Chapter 1. Introduction and Background

In the field of agriculture, it has become necessary to adopt new technology to remain globally competitive. Automating the farming process reduces labor costs and, with well-developed technology, can increase production. At many large industrial farms, it has become common practice to use automated planting and harvesting equipment [3,6,7] as well as to integrate GPS and satellite imaging capabilities into crop production and monitoring operations. While the data garnered from satellite images can be useful in determining the overall health or ripeness of a field of plants, the images have insufficient resolution to monitor smaller areas of a field or individual plants. There are numerous benefits to be gained by having the ability to monitor small areas. One such benefit that will be explored in this thesis is that crops can be harvested at their peak maturity instead of when the average maturity of the field is acceptable. This harvesting technique would create less waste from plants that are over- or under-mature at harvest time and must therefore be discarded.

1.1 The Broccoli Market

Over the past few decades broccoli has become an increasingly popular vegetable in the United States. According to the USDA, from 1975 to 2004 the amount of land used to grow broccoli increased from 49,650 to 137,900 acres [8]. This translates to approximately half a billion dollars of broccoli being sold annually in the United States. Of this broccoli, around 94% is sent to the fresh market, while only around 6% is processed. This is a huge shift from 1985, when about half of the broccoli harvest was processed. Fresh market broccoli must have a stem length around eight inches; it must also be free of any flowers and not have too much separation between the florets—hence the window to harvest fresh market broccoli is limited. Processed broccoli is chopped into smaller pieces and sold frozen or in products such as canned soup. There are clearly stringent requirements for harvesting fresh market broccoli that must be kept in sellable condition and sold at peak maturity.

Broccoli is considered a high-value crop. Each flower head may be sold for over a dollar, which means that farmers are often willing to invest more in broccoli production then they would

in low-price field crops, such as corn and wheat. Broccoli is also one of the most expensive crops to harvest because it requires a large labor force and either multiple harvesting iterations or a significant loss in produce available for the fresh market. The return price for broccoli was $7.90 per box (22 lbs.) in 2004 [15]. At this time the total cost to harvest each box of broccoli was $5.80 and $2.78 of the harvest price was from labor expenses.

1.2 Selective Broccoli Harvesting

If the broccoli is selectively harvested over multiple iterations, then the heads are harvested only when they are of marketable maturity and very little of the crop must be discarded due to the broccoli heads being under- or over-mature. However, for the selective harvest, the farm workers must analyze each head of broccoli and only harvest the heads that are of acceptable maturity. A selective harvest requires farm workers that are trained to identify mature broccoli and to harvest it. Selective harvesting also requires farm workers to re-examine each row of broccoli over a period of days or weeks so that only the plants that are viable for sale are picked. This greatly increases the cost of harvesting the broccoli.

1.3 Single-pass Broccoli Harvesting

For single-pass harvesting, the amount of labor is reduced as farm workers need only go through a field of broccoli once. However, for most broccoli varieties, fewer then 50% of the heads are of a suitable maturity for the fresh market at any one time [17], which greatly reduces the yield of the marketable produce. As single pass harvesting is still the most common type of harvesting, a good deal of research has gone into determining which varieties of broccoli mature at the most constant rate and how to further standardize the maturation time among plants. The most common production practices include producing transplant plugs in a greenhouse and then transplanting them into the field instead of direct seeding, as well as using plastic covers (plasticulture) to suppress weeds and conserve moisture. However, transplantation more then doubles the cost of planting [16] and plastic covers restricts cultivation and need to be periodically replaced. The broccoli heads must also be closely monitored so that the harvest can take place when the greatest number of heads are mature.

1.4 Overview of Research

In the fall of 2005, images were taken of the Gypsy *skata* broccoli plants at Virginia Tech's Kentland Farms, Figure 1. Because the exact nature of this research project was not then known, only thirteen of the over sixty images taken in the fall were appropriate for use in determining broccoli maturity. Of the other images, many were of insect infestation or leaf damage and the few that do contain a broccoli plant are at an angle or distance that made it impossible to gather any useful data about the maturity of the broccoli head. While the broccoli images were being recorded, an experienced broccoli researcher determined the maturity of each broccoli head. This information was documented and later used to verify the results of the maturity analysis algorithms.

Figure 1. Broccoli Field at Virginia Tech's Kentland Farm

To determine the maturity of broccoli, computer vision analysis was performed. As a starting point, it is first necessary to locate the broccoli head within the image. This was done on all of the recorded images using a combination of a Hough transform and contrast texture analysis. Of the thirteen images, seven were of plants classified as mature and six were of plants classified as immature. Texture feature based analysis techniques that could accurately distinguish between these mature and immature broccoli heads were developed. Theses

techniques were used in the creation of a software algorithm, implemented in LabView, which can determine the maturity level of a broccoli plant on a numeric scale.

In the May and June of 2006, hundreds of additional images were taken of the spring broccoli crop, Major *seminis*. Due to unfavorable weather conditions during transplantation, this crop of broccoli was left on a truck for three weeks until environmental conditions were appropriate for planting. This delay in transplantation, as well as a few days of unseasonably hot weather soon after the broccoli was planted, caused many of the plants to suffer from pre-mature floret loosening (Figure 2) and bolting, or to develop cat-eye (Figure 3), or starring patterns on the head [2]. Much of the broccoli that did not develop these conditions suffered from either non-uniform maturation of the head or did not develop to a size suitable for market [2]. Because of these conditions, almost all of the spring broccoli images were of broccoli plants that were unsuitable for market sale, and would have been culled by a qualified human harvester. These images were thus removed and the computer vision analysis was performed only on the thirteen original images.

While the small number of images do make the results found in this thesis less statistically significant, there is a better chance that the results will be able to accurately distinguish between normal broccoli plants in the future, as the data from abnormal broccoli plants will be excluded.

Figure 2. Broccoli plant with loose florets and beads

Figure 3. Cat-eye or starring pattern on broccoli head

Chapter 2. Literature Review

In the 1980's and early 1990's a plethora of research was conducted on broccoli harvesting [4, 10, 12, 13, 19]. Single- and multi-pass harvesters were developed as well as computer vision software that could distinguish, with varying levels of success, between immature, mature, and post mature broccoli. However, it does not appear that the vision software was ever integrated with a multi-pass harvester.

2.1 Broccoli Harvest Research

2.1.1 Single-pass Harvesting

In 1988, Casada *et al* [4] developed a single pass harvester that used a rotating circular saw mounted on a tractor to harvest broccoli heads. They were able to attain harvest rates of 60-90 heads/min with minimal damage to individual heads. Also in 1988, Walton and Casada evaluated assorted varieties of broccoli to determine the percentage of each variety that would be mature for a single harvest [17]. While they found that some broccoli varieties did have a majority of heads mature for a single harvest, most of the varieties had less then 50% of heads mature at any one time. The maturation rate for broccoli is heavily dependent on weather, soil conditions and other factors, so it may change greatly from year to year.

2.1.2 Selective Harvesting

Wilhoit and Vaughan [19] developed a pneumatically powered selective harvester in 1988. The harvester consists of a cylinder and blades that bend back the foliage surrounding the head of broccoli. The machine then severs the head and transports it to the field pack equipment. Shearer *et al* [10,13] also developed a selective harvester that consists of a cut-off saw mechanism on an actuated swing arm, which is mounted to the side of a tractor. The tractor operator would identify a mature broccoli plant and then swing the arm out to sever the broccoli head from the rest of the plant. In 1990, Shearer *et al* [12] ran field trials with the multi-pass selective harvester and were able to achieve a harvest rate of 40 heads/min with 80% of the heads selected by the tractor

operator being of the required maturity. While these selective harvesters show promise, they all require an operator, which increases the size of the harvester as well as the cost of labor.

2.2 Computer Vision Research

Many of the same researchers who developed mechanical harvesters in the late 1980s, developed computer vision algorithms for determining the maturity of broccoli heads in the early 1990s. In 1990, Wilhoit *et al* [20] correlated the grey-level run lengths from an image of a broccoli plant with broccoli head sizes. They found an exponential relationship between the run function (a measure of texture) of an image and the area of the broccoli head. They used the size of the broccoli head to determine if the plant was of the appropriate maturity to be harvested for fresh market sale. The images used for Wilhoit *et al*'s analysis were all taken in a highly controlled environment in an indoor laboratory. The plants were all equally front lit and all of the images were recorded at the same distance from the broccoli heads.

Soule and Sides [16] researched a device the used the light reflected from a broccoli head to determine the maturity. However, this testing did not produce acceptable results and they concluded that, due to the harsh lighting conditions of an outdoor field, the property of color reflectance from a head of broccoli is not a feature that can be used to determine maturity.

Qui and Shearer [9] used a Discrete Fourier Transform (DFT) to assess the maturity of a head of broccoli. A single grey-scale line scan was taken of each broccoli head and a DFT was performed of the resulting graph. Qui and Shearer were able to differentiate between various levels of maturity by using the average response values from various frequency bands. The maximum accuracy for distinguishing between the maturity-levels of different varieties of broccoli was found to be 85.0% using the DFT method. In 1994, Shearer *et al* [14] used the same apparatus from the DFT study to determine the maturity of broccoli using a co-occurrence texture analysis method. Again, a single grey-scale line scan was taken of each broccoli flower head for this study. A co-occurrence matrix, the mechanics of which will be explained in detail later in this paper, was created from the line scan and then texture features, proposed by Haralick *et al* [5], were applied to determine the maturity. Shearer *et al* had an 83.1% success rate for distinguishing between mature and immature broccoli heads for multiple varieties.

Both Qui *et al* and Shearer *et al* used an elaborate apparatus with adjustable cameras, lights, reflectors, and diffusers to record the images for analysis; they also recorded the images at night to reduce the effects of ambient light. Because such a setup is impractical on an actual harvester, this paper explores the use of texture analysis to distinguish between images of immature and mature broccoli heads that were taken in the field under a variety of natural lighting conditions, camera distances and angles.

