolated from the four nearest centers. The expression for the intensity of the rendered area of a macrodetector is thus as follows:

$$I(x, y) = I1\left(\frac{1-x}{2}\right)\left(\frac{1-y}{2}\right) + I2\left(\frac{x+1}{2}\right)\left(\frac{1-y}{2}\right) + I3\left(\frac{x+1}{2}\right)\left(\frac{y+1}{2}\right) + I4\left(\frac{1-x}{2}\right)\left(\frac{y+1}{2}\right)$$

Note that at any detector center (i.e., any corner of the rendered area, such as $x=0.5$, $y=0.5$, etc.), the intensity matches that of the corresponding detector. As a result, the corners and edges of adjacent rendered areas always match in intensity. This continuity gives the bilinear model a significant visual advantage over the plane model. In addition, the total intensity over the area of each macrodetector is preserved.

Performance

Images have been divided into three categories to demonstrate various aspects of the performance of these algorithms. These categories are continuous tone, text/line art, and halftone. A representative image from each category is shown in its original form and as processed for each of three algorithms: BBJJ for PSDs, BBJJ for standard detectors, and gray (i.e., the image area corresponding to a given is rendered with the average intensity incident on the detector - the rendering method used in typical scanning systems [5]).

Continuous Tone Images

Continuous tone images are those in which the intensity can vary continuously over the image. This class of images is a particularly good test of the BBJJ algorithms, because it entails resolving edges of a variety of contrasts and as well accurately rendering non-edge areas. The original image is shown below:

FIGURE 19.

Original Continuous Tone Image



This image is now scanned by an array of 10 pixel by 10 pixel standard detectors, with the average intensity for each detector rendered over that detector (gray model), resulting in the following image.

FIGURE 20.


Continuous Tone Image - Gray Model for 300 dpi Standard Detector Array



The BBJ for standard detectors is then applied to this input:

FIGURE 21.

Continuous Tone Image - BBJJ algorithm for 300 dpi Standard Detector Array



The same image is now shown for the BBJJ algorithm applied to a PSD array of the same resolution:

FIGURE 22.

Continuous Tone Image - BBJJ algorithm for 300 dpi PSD Array



To provide a fair comparison with regards to the amount of information available to the PSD, the gray and standard detector BBJJ are shown for a standard detector scanning at twice the resolution:

FIGURE 23. Continuous Tone Image - Gray Model for 600 dpi Standard Detector Array



FIGURE 24. Continuous Tone Image - BBJJ algorithm for 600 dpi pixel Standard Detector Array



Text and Line Art Images

A second important class of images is text and line art. This type of image consists of relatively thin, high-frequency features on a background of constant or slowly varying intensity. A sample of characters is shown processed by both versions of the BBJJ algorithm at a variety of resolutions. The spots-per-inch (spi) figure given is for a 10-point font. For a 20 point font, the spi would be halved, for a 5 point font, it would be doubled.

An important consideration is that the standard grayscale model often seems to work remarkably well on text images, even picking out some edges in certain characters. This occurs because a significant number of the edges encountered in text are either horizontal or vertical with respect to the page. As a result, they often fall along the lines of the standard detector. Edges then appear sharp since they follow the edges of iso-intensity

detectors. This effect quickly degenerates when the image is skewed, since the majority of edges are then no longer horizontal or vertical.

FIGURE 25.