Chapter 3. Broccoli Head Localization and Sizing

3.1 Initial Image Experimentation

Various image processing techniques were evaluated to determine their efficacy in locating the broccoli flower heads within an image of an entire broccoli plant. For efficiency, these techniques were reviewed using commercial software or software written by other researchers in order to quickly gauge their possible utility. After a promising method was determined, software was tailored to the specific method of detection of the mature flower heads.

3.1.1 Pattern and Color Matching

The first method of image processing used to detect individual heads of broccoli was basic pattern matching. Initially, this seemed to be a logical choice, as the broccoli flower head, to the human eye, appears to exhibit characteristics that are distinctive from the foliage that surrounds it. However, preliminary tests using pattern matching software available with the LabVIEW Vision package gave poor results

Figure 4 shows the results from the pattern matching software. The software was given a partial picture of a broccoli flower head, taken at approximately the same distance as the original image, as a template. Instructions were then given to find the three closest matches in Figure 4. As can be seen, the results were far from promising. Similar results were produced with the pattern matching software from all available broccoli images except the image from which the template was taken. On that image, the software correctly located the template but also recognized various leaves as having the same pattern. The reason for this is unknown. I speculate that the pattern of veins in the leaves, shown in more detail in Figure 5, may mimic the bunching pattern of beads in the broccoli head closely enough to confuse the software.

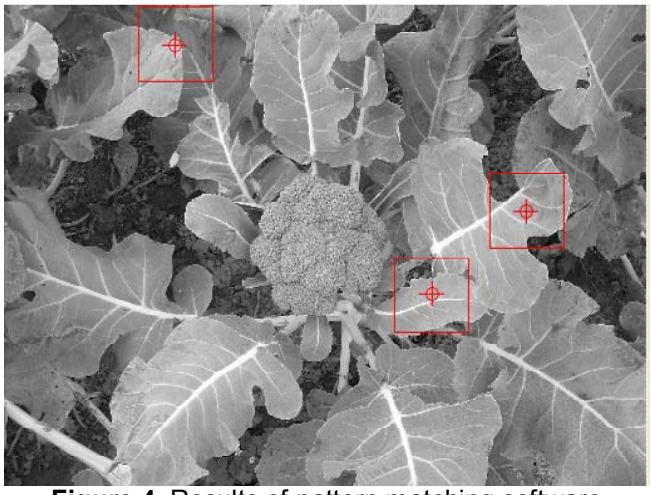

Figure 4. Results of pattern matching software used to find the broccoli flower head.

Figure 5. Close-up of broccoli leaf to show vein pattern.

Color matching was also studied using the RGB, HSI and HSV color representations. This method appeared promising with images taken under the same lighting conditions and with broccoli at the same stage of ripeness. However, as also found by Soule and Sides [16], for practical use there was too much variation in the broccoli heads from image to image to make this method a reliable way to locate the flower head.

3.1.2 Hough Transform

Preliminary tests were also run on the broccoli images using a Hough transform to extract geometric features that approximate lines. The Hough software was able to successfully detect the main stems running through each leaf. As the steams radiate out from the head in a relatively uniform pattern, using this method to locate the head appeared to be a promising solution and was thus pursued. The results of this work are detailed throughout the rest of this chapter.

3.2 Image Processing

Images of broccoli taken in the field under natural lighting conditions are not ideal for image processing. Such images are inherently noisy and the stems vary significantly from straight lines or other simple curves. To simplify the problem, a threshold and then a Canny Edge Detector were used to create a binary image which highlights the leaf stems. This allowed the direct application of a Hough transform algorithm to fit lines to the stems. After overlaying the stems with the Hough-transform-produced lines, the intersection of the lines was found and used to locate the broccoli head.

3.2.1 Thresholding

A basic threshold is one of the most simple image processing techniques. To threshold the grayscale broccoli images, all of the image pixels with intensity above a user defined threshold were set to 255, or white, and all of the pixels with intensity below the threshold were set to 0, or black. This process was used to highlight the broccoli leaf steams which, as can be seen in Figure 6, are the majority of the lightest area in a typical broccoli image. Figure 7 is the same image after the threshold has been applied.

Figure 6. Image of a broccoli plant

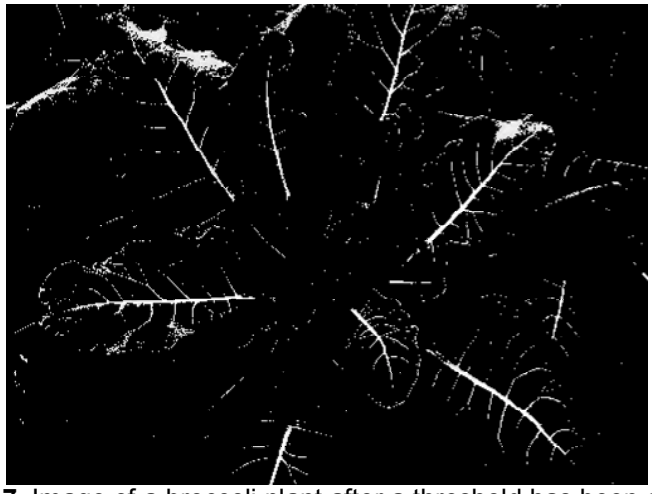

Figure 7. Image of a broccoli plant after a threshold has been applied.

However, a threshold alone left too many stray white pixels that might be misconstrued by the Hough transform as additional lines. It was therefore necessary to apply a Canny Edge Detector to the image to further highlight the long stem lines.

3.2.2 Image Convolution

A good deal of image manipulation is performed using convolution filters to enhance or suppress certain features in the image. Convolution involves multiplying each element in the convolution matrix by the corresponding element in a sub-array of the digital image. The elements are then summed and divided by the sum of the elements of the convolution matrix. The center value of the sub-array of the digitized image is then replaced with the resulting value. This process is

continued until the entire digitized image is replaced with convolved elements. Figure 8 shows a smoothing operation.

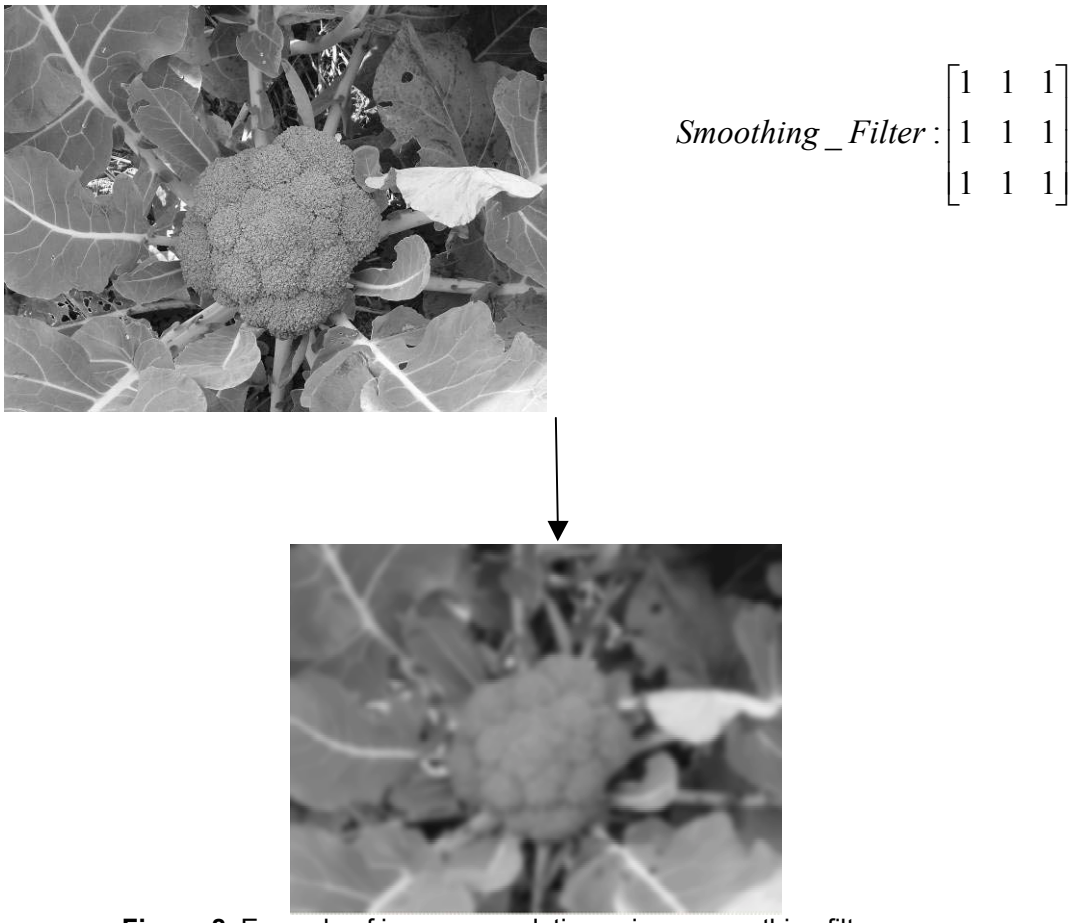

Figure 8. Example of image convolution using a smoothing filter.

Depending on the magnitude and sign of each element of the convolution matrix, edges in different directions can be extracted. For instance, if all the elements in the first column are one, the elements in the center column are zero and the elements in the third column are negative one, then vertical edges will be emphasized.

There are many different ways to handle the edges of the image when convolving it with a filter. In the example in Figure 8, the edge elements were averaged with the surrounding five elements instead of the surrounding eight elements used for interior elements.

3.2.3 Canny Edge Detection

After a threshold had been applied to the image, the image was then processed using a Canny Edge Detector, which finds prominent edges in the image. The specific goal is to find edges representing the leaf stems that have been highlighted by the threshold. The first step in a Canny Edge Detector is a Gaussian filter that smoothes the image and removes noise. The smoothing is done by convolving the image using a filter, or mask (an example of which is shown in the previous section). The size of the convolution mask is defined by the user; the larger the mask, the greater the reduction in noise. Too much smoothing may obscure non-noise edges; therefore the size of the filter mask must be carefully selected based on the application.