Original Text Sample


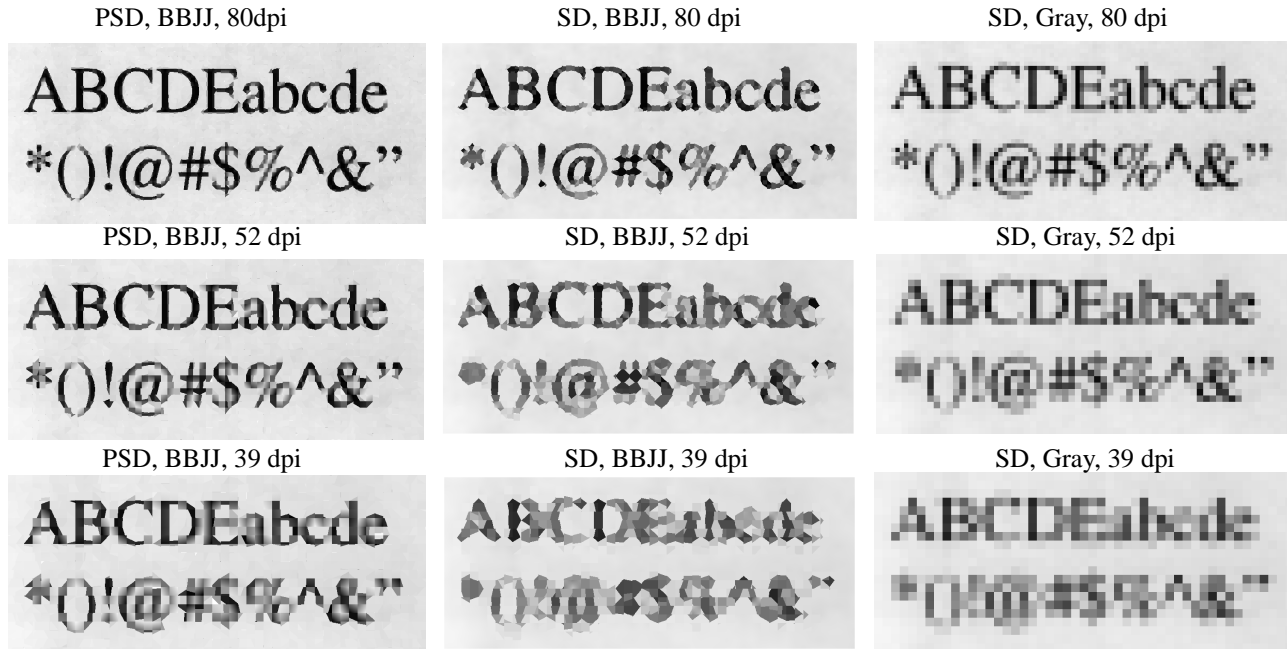


FIGURE 26. Text Sample Processed at a Variety of Resolutions



FIGURE 27. Text Sample Processed at a Variety of Resolutions (continued)



PSD: Position Sensitive Detector
SD: Standard Detector

Halftoned Images

With halftone images, there are three regions of performance with regard to the scanner spot to printer spot (s:p) ratio. The first region is where the scanner spot size is on the order of the printer spot size, where the ratio ranges from about one half to two. In this region, any gray-preserving algorithm reproduces the visually perceptible characteristics of the original image quite accurately. However, the BBJJ algorithm (for the PSD) has the advantage of being able to resolve individual halftone dots at this resolution. Though the effect of this is not readily visible if the image is rendered in gray, it can be binarized (converted from gray into black and white) without significant loss of quality. The gray model cannot be binarized in this manner without a serious loss in quality. The gray and binarized version of each algorithm are shown below for an image with a printer spot size of 4 pixels and a scanner spot size of 7 pixels. Note, however, that the effective s:p ratio for *macrodetectors* in the standard detector BBJJ case is twice that for single detectors (i.e., 14:4 instead of 7:4).

FIGURE 28. Original Halftoned Image




FIGURE 29. Grayscale Results of Algorithms on Halftoned Image (2:1 s:p ratio)






FIGURE 30. Binarized Results of Algorithms on Halftoned Image (2:1 s:p ratio)



In the second region (shown below), the scanner spot is three to five times the printer spot size. This is the most volatile of the three regions in that slight changes in the scanner spot size or the alignment of the image can produce severe Moire (aliasing) patterns (see Hall [5] and Rozenfield [12] for descriptions and explanations of this effect). This result occurs in all three of the algorithms, and is not any more pronounced for the BBJJ


algorithm than with the simple gray model. The result of each algorithm with a scanner spot size of 15 pixels is shown below.