In the next step of the Canny Edge Detector, the gradient of the image is calculated in both the x and y directions. This is done by finding the average difference in grey-scale intensity between two neighboring pixels in the same column and the two pixels to their side for the x gradient. For the y gradient, the difference in intensity between two neighboring pixels in the same row and the two pixels above or below is found. The magnitude and angle of the gradient is found for every pixel using the formulas

$$M(i,j) = \sqrt{X(i,j)^2 + Y(i,j)^2}$$
$$\theta(i,j) = \arctan\left(\frac{X(i,j)}{Y(i,j)}\right) \quad , \qquad (1,2)$$

where X(i,j) is the x-direction gradient matrix and Y(i,j) is the y-direction gradient matrix.

The algorithm looks for the most prominent edges; therefore it selects pixels with the largest magnitude that are in the direction of the gradient as the edge points. All other pixels are set to the background intensity. This technique is called "nonmaxima suppression."

In the final step of the Canny Edge Detector, the upper and lower thresholds, defined by the user, are used to remove false edge fragments. All pixels with an intensity level higher than the upper threshold are considered to be true edges. Pixels that are connected to true edge pixels and have intensity greater than the lower threshold are maintained as edges. This way, if the noise removal obscures some of the true edges, they can still be recognized as long as they are connected to known edges and have a reasonably high intensity level.

Results from the Canny Edge Detection can be seen in Figure 9.

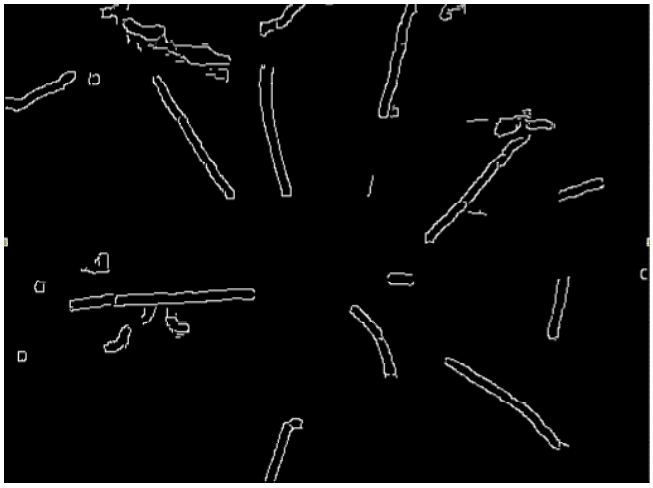

Figure 9. Image after a threshold and Canny Edge Detection

3.3 Broccoli Head Localization

The flower head was approximately located using a Hough transformation of the processed image. The Hough transform produces lines that lie along the stems of each leaf. The intersections of the lines produced by the Hough Transform are used to locate the flower head, which is then sized based on analysis of the contrast.

3.3.1 Hough Transform

A Hough line detector was applied to the image to find the stems. The stems represent most of the long line segments in the image.

A Hough transform line detector operates by finding every white pixel in a binary image and assuming that it could be part of a line going in any direction. The Hough transform was programmed using polar coordinates to avoid having to record an infinite slope for vertical lines. An accumulator array is created that has a row for every perpendicular radius from the origin that a line could be and a column for every angle from the origin at which the radial line could be oriented. Because the accumulator matrix must be of finite size, it is necessary to discretize all the possible angles and radii into the row and column values. For each white pixel, all possible

lines across the image that would go through that point are computed. The length of a perpendicular radius between the origin and a single white pixel is found. The corresponding angle between this radial line and the x-axis is also found. The accumulator is then incremented by one at the intersection of the column with the same value as the length of the perpendicular radius and the row with the same value as the angle between the radial line and the x-axis. This is repeated for every possible perpendicular radius and corresponding angles. An accumulator for one of the broccoli images is show in Figure 10. After this array has been calculated for every white pixel, a threshold is determined. The lines with the same perpendicular radius and angle as the coordinates corresponding to the elements in the accumulator which are greater than this threshold are plotted. These lines correspond to the straight lines in Figure 11.

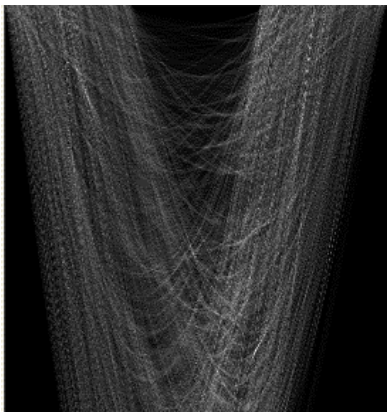

Figure 10. Display of an accumulator array.

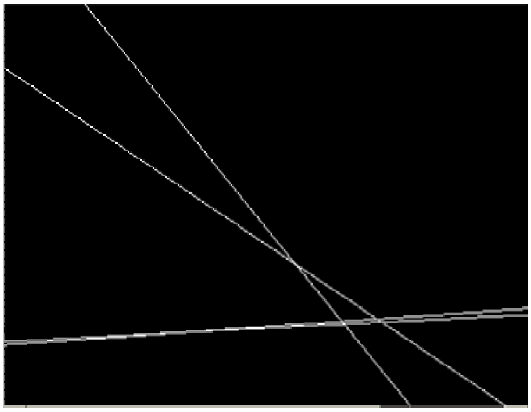

Figure 11. Lines corresponding to straight edges in an image.

As the threshold value is increased, fewer and fewer lines are produced. The computer vision program calculates the number of intersections between Hough lines and sets the threshold to restrict the number of simultaneous intersections to three. Using this method, Hough lines corresponding to the most distinct stems are preserved, while lines corresponding to random white pixels that happen to occur in a line are suppressed.

3.3.2 Flower Head Detection

The original assumption that the majority of the stems lines would cross over approximately the center of the broccoli flower head did not hold true for all of the images. Figure 12 is an image that displays the expected results: most of the stem lines intersect over the broccoli head. Figure 13 shows a more common result: the stem lines intersect around the broccoli flower head, often coming close to enclosing it.

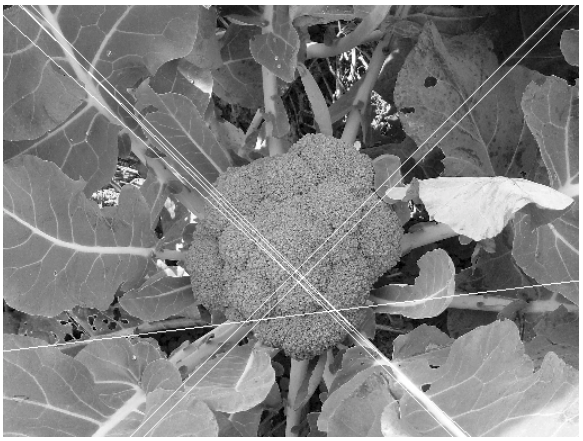

Figure 12. Expected result; stem lines intersect over broccoli flower head.

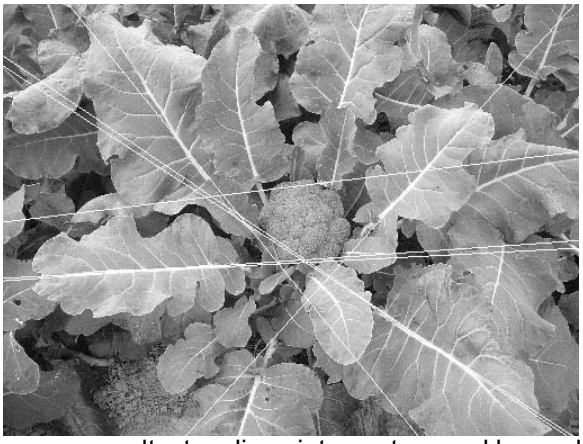

Figure 13. Common result; stem lines intersect around broccoli flower head.

Based on these results, a modification of the original flower head locating strategy was developed. An array the same size as the image was created and initialized to zero. When lines intersected during the Hough line detection, the element at the intersection point in the secondary array was incremented by one. After the Hough detection was finished, the locations of all the stem intersections were averaged and the resulting coordinates were considered to be the location of the broccoli head. Figures 14 and 15 show two examples of the results of this head location algorithm. The results from averaging the stem line intersections were reasonably accurate.

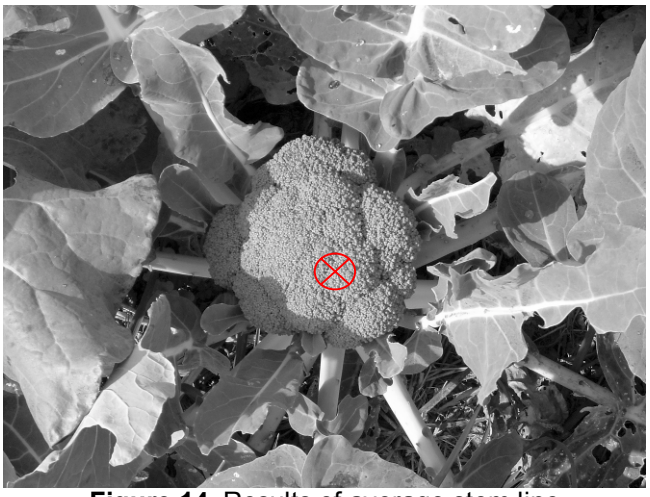

Figure 14. Results of average stem line intersection algorithm to find close broccoli head.

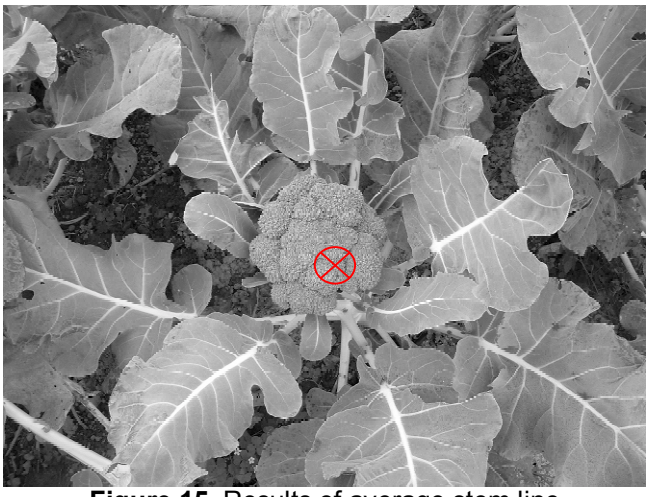

Figure 15. Results of average stem line intersection algorithm to find distant broccoli head

As can be seen from the above images, the flower head finding algorithm will work on images taken at different distances and with varying amounts of foliage. The algorithm also works for images with shadows or other lighting anomalies, as well as with some variation in flower head size, location, and angle.