FIGURE 31. Grayscale Results of Algorithms on Halftoned Image (4:1 s:p ratio)



The third region (shown below) is where the scanner spot is about an order of magnitude or more times greater than the printer spot size. At this resolution, at least one entire halftone cell is under each detector, resulting in the average gray value of the detector being the original gray value that was rendered with the halftone distribution. The image looks increasingly like a continuous tone image as the s:p ratio increases, and the performance goes up accordingly. The BBJJ algorithm is quite effective in this region because it can find edges between the regions of these halftoned grays, just as it could with continuous tone images. The result of each algorithm with a scanner spot size of 30 pixels is shown below. Note that the blurring effect in the border of the standard detector BBJJ case occurs because there are no detectors outside this region with which a macrodetector can be composed. As a result, a half-detector width region on each border of the image is rendered using the non-edge model. The effect is particularly pronounced here because the area of a detector is quite large with respect to the image size.

FIGURE 32. Grayscale Results of Algorithms on Halftoned Image (7:1 s:p ratio)



Discussion

It is worth reviewing at this point the major goal of hyperacuity sensing - the correct rendering of *visually significant* features of the original image. It is not possible to re-render a 600 spi checkerboard with the output from a 200 spi PSD array with the BBJJ algorithm, nor is it necessary in order to satisfy the goal above - the algorithm would produce the same uniform gray that the human eye would see from a normal viewing distance. This goal is not sufficient for all applications of image scanning. For example, the content of an image in which the halftoning pattern itself contains encoded information cannot be retrieved. However, for the vast majority of images that are scanned, it is visual significance that is crucial and not pixel for pixel accuracy. In other words, if text is legible on a page, it should still be legible in the scanned image, but if there is a stripe that appears gray which is actually rendered with a high-resolution halftone on the page, it is sufficient to render it with a single gray tone. With this expanded view of the standard in mind, the performance of the BBJJ algorithm can be more practically judged.

Advantages

Though the performance of the BBJJ algorithm was shown on several distinct classes of images, it is crucial to note that no parameters were set differently for the different images. In other words, if all of the images were put together onto a single page, the algorithm would have the same performance as shown above for each part of such a conglomerate image. This is because the algorithm uses only extremely local context information - within a radius of one detector/macrodetector - to determine the representation for a given detector. As opposed to traditional edge detection and resolution methods which attempt to parametrize edges spanning large areas in the image (see, for example, Haralick [6] and Hall [5]), the BBJJ algorithm only represents and renders edges on a local scale. As a result, the errors it makes are also only on a local scale. In addition, because the algorithm works in a single pass over the image, local mistakes cannot propagate outside the detector area. The single-pass aspect is also beneficial from a computational perspective, in that the time required to process an image is linearly related to the size of the image and independent of the image's complexity.

Another important aspect of the algorithm is that it is intensity-conserving on the detector level for the PSD case and on the macrodetector level for the standard detector case. In other words, the total intensity of the algorithm's output representation for a given detector/macrodetector is always equal to the total intensity incident on the detector/macrodetector. This fact can be easily verified by examining the algorithms above. Conserving the intensity is crucial in preserving the appearance of the image. At a large enough viewing distance for a given size of image area, two areas that contain the same total intensity appear the same to the eye. The output of the BBJJ algorithm thus appears very similar to the gray model at the viewing distance corresponding to the area of one detector/macrodetector, regardless of how poorly the algorithm chooses and parametrizes the edge and continuous models. The key benefit of this is that at some viewing distance, the output of the algorithm will appear indistinguishable from the original. Though details may be lost and moire patterns may appear, the conservation of gray at the detector level prevents the algorithm from making gross errors over a large region of the image.

The examples in the performance section have shown several situations where the BBJJ algorithm can reconstruct visually significant details in the image. What may not be evident, though, are the inherent benefits of extracting an edge representation from the data. The final output of the BBJJ algorithm for each detector-sized area is a set of four parameters, describing either the position, orientation, and intensity levels of an edge or the coefficients for a continuous description (either plane model or bilinear model). The fact that this is an analytic description of the image under the detector's area makes operations such as scaling and rotation simple: the representation can be held constant, while the area of the output image corresponding to each detector can be scaled, rotated, etc. As long as a transformation can be found between a point on the original location of the detector to the new coordinate system of the rendered image, the rendering can be performed on these transformed coordinates. This is not the case in the typical scanned image representation, in which only the average intensity for each detector is extracted from the scan (i.e., a grayscale representation). While edges in the image may appear smooth when each detector corresponds to one or a few pixels in the image, the square shape of each detector's area becomes apparent as the image is scaled (see, for example, Figure 20).

In addition, the edge/continuous model representation can be used to obtain "segmentation" information about an image, where segmentation refers to the process of dividing an image into several regions according to the image content (i.e., regions of text, regions of halftoned images, etc.). Because the system developed in this study did not require segmentation for its operation, it was not a primary concern. However, it was found that information resulting from the BBJJ algorithm could be quite useful in performing segmentation. It was observed that the output of the algorithm for text and line art images contained a large fraction of detectors with edges, while continuous tone images had a much smaller fraction. The fraction of edge-containing detectors for a given area of an image could thus be used to determine whether it contained text or illustrations. This information could be very useful if the BBJJ algorithm was being used as an image pre-processor for a system that did require segmentation, such as an OCR system.

Problems

There are several issues concerning the applicability of the BBJJ algorithm in scanning systems. The first of these concerns the output to scan (o:s) ratio. If the area of each detector is rendered on a group of n by n pixels in the output device, this ratio is n . Clearly, if this ratio is not at least two, the BBJJ algorithm provides no benefit whatsoever, since a sub-detector edge cannot be rendered in a single pixel of the output device. Since the output device of interest is often a printer, the ratio must be even higher, since most printers (film and color printers excepted) can produce only binary intensity for a given pixel and need to compose several pixels to produce a gray value. This means that for a significant performance advantage, the output device must have a resolution several times higher than that of the standard 300 spi scanner. This implies that the algorithm is not very useful for images scanned on a standard 300 spi scanner and printed on a 600 spi binary printer. However, printer technology is advancing rapidly in terms of both ability and cost, and much higher resolution printers will soon be readily available. This concern can also be countered, though, by considering the many additional abilities of a system using BBJJ. For instance, if the image was to be scaled by some factor, the o:s ratio would be increased by the scale factor and the details reconstructed by the algorithm would be of direct benefit in the output image. Another approach is to consider

the possibility of using cheaper, lower resolution scanners to deliver the same or better image quality using the BBJJ algorithm. This method also increases the effective o:s ratio.

The other class of problems involves limitations and artifacts produced by the BBJJ algorithm. Because the system requires clear context values to be able to correctly render edges, an image that has a very high spatial density of edges or edge-like features cannot be accurately rendered: the corresponding regions will be rendered as continuous approximations to the variation. As discussed earlier in the description of the algorithm, this will only occur when more than two edges occur within one detector/macrodetector. The scanning system thus must have a resolution such that desired features (i.e., a given size/font of text) do not exceed this constraint.

An additional problem with the algorithm arises from the fact that edges are always modeled as linear. If there is an edge of high curvature that passes through several adjacent detectors, the algorithm will produce the best linear fit for each detector to its corresponding region of the curve. However, in some cases this results in the edges being mismatched at detector boundaries. This same problem can occur along a relatively straight edge if there is noise in the image or the detector measurements. This problem can be corrected by doing a second pass which constrains the edges of neighboring edge-containing detectors to match. However, this would cause two significant problems - first, it would increase the time required to run the algorithm, and second, increase the effective context area on average to twice its current size (since all of the context of the relevant neighbors would be taken indirectly into account).

Future Directions

The most important next step to be undertaken in this study is the production of full-sized PSD arrays and the investigation of their properties. It must be determined whether the effects of noise and the system required to support such an array supersede the benefits of the BBJJ PSD algorithm.