3.3.3 Flower Head Sizing

After the approximate location of the center of the flower head is determined, the boundaries of the head must be found. One of the texture features referenced later in this paper, contrast, showed a marked difference between the broccoli head and the surrounding foliage and background. The contrast was found from a co-occurrence matrix, explained in section 4.2.1, of an area of interest. The co-occurrence matrix was then subjected to the contrast equation,

$$contrast = \sum_{n=0}^{N_g-1} n^2 \left\{ \sum_{i=1}^{Ng} \sum_{\substack{j=1 \\ |i-j|=n}}^{Ng} p(i,j) \right\}, \tag{3}$$

where N_g is the number of grey level intensity values and $p(i,j)$ is the value of the element in the i^{th} row and j^{th} column of the co-occurrence matrix. The equation for contrast emphasizes large intensity changes by multiplying the number of times each intensity change occurs by the square of the magnitude of that change.

The contrast was much higher for the broccoli head, as the grey scale intensity value of the pixels was constantly changing; but the grey scale intensity values of the pixels that make up the leaves and other background material stay much more constant.

The contrast values of four small (30x30 pixel) areas of interest, around the supposed center of the broccoli head as located by the Hough transform, in the image were calculated and compared to a user defined threshold contrast value. The area with the highest contrast value was defined as the starting point for the contrast measurement. If none of the areas had a contrast value greater then the user defined threshold, the head was deemed unacceptable for harvest. This situation usually occurs when the head is smaller then the 30x30 pixel area or has a very low contrast, and is therefore immature and should not be harvested.

This initial four area contrast measurement was included because occasionally the Hough transform will define the very edge of the head as the center. In this situation, the contrast is lower then the threshold because half of the area of interest is of low contrast leaf or ground.

In Figure 16, the red circled cross represents the head center, as found by the Hough transform. The four blue boxes represent the initial four areas of interest where the contrast is calculated.

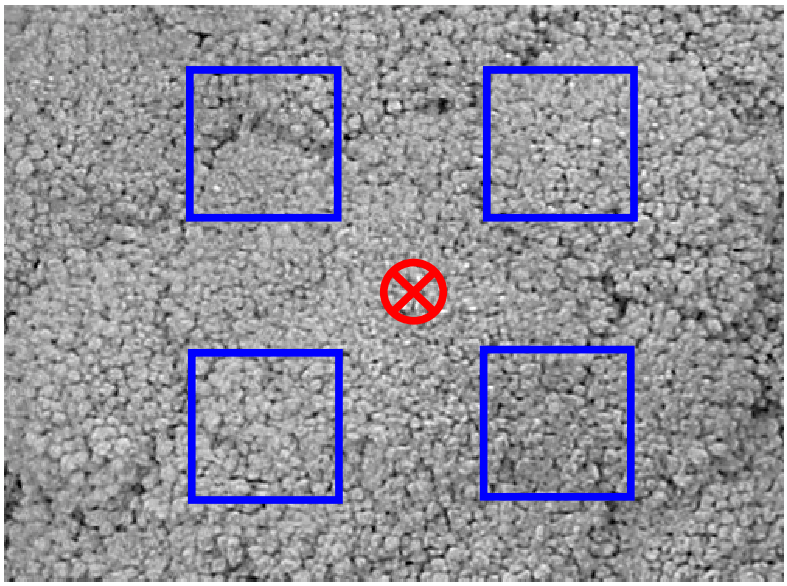

Figure 16. Locations of initial contrast analysis of the broccoli head

After the area of interest with the highest contrast has been located, four additional 30x30 pixel areas of interest at the four corners are also calculated. In Figure 17, the blue box is the initial area of interest with the highest contrast; the four surrounding orange boxes are the next four areas of interest where the contrast is calculated. If any of these four areas has a contrast that is lower then the threshold, the inner corner of that area is defined as one of the edges of the head. If an area has a contrast that is higher then the threshold, then the contrast of another 30x30 pixel area, the green boxes, at the previous area of interest's far corner is calculated. The boxes continue to emanate out from the center until each one reaches a location where the contrast is less then the threshold value.

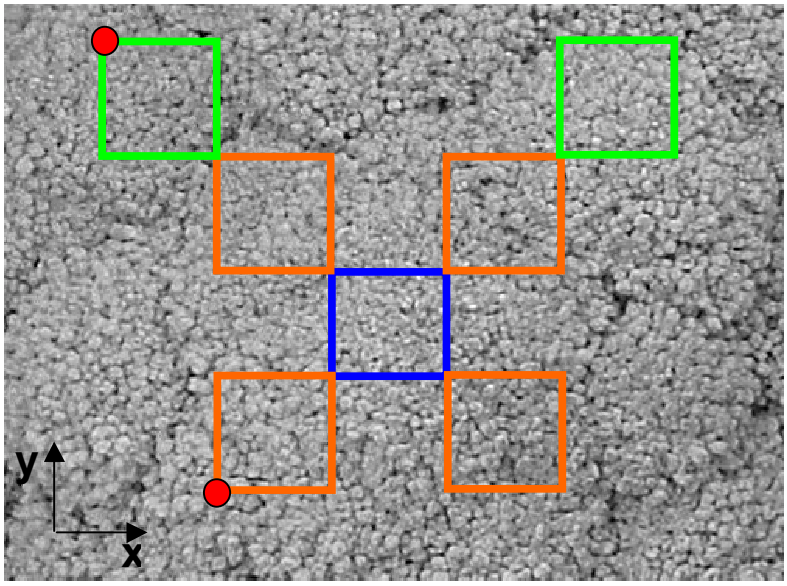

Figure 17. Areas of interest emanating out from initial high contrast area

Once the edges of the broccoli head have been reached in all four directions, the location of each edge is defined by the more conservative area of interest edge point. For instance, in Figure 17, if the two red circles represent two of the broccoli head edges then the x-coordinate of the lower circle would be defined as the left edge of the head.

Using this method it does not matter what part of the head the Hough transform program locates. As long as the assumed center point is somewhere within (or very close to) the broccoli head, an area of the head, large enough for textural analysis, should be found. Finally, a box, shown in Figure 18, is inscribed within the four boundary edges. This sub-image of the broccoli head is then used for textural analysis to determine maturity.

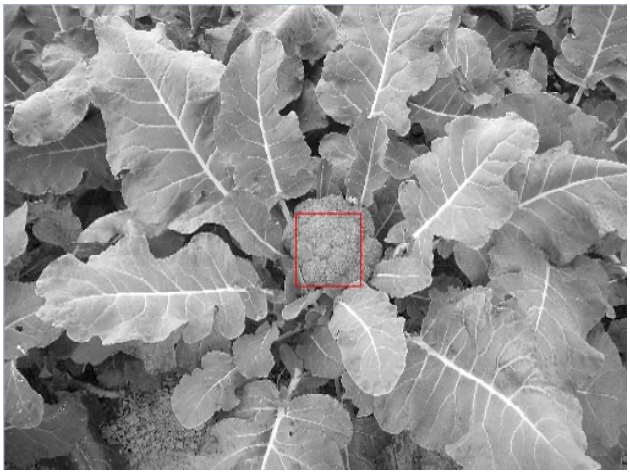

Figure 18. Approximate extent of broccoli head, as found by contrast measurement

Chapter 4. Texture Analysis from a Co-occurrence Matrix

As a broccoli head matures, the color of the head and the size and spacing of the individual florets develop in a relatively predictable manner. While the maturity of broccoli at different stages can be distinguished with little difficulty by an experienced human, determining the maturity with a computer vision algorithm presents many difficulties. Figures 19, 20 and 21 show immature, mature and post-mature broccoli heads. Although the computer vision problem is difficult, a robust algorithm may prove to be more accurate then a human at harvesting broccoli at its optimal maturity.

Figure 19. Immature broccoli head

Figure 20. Mature broccoli head

Figure 21. Over-mature broccoli head

4.1 Standard Deviation of Texture

During initial testing, it was found that when a line scan was taken from the broccoli head, the line scan of a mature broccoli head, such as the one shown in Figure 22, generally deviated more from its local mean then the line scan of the immature broccoli head, such as the one shown in Figure 23. This appears to be the result of the loosening of florets that takes place as the broccoli matures. When the broccoli is immature, the florets are tight and packed closely to one another; however, as the plant approaches harvest the florets loosen and separate as they prepare to bloom into individual flowers.

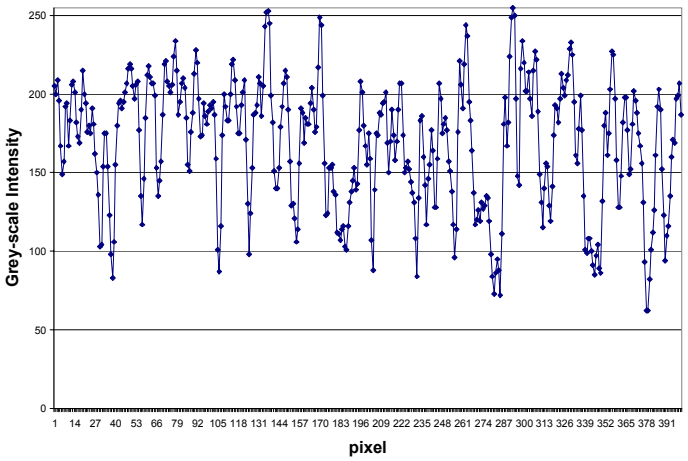

Figure 22. Grey-level intensity values of mature broccoli florets

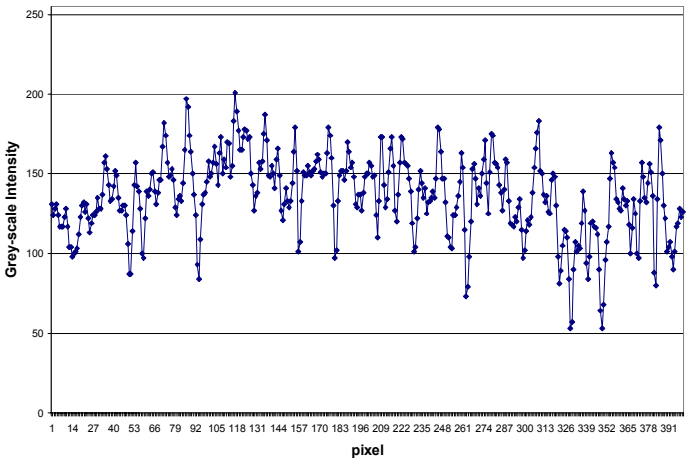

Figure 23. Grey-level intensity values of immature broccoli florets

Based on this observation, it was inferred that the standard deviation of a line scan of a mature plant would be greater then that of an immature plant. The standard deviation of every row in each of the broccoli images was calculated using the equation for standard deviation, σ,

$$\sigma = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(x_i - \bar{x})^2} \quad , \tag{4}$$

where N is the number of samples, x_i is the intensity of the current pixel and $\bar{x}$ is the average intensity of all the pixels in that row, to calculate the standard deviation of each row. The mean and the standard deviation of the standard deviations of each row were then computed. While this method showed promising results, they were not always statistically significant.

From subsequent inspection of the broccoli images, it appeared that some of the images had shadows or other variables in the light intensity that had not been taken into consideration. To compensate for these variables a dynamic standard deviation was calculated for every pixel row of every image. The value $(x_i - \bar{x})^2$ was calculated separately for each x_i using $N/2$ pixels on either side of x_i to calculate $\bar{x}$. The $(x_i - \bar{x})^2$ values were then summed, divided by N and the square root was taken to find σ.

The results of the dynamic standard deviation are tabulated in Appendix A. While N was between five and twenty pixels, the standard deviation increased significantly with every

additional pixel. However, when N was larger then approximately twenty, the standard deviation remained relatively constant as new pixels were added. As shown in Table A1, and summarized in Table 1, the mean of the standard deviation of the line scans for the mature broccoli hovered around 31 with a standard deviation of about six for $N>20$, while the standard deviation of the line scans for the immature broccoli was around 18 with a standard deviation of around three for the same conditions.

Table 1. Mean and Standard Deviation of Standard Deviations of line scan of Broccoli Heads when $N>20$

	Mature	Immature
Mean Standard Deviation	31.63	18.05
Standard Deviation	6.89	2.95

A decision boundary of $\sigma=22$ was defined (twenty-two is approximately one and a half standard deviations below the mean of the mature broccoli and approximately one and a half standard deviations above the mean of the immature broccoli) which was used to differentiate between mature and immature broccoli heads.

When the σ from all of the heads were classified using this decision boundary there was a 0% probability of error for classifying mature broccoli and a 16.7% probability of error when classifying immature broccoli. In other words, all of the mature broccoli will be correctly classified and therefore harvested, but 16.7% of the immature broccoli will also be harvested, instead of being left in the field to continue maturing.

Given these results, a dynamic standard deviation appears to distinguish relatively well between mature and immature broccoli. However, it should be taken into consideration that the entire head of broccoli does not mature at a constant rate; but the broccoli needs to be harvested before any section of the head is overly mature, and therefore an error in the favor of harvesting a less mature broccoli head is to be preferred over an error that leaves a mature head in the field.

4.2 Textural Features using a Co-occurrence Matrix

The broccoli head images were converted into grey-scale co-occurrence matrices and analyzed using texture quantization equations as defined by Haralick *et al* [5]. A numeric

maturity scale was created that defined the maturity of a broccoli head based on the texture equations.

4.2.1 Grey-scale Co-occurrence Matrix

Grey-scale co-occurrence matrices, as proposed by Haralick *et al*, were created for each of the broccoli head images at a series of displacements, ***D*****,** and at a series of angles**,** ***θ***. A co-occurrence matrix is an *NxN* matrix, where *N* is the number of grey levels being used. Similar to an accumulator matrix in a Hough Transform, each element in this matrix acts as a bin that counts the number of times a pixel with the grey-level value of the row index is a distance, D_i, and an angle, θ_i, away from a pixel with the grey-level value of the column index. Every element in the co-occurrence matrix is then divided by the total number of elements in the matrix in order to normalize it for comparison with co-occurrence matrices from other images.

The following example may help to explain the concept of the gray-scale co-occurrence matrix. In Figure 24, a possible image, converted to numeric grey-scale values is shown. Only four levels of gray are used to simplify the example. Each element in the matrix corresponds to one pixel in a grey-scale image with grey-level intensity values between zero and three. Figure 25 is the co-occurrence matrix of Figure 24 for a displacement, $D=1$ and angles, $\theta=0^o, 180^o$. Each element in Figure 24 was compared to the element horizontally to the right ($\theta=0^o$) and horizontally to the left ($\theta=180^o$). The elements in Figure 25, with row and column values corresponding to the grey-level intensity value of the original element and the right and left neighbors, were each increased by one. The co-occurrence matrix is symmetric about the diagonal because each element is compared to its neighbors on either side.

$$I(x,y) = \begin{bmatrix} 0 & 1 & 3 & 3 \\ 3 & 2 & 1 & 2 \\ 0 & 0 & 3 & 2 \\ 0 & 2 & 1 & 0 \end{bmatrix}$$

Figure 24. A matrix, representing the grey-scale intensity values of an image

$$P(i,j,D,\theta) = \begin{array}{c|cccc} & 0 & 1 & 2 & 3 \\ \hline 0 & 2 & 2 & 1 & 1 \\ 1 & 2 & 0 & 3 & 1 \\ 2 & 1 & 3 & 0 & 2 \\ 3 & 1 & 1 & 2 & 2 \end{array}$$

Figure 25. The co-occurrence matrix of Figure 24, before normalization

After the co-occurrence matrix is normalized, each element is the percentage of times that a pixel with grey-scale intensity value *i* will be a displacement, *D*, and angle, θ, from a pixel with grey-scale intensity value *j*.

4.2.2 Grey-Scale Reduction

A program was written to create co-occurrence matrices for a given image and a given vector of displacements using $\boldsymbol{\theta}$ = 0º, 90º, 180º, 270º. The image was converted into an array of grey-scale values and then reduced to the number of grey-scale values, *R,* specified by the user. The reduction is done using the equation

$$RI(x,y) = \frac{I(x,y)}{255/R}, \tag{5}$$

where *I(x,y)* is the original image matrix, 255 is the assumed number of grey-scale values in the original image and *R* is the number of grey-scale values in the reduced matrix *RI(x,y)*. The elements in *RI(x,y)* are then rounded to the nearest integer.

This reduction process reduces the number of operations necessary to create the co-occurrence matrix as well as the size of the co-occurrence matrix. It also reduces the number of operations needed to compute each of the texture features (as defined later). However, reduction in the number of grey-scale values also obscures more subtle texture patterns and creates artificially large gaps between grey-scale values, which were close in the original image, due to rounding. It is therefore necessary to weigh the efficiency of reduction against the cost. It was decided that the original texture features used to differentiate between mature and immature

broccoli would be computed using the full 255 levels of grey as there was no great constraint on time or computational power. However, when the broccoli images are being analyzed in the field, it may be necessary to reduce the number of grey-scale levels to operate in real time without slowing the harvest. More testing will need to be performed to determine the optimal blend of accuracy and efficiency.

4.2.3 Texture Analysis

Once the co-occurrence matrices were produced, they were analyzed using a series of texture analysis equations. A complete list of these equations is in Appendix B. These equations, as well as their use with the co-occurrence matrix were proposed by Haralick *et al* in the paper *Textural Features for Image Classification*. Some of these equations can characterize the grey-scale transitions and their complexity while other equations quantify specific textural features, such as contrast, structural patterns, or homogeneity. These equations were used with reasonable success by Shearer *et al* [14] to determine the maturity of broccoli based on a line scan across the head. However, as mentioned in the literature review, this research was done in 1994, when computing power was extremely limited compared to what is available today. Due to this constraint, Shearer *et al* was limited to analyzing a single line of pixels where it is now easily possible to analyze the entire head. Shearer also found it necessary to use a rather elaborate lighting system to record his images in the field at night. Because such a system would be highly impractical for use on an actual harvester, the current research proposes increasing the dimensions that were analyzed for texture as well as increasing the resolution of the image and the co-occurrence matrix. It was hoped that this would produce defining results under less controlled conditions.

The co-occurrence matrix was analyzed using the thirteen equations proposed by Haralick. The texture features, angular second moment, inverse difference moment, contrast and difference variance, were selected because they each clearly differentiated between mature and immature heads with no error. Appendix C provides numerical results from all thirteen of the texture equations.

4.2.4 Angular Second Moment

The first algorithm that showed a difference between immature and mature broccoli was the angular second moment, f_1, which is also a measure of homogeneity, and is defined by

$$f_1 = \sum_i \sum_j \{p(i,j)\}^2 , \quad (6)$$

where *p(i,j)* is the value of the element in the i^{th} row and j^{th} column in the co-occurrence matrix.

As each *p(i,j)* value is the percentage of times that a pixel, with the grey-scale value *i,* is a certain displacement from another pixel, with a grey-scale value of *j,* the angular second moment indicates how much regularity there is in texture. For instance, if many of the pixels in an image are of a constant grey-scale intensity value, Figure 26, or if there is a highly repetitive pattern of pixel intensities, Figure 27, then the co-occurrence matrix of that image will have a large angular second moment because elements in the co-occurrence matrix that correspond to those changes in intensity values will be large. As might be expected, the angular second moment of the immature broccoli heads is large due to the low level of variation in grey-scale intensity values of the small, tightly packed florets.

However, if an image has a less repetitive texture, Figure 28, there will be fewer elements in the co-occurrence matrix with a high value, as there are fewer repetitive grey-scale intensity value changes. For a mature head, the intensity values of the image change rapidly and without constant repetitiveness as the florets are of uneven size and loosely packed with dark areas sometimes showing around their perimeter. As there is no frequently occurring intensity change, all of the elements in the co-occurrence matrix are of a low value which corresponds to a low angular second moment. Appendix D contains a comprehensive collection of basic texture images and their values for each of the four texture equations.

Figure 26. All pixels are of the same intensity. This image has a very high angular second moment

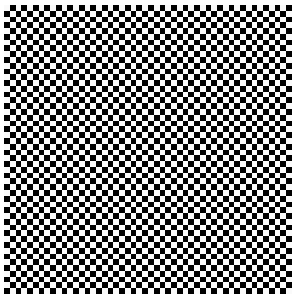

Figure 27. Pixels alternate between intensities of 256 and 125.
This image has an average angular second moment.

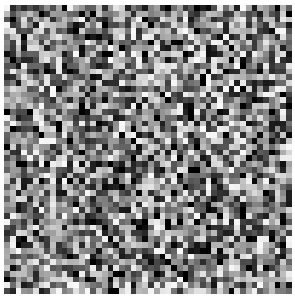

Figure 28. Pixels with intensity values between 256 and 0 are randomly scattered.
This image has a very low angular second moment.

A homogenous texture has a highly repetitive structure. As shown in Figure 30, the immature broccoli head has a very uniform and repetitive structure as the florets are tight and closely packed together. The mature broccoli head, shown in Figure 29, has larger florets that are more loosely spaced. Even though the pattern of mature florets looks regular, they are not as uniform in size or shape (spatial variation in elements) as the immature florets. This is because the broccoli head does not mature in its entirety at a constant rate and the mature florets do not have a strictly uniform size and dispersal pattern (variation in spatial arrangements). The homogeneity may be even better for distinguishing broccoli heads that are partially mature then the heads that are totally mature, as the homogeneity of the partially mature will be lower then the homogeneity of both the fully mature and the immature. But as it is necessary to harvest heads that are both fully mature and ones that have sections that are mature, homogeneity should be a good tool.

Figure 29. Grey-scale image of mature broccoli head.
Note the large, loosely spaced florets.

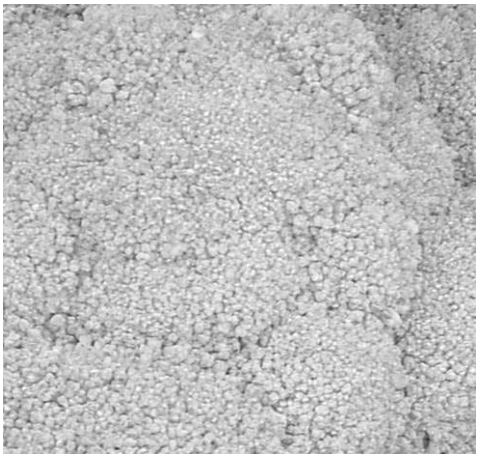

Figure 30. Grey-scale image of immature broccoli head.

4.2.5 Inverse Difference Moment

Another textural feature, the inverse difference moment, f_5, which is also a measure of homogeneity, is calculated using the equation

$$f_5 = \sum_i \sum_j \frac{p(i,j)}{1+(i-j)^2}, \qquad (7)$$

where *p(i,j)* is the value of the element in the i^{th} row and j^{th} column in the co-occurrence matrix and *i* and *j* are the index values of the row and column coordinates.

The numerator of this equation is similar to the angular second moment as it is larger with repetitive intensity value changes and smaller with more random changes. However, for broccoli head analysis, the denominator almost completely dictates the outcome of this equation. The denominator is the square of the difference between the grey-scale intensity levels (increased by one to avoid division by zero). Therefore, while the angular second moment will be larger for repetitive texture in general (such as for the image shown in Figure 31) the inverse difference moment only has a high value for texture that has small changes in intensity level (such as the image shown in Figure 32). Both measurements of texture will be high for a single intensity image, such as the one shown in Figure 33, as the texture is perfectly repetitive and has no change in intensity between pixels.

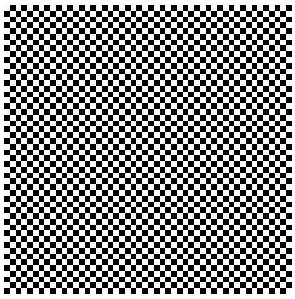

Figure 31. Pixels alternate between intensities of 0 and 255. The image has an average angular second moment but the inverse difference moment of this image is zero.

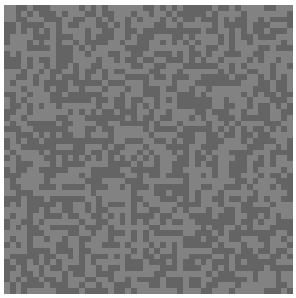

Figure 32. Pixels with intensity values of 100 and 130 are randomly scattered. This image has a low angular second moment and an average inverse difference moment.

Figure 33. All pixels are of the same intensity. This image has a high inverse difference moment and a high angular second moment.

If the florets of the mature broccoli head were all evenly spaced with the same size and grey-scale intensity values, then a large angular second moment would be expected. However, the inverse difference moment would be small unless the transition in grey-scale intensity values from the center of one floret to the center of the next happened in very small intensity value changes, so as to reduce the size of the denominator.

4.2.6 Difference Variance

The difference variance, $f_{10,}$ also consistently differentiated between mature and immature broccoli heads. The difference variance is the variance in the equation

$$p_{x-y}(k) = \sum_{i=1}^{Ng} \sum_{\substack{j=1 \\ |i-j|=k}}^{Ng} p(i,j), \tag{8}$$

where *Ng* is the total number of grey-scale intensity values. This equation creates a vector, where each element has a value corresponding to the percentage of grey-scale intensity transitions with a magnitude equal to the row index of that element. For instance, the value of element in the 10^{th} row in vector $p_{x-y}(k)$ would be the percentage of times that a pixel with intensity value x was a displacement, D, away from a pixel with an intensity value of $x \pm 10$.

The variance of $p_{x-y}(k)$, σ^2, is

$$f_{10} = \sigma^2 = \frac{1}{k}\sum_{i=1}^{k}\left(p_{x-y}(k) - \overline{p}_{x-y}\right)^2 , \tag{9}$$

where k is the number of elements in the $p_{x-y}(k)$ vector. Therefore the difference variance is a measure of how much variation there is in the magnitudes of intensity transitions. For example, if there is a fairly even spread in the magnitudes of intensity transitions, i.e. two pixels, a displacement D apart, are as likely to be 5 grey-scale levels apart as 155, then σ^2 would be relatively small: Figure 34. However, if certain magnitudes of intensity transitions occur much more frequently then other transitions then a large difference variance will be expected: Figure 35.

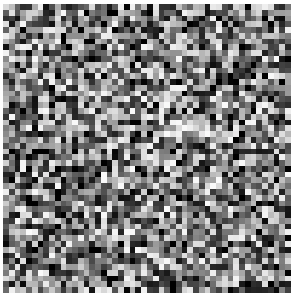

Figure 34. Pixels with intensity values between 255 and 0 are randomly scattered. This image has a very low difference moment.

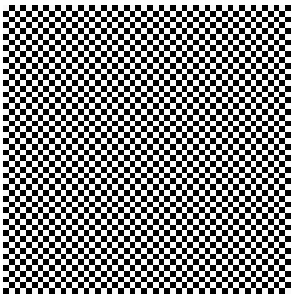

Figure 35. Pixels alternate between intensities of 0 and 255. The image has a very high difference variance, an average angular second moment and the inverse difference moment of this image is zero.

Again, the co-occurrence matrices of the immature broccoli heads have a higher difference variance then the mature heads because their florets are small and of relatively constant intensity values, so intensity transitions with small magnitudes happen frequently, while large intensity transitions are rare. The mature broccoli heads have greater variation in the magnitude of intensity transitions, so the elements in $p_{x-y}(k)$ are all closer to the mean, which creates a smaller difference variance.

4.2.7 Contrast

The last textural feature that consistently differentiated between mature and immature broccoli heads was contrast. The contrast, f_2, is represented by the equation

$$f_2 = \sum_{n=0}^{N_g-1} n^2 \left\{ \sum_{i=1}^{Ng} \sum_{\substack{j=1 \\ |i-j|=n}}^{Ng} p(i,j) \right\}, \tag{3}$$

where, in the brackets, the row and column values are only added to the summation if the absolute value of the difference between the row and column indices is equal to the current value of n.

The equation for contrast emphasizes large intensity changes by multiplying the number of times each intensity change occurs by the square of the magnitude of that change. An image with few or no intensity changes, such as the one shown in Figure 36, would have a low contrast, while an image with many large intensity changes, such as the one shown in Figure 37, would have a high contrast. Figure 38 has a small contrast as the pixel intensities change only once on each row.

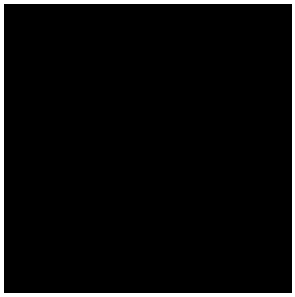

Figure 36. All pixels are of the same intensity. This image has a contrast of zero.

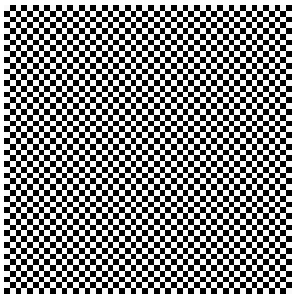

Figure 37. Pixels alternate between intensities of 255 and 125. This image has very high contrast.

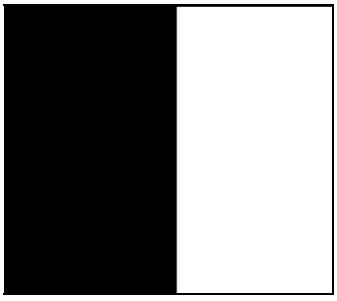

Figure 38. Half the image is 255 and the other is 0. This image has very low contrast.

The immature heads, which have a fairly constant intensity, have many low magnitude intensity changes which create a small contrast. Conversely, the mature heads have larger intensity changes, as the individual florets are much lighter then the dark space separating them and this leads to a much larger contrast.

4.2.8 Decision Rule

After the texture features, angular second moment, inverse difference moment, contrast, and difference variance were found for each of the images, it was determined that the overall greatest difference between mature and immature heads for each of the features occurred when the displacement, *D,* equaled 15. Table 2 shows the percent difference between mature and immature heads for every analyzed displacement, taking the standard deviation into account. It was first thought that if *D* was selected to be the approximate size of a floret, then the structure of the head could be found using the textural equations. However, as the florets were very rarely of a uniform size within one head, let alone among different heads, this assumption proved to be of little value.

Table 2. Percent difference between textural features of immature and mature heads

Displacement	Angular Second Moment	Contrast	Inverse Difference Moment	Difference Variance
1	4.06	1.88	2.66	1.85
3	9.87	13.93	5.05	8.94
5	13.30	17.58	11.68	16.76
7	13.66	16.62	11.36	17.54
10	13.66	16.62	11.36	17.54
12	14.12	15.39	11.37	17.98
15	15.00	14.65	10.82	18.43

Using a small displacement did not always produce significant differences in the four textural features due to the variations in floret size. For each image, depending on the floret sizes of the head, a small displacement ($D<10$) could be comparing the intensities of two pixels on the same floret, the same location on different florets, or comparing a pixel on a floret with a dark pixel showing the separation between florets. Because the florets varied in size and placement within a single head, using a large displacement for the co-occurrence matrix means that the pixels that are being compared will be a random sampling of elements from the heads.

In other words, if a small displacement is used to create the co-occurrence matrices, one co-occurrence matrix may represent the intensity change on individual florets while another co-occurrence matrix represents the intensity changes between florets. Such a comparison would be meaningless. Having a displacement that is larger then any obvious structure in the head ensures that the overall texture of the head is being analyzed for comparison instead of localized texture segments. It was impractical to use a displacement larger then fifteen due to the nature of the co-

occurrence matrix. To calculate the co-occurrence matrix a border the size of the displacement must be disregarded from each edge because pixels are being compared to pixels a displacement away. A displacement larger then fifteen resulted in sizeable percentage of the head image being ignored.

Figure 39 is a three dimensional plot that compares mature and immature heads from the fall broccoli images using each of the four selected features. While a linear decision boundary can be drawn in each dimension to separate the two classes, it is obvious that there are not discernable clusters for each class. It can be inferred that, as the broccoli head matures, its textural features slowly moves from a low contrast with large values of the other three features to a high contrast with smaller values for the other features. In Figure 39 the color of each point indicates its contrast; red is the highest, violet is the lowest.

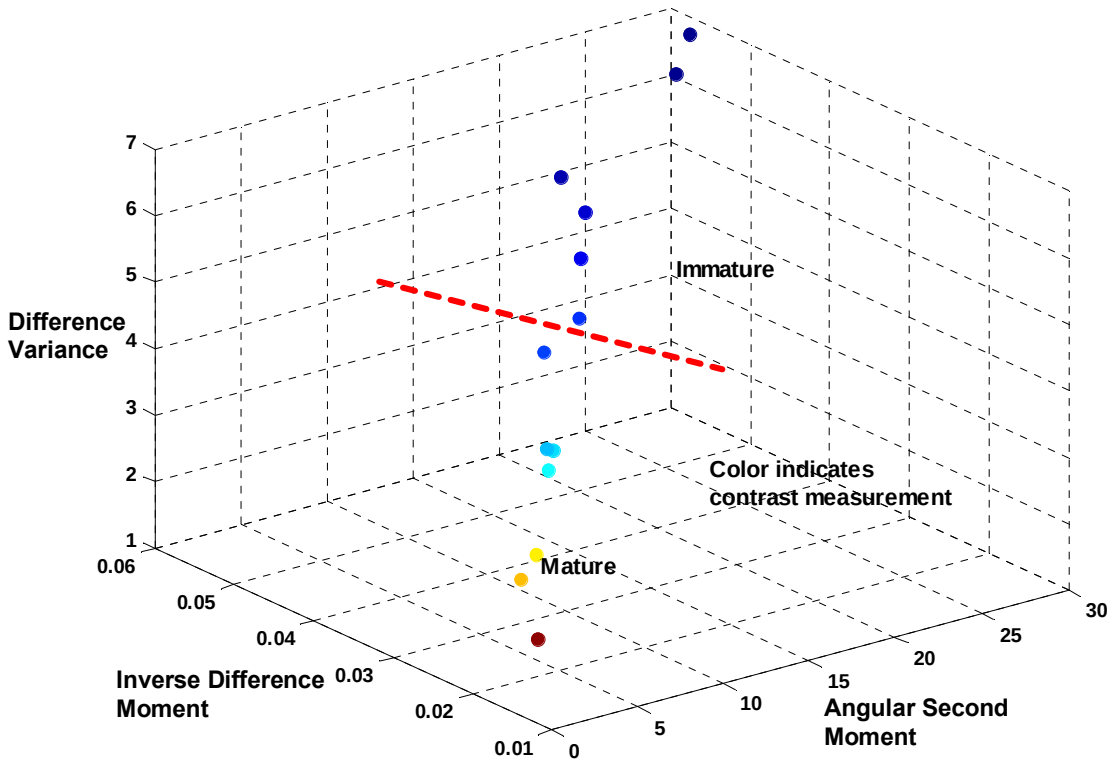

Figure 39. Three dimensional plot showing the linear separability of mature and immature heads by the four texture features.

4.2.9 Maturity Scale

The linearity of the sample measurements in the three dimensional space is what led to the creation of a numeric scale to indicate maturity. A value of zero on the scale would define the most mature (over-mature) heads of broccoli, while a value of thirty on the scale would define the least mature head of broccoli. The values for the maturity scale were determined with the use of the basic texture images and the broccoli images analyzed for this thesis. The ranges for each of the four texture features were normalized to create the maturity equation, *M*,

$$M = 0.100731 * ASM + 122.4640 * IDM + 0.54146 * DV - \frac{100{,}000}{C}, \qquad (10)$$

where *ASM* is the angular second moment, *IDM* is the inverse difference moment, *DV* is the difference variance and *C* is the contrast of a head. This scale was defined using the images of Gypsy broccoli. Its validity with different varieties of broccoli will need to be tested and it may be necessary to alter the equation, or change it all together for different varieties.

Based on a conversation with an experienced broccoli researcher [2], the decision boundary for harvesting broccoli assessed by Equation 8 was defined as 8.67 for the broccoli at Kentland Farms. The maturity equation was completely successful at correctly classifying the thirteen broccoli images as mature or immature. Needless to say, many more broccoli samples will need to be analyzed to verify this equation. The dynamic standard deviation was not used in this equation because it did not successfully classify all of the sample broccoli heads.

The equation is versatile, as it will allow farmers to define the maturity threshold that is most appropriate for the market where they are selling broccoli. For instance, a farmer growing broccoli for a local farmer's market might choose to harvest it at its peak maturity. While a farm that will ship the broccoli across the country may choose to harvest while the broccoli is still immature so that the broccoli reaches the market at its peak maturity.

Chapter 5. Conceptual Design of an Autonomous Broccoli Harvester

As mentioned in the literature review, numerous conceptual designs, prototypes, and field ready mechanical harvesters are in existence. The capabilities of these harvesters range from basic labor-saving devices that provide a moving platform for human harvesters to deposit cut broccoli to designs for selective broccoli harvesters that only rely on humans to determine if the plant is mature. While all of these machines have their uses, none have all of the components necessary for a truly autonomous broccoli harvester.

The design for a conceptual harvester consists of three main parts: the manipulator/harvester, the platform and the broccoli storage receptacle, proposals for each will be discussed in the following sections.

5.1 Conceptual Design of Manipulator/Harvester

The first element of the design, the manipulator/harvester, is proposed to be a multi-purpose arm, shown conceptually in Figure 40. A camera mounted on the end of the arm will be used to capture both the images. The same images will be used to locate the head and to perform maturity analysis. As shown in Figure 40, the camera, a, will be mounted in the center of the harvester mechanism.

The second purpose of the arm is to cut the broccoli and transport it to the storage receptacle. To accomplish this, there is a large radius ring, b, mounted around and underneath the actual cutting mechanism. This leafbreaking ring, proposed by Wilhoit and Vaughan [19], is repeatedly moved up and down around the broccoli head to break away the plant's leaves. The manipulator/harvester arm is mounted on a telescopic joint that allows the leafbreaking ring to move up and around the plant, c. Farm workers commonly cut away the unwanted broccoli foliage using a knife or a hand-held ring similar to the ring proposed. Once the foliage is out of the way, the cutting and grasping mechanism holds the head of broccoli and severs the stem.

The cutting/grasping mechanism consists of a set of four tines, d, mounted below the camera that open and close on an automated scissor joint to grasp beneath the head of broccoli. A small, ultrasonic device is located next to the camera to determine the distance from the manipulator to the broccoli so that the tines can be closed at the correct distance. Mounted eight

inches, which is the desired length of the broccoli stem, below the tines, but powered by the same automatic joint, is the actual cutting device, e. Most of the research done on mechanical broccoli harvesters [4,10,13] assumes that either a reciprocating blade or rotating saw will be used to sever the head of broccoli. However, this type of powered cutting device adds another dimension of complexity to the harvester, therefore, a blade and plate pair, similar to the horizontal scissor blades proposed by Wilhoit and Vaughan [19], are recommended to sever the head. Some research will have to be done to determine the amount of force necessary to sever the broccoli stem as Wilhoit and Vaughan do not mention what they found to be required.

Finally, the arm is mounted on a revolute joint that allows it to rotate and deposit the broccoli head in the storage receptacle. The arm has a total of three powered joints. A telescopic joint which allows the grasper/cutter to be lowered down to the level of the broccoli, a scissor joint to close the grasper tines and cutting blade, and the revolute joint that allows the broccoli to be moved from above the plant to the storage bin.

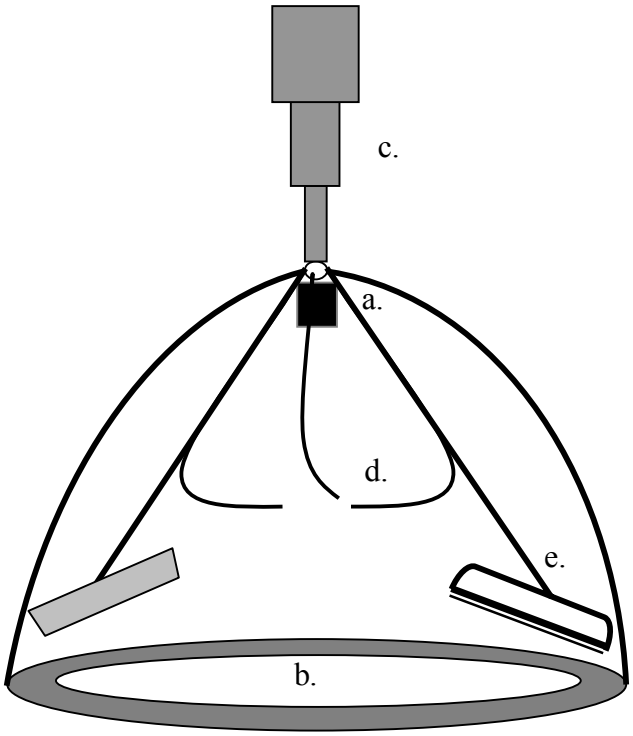

Figure 40. Conceptual design of manipulator/harvester mechanism. Note, figure is not to scale and some parts have been removed to show interior.

This configuration of joints was designed under the assumption that the distance between the row of broccoli and the autonomous harvesting vehicle remains constant throughout the harvesting process. Specifically, the distance between the row of broccoli and the harvester will be the length of the harvester/manipulator arm. This assumption is reasonable because the harvester's location relative to that of the rows must be precisely known in order for the harvester to avoid crushing any of the plants.

5.2 Platform

When dealing with autonomous vehicles, safety is always of the utmost concern as there is no human operator to use judgment in dangerous situations. If a large harvester, complete with multiple harvesting blades, was to run into a faulty bit of code or mistake an animal, or worse a human, for a plant, the results could be similar to something out of a bad science fiction movie. To prevent such a situation from occurring, many precautions will be taken to assure the safety of the harvester. First, the harvester will be capable of going no more the a few miles an hour, thus limiting its ability to chase anyone menacingly down a street. Second, the platform for the harvester will be kept as small as possible, which will limit potential damage if it does run amuck. Finally, the harvester will have a large and easily noticeable emergency stop button, so that it can be easily disabled by a non-trained operator, should it be necessary.

The conceptual platform is based loosely on Virginia Tech's Intelligent Ground Vehicle Competition (IGVC) Team's robot Johnny-5, shown in Figure 41. The platform has two driven wheels and one caster wheel, which create a kinamatically stable base. All of the wheels will be wider then normal, to afford the harvester more purchase on the loose and rocky field soil. Due to the high power requirements necessary to drive the harvester, tow the storage receptacle, operate the arm and power the automating software in addition to the necessity of the harvester operating for long periods in the field without access to electrical outlets, it will be necessary to power the harvester with either an internal combustion engine or by batteries that can be recharged in the field with a generator.

Figure 41. Virginia Tech's autonomous land vehicle, Johnny-5.
A similar design will be used for the harvester platform.

The platform will have a footprint less then 30" wide, which is the space available between rows of broccoli plants. As the harvester would severely damage any plants that it hit, accurately controlling the path of the harvester is another critical problem. If differential GPS is available in the field, it may be possible to drive the harvester through a predefined path. Otherwise, it will be necessary to use additional sensors to held guide the harvester between rows.

Figure 42. Rows of broccoli. There is approximately 30" of space between rows of fully grown broccoli plants at Kentland Farms.

5.3 Broccoli Storage Receptacle

After broccoli is harvested, it is necessary to remove the field heat from it as quickly as possible in order to preserve the vegetable. This is commonly done by packing the broccoli in ice and keeping it cool until it reaches market. As it is beyond the ability of the harvester to manipulate the broccoli head and ice, a receptacle that has refrigeration capabilities, even as simple as an insulated cart filled with ice water slush might be acceptable.

The conceptual design for the broccoli storage receptacle is a simple cart that can be towed behind the harvester. The smallest possible opening to allow the head of broccoli to enter the receptacle should be used to minimize heat loss.

Depending on the layout of the field and how much broccoli will need to be harvested every day, it may be beneficial to have the harvester automatically disengage a full receptacle and pick up an empty one at pre-specified locations.

Chapter 6. Future Research

This thesis details a solution to only one part of the problem of autonomously monitoring and harvesting crops. A great deal of research and development must be done to implement the conceptual designs presented here. In addition, there is still a plethora of research that can be done in monitoring the health and development of a broccoli plant to ensure maximum production. Aside from broccoli, there are numerous other high-value crops such as bell peppers, radishes and many types of fruit that could be made more profitable by automating the many processes associated with the plant's growth and harvest.

6.1 Additional Maturity Algorithm Research

Additionally research will need to be done in order to validate the maturity algorithm proposed in this paper. There are a number of alterations that might increase the ability of the head location algorithm to locate the broccoli head and increase the accuracy of the maturity algorithm when it determines the maturity of the broccoli plant.

One of the best methods for increasing the ability of the broccoli head location algorithm to find the head is to make the many stages of this process dynamic, so that they can better adapt to different images of broccoli plants. The first dynamic modification is in the threshold process. Instead of performing a threshold at a user defined value, the algorithm could be adapted to define the threshold based on the number of particles with shapes similar to that of a leaf stem. Basically, an original threshold of a very high value would be defined. The elongations of all of the particles above a certain size would be calculated and the number of particles with a size and elongation similar to a typical stem would be recorded. If the number of stem-like particles is fewer then expected for a broccoli plant then the threshold value is decreased and the number of stem-like particles is recomputed. This process would continue until there were an adequate number of stem-like particles in the image.

Another modification that could be made to the broccoli head location algorithm is with the use of contrast analysis to locate the edges of the head. Currently, the algorithm defines an edge of the broccoli head as soon as it locates a 30x30 pixel area of interest that has a contrast lower then the threshold. This can create a problem if the part of the broccoli head located with

the Hough transform is far from the true center of the broccoli head because the areas of interest move out from the Hough-determined head location in a diagonal fashion. As the areas of interest move outwards, they may encounter a vertical or horizontal edge of the broccoli head and define this as the true edge. However, the broccoli head may extend further in either the vertical or horizontal direction. The area that is later used for maturity analysis is defined by both the x and y coordinates of the edge of the broccoli head, so if the algorithm locates the true vertical edge but not the true horizontal edge, this will limit the area that can later used. If instead, the area of interest is moved in both vertical and horizontal directions after it reaches an area of interest with a low contrast, the extent of the broccoli head can then be more precisely determined.

Another error that occasionally occurs when using contrast to determine the size of the head is that, sporadically on the broccoli head, especially for immature plants, there will be an area with a contrast lower then the threshold value. Once the algorithm reaches this section, it will define the low contrast area as the edge of the broccoli head instead of continuing to seek the true edge. The algorithm could be modified to check the contrast of a few additional 30x30 pixel areas after it has reached the first low contrast region, if one of those additional areas of interest has a contrast above the threshold then the algorithm would continue moving outward until it found the true edge.

The texture analysis method of determining the maturity of the broccoli head should also be tested more extensively to see if it could be improved. Some of the equations used to distinguish between immature and mature heads of broccoli might be dependent on other equations and therefore not provide additional information. If this is so, the dependent equations could be removed from the maturity algorithm to decrease processing time without reducing the accuracy of the algorithm.

Finally, additional methods of determining maturity should be considered. For instance, neural networks or correlation methods might be able to determine if a plant is mature by comparing it to samples of plants known to be mature. As humans seem to determine maturity using a similar method, we “learn” what a mature plant looks like and then define plants with similar appearances as mature; these techniques may prove very useful.

6.2 Platform Prototype and Guidance Program

Now that techniques have been developed to locate the broccoli head and determine if it is mature, the next logical step will be to develop a prototype autonomous broccoli harvester such as the one described in the previous chapter. Field trials will be necessary to verify that the computer vision algorithms work as well in practice as they do with manually acquired images. As the autonomous harvester design is purely conceptual, building and implementing it will be a complex task.

It will also be necessary to program the harvester to drive up and down the rows of broccoli. There is approximately 30" of space between the rows, which limits the size of the harvester and necessitates a high degree of accuracy for the navigation algorithms.

One possibility is a combination of a laser range finder or camera (to avoid the plants) and dead reckoning using encoders to determine the harvester's relative location. This method suffers from two drawbacks. First, the broccoli has abundant foliage which will likely spill into the drive row. It will not cause undue harm to the plant for the harvester to run over these leaves, but a laser range finder or camera may be unable to identify a navigable path if there are too many leave in the way. Second, while having larger wheels will make the harvester more stable, it will still slip on the lost soil and uneven ground and this will cause an encoder error.

Differential GPS would probably have accuracy acceptable for the navigation of the slow moving harvester. However, unless the farm owner is already utilizing DGPS for some other application, the expense of operating the harvester may be prohibitive.

A good deal more research must be conducted before a reliable field navigation system is fully developed. It may be necessary to select the navigation system based on the layout and capabilities of the individual farm where the harvester is to be implemented.

6.3 Future Broccoli Research

Another aspect of broccoli monitoring that should be explored is the possibility of developing vision algorithms that can detect insect infestation and damage. This particular type

of research presents some additional difficulties due to broccoli's ability to thrive despite damage to leaves. For instance, broccoli generally has two separate growth cycles for leaves. During the first growth, up to fifty percent of the leaves can be destroyed by insects without negatively affecting the broccoli production. In fact, broccoli is shown to be most productive if approximately 22% of the first leaf growth is destroyed [1]. Therefore, an algorithm to determine if insect damage warrants pesticides will have to be able to distinguish between first growth