Another interesting direction to proceed is the extraction of “macro” information from the BBJJ algorithm’s output representation. A method for image segmentation has already been presented: it has been tested in a very basic sense; it would be worthwhile to fully develop this idea and determine the quality of segmentation that can be performed. Similarly, it may be possible to deskew a text image using information from the BBJJ algorithm. As mentioned earlier in the discussion of the algorithm’s performance on text images, most fonts tend to have a large number of vertical and horizontal lines in them. As a result, the distribution of the θ parameter in edge-containing detectors should have peaks for the horizontal and vertical directions of the text. Additionally, it is known that these peaks should be 90 degrees apart from each other. The corresponding pair of peaks should thus be fairly easy to find. Because the smaller of these thetas correspond to the orientation of the text lines, the image can be deskewed by rotating it back by this angle. Both the segmentation and the deskewing algorithm are worth significant further investigation.

Conclusions

It was found through this study that hyperacuity sensing can be beneficial in reconstructing visually significant sub-detector-sized features in scanned images. While performance varied depending on the size and spatial density of image features in relation to the size of the detector elements, the hyperacuity sensing algorithm performed at least as well as the standard gray algorithm, and in most cases was able to accurately render many sub-detector sized features. In addition, it was found that composing information from four standard detectors to simulate the hyperacuity sensing properties of the PSD detectors was fairly successful. Significant improvements over the original scanned result (standard detector data) could be attained through processing by the BBJ algorithm. Since this latter approach works with the information available from existing detector elements, it could be easily applied in software to the output of current scanning systems.

References

- [1] Bloomer, J., and Abdel-Malek, A. "Visually Optimized Image Reconstruction," *SPIE vol. 1249: Proc. Human Vision and Electronic Imaging: Models, Methods, and Applications*. 12-14 Feb., 1990. pp. 330-335.
- [2] Fahle, M. "Parallel, Semi-Parallel, and Serial Processing of Visual Hyperacuity." *SPIE vol. 1249: Proc. Human Vision and Electronic Imaging: Models, Methods, and Applications*. 12-14 Feb., 1990. pp. 147-159.
- [3] Fahle, M. "Visual Hyperacuity: Parallel vs. Serial Processing." *Invest. Ophthalm. Vis. Sci. Suppl.* 28, pp. 361-365, 1987.
- [4] Green, W. B. *Digital Image Processing, 2nd ed.* Reinhold: New York, 1989, pp.118-120.
- [5] Hall, Ernest. *Computer Image Processing and Recognition*. Academic Press: New York, 1979. pp. 93-94.
- [6] Haralick, R. and Shapiro, L. "Chapter 10: Image Segmentation" and "Chapter 11: Arc Extraction and Segmentation." *Computer and Robot Vision*. Addison-Wesley: New York, 1992.
- [7] Lennie, P. and D'Zmura, M. "Mechanisms of Color Vision." *CRC Critical Reviews in Neurobiology* 3, 1988. pp. 333-400.
- [8] Marr, D., and Hildreth, E. "Theory of Edge Detection." *Proc. Soc. Lond. B.* 207, 1980.
- [9] Martinez-Uriegas, Eugenio. "Spatiotemporal Multiplexing of Chromatic and Achromatic Information in Human Vision." *SPIE vol. 1249: Proc. Human Vision and Electronic Imaging: Models, Methods, and Applications*. 12-14 Feb., 1990. pp.178-199.
- [10] Nishihara, H. K., and Poggio, T. "Hidden Cues in Random Line Stereograms," A.I. Memo No. 737. MIT: Cambridge, August 1993.
- [11] Pearlman, W. A. "A Visual System Model and a New Distortion Measure in the Context of Image Processing." *Journal Opt. Society of America, Vol. 68(3)*. March 1978. pp.374-385.
- [12] Rozenfeld, A. and Kak, A. *Digital Picture Processing*. Academic Press: New York, 1976. pp. 79-81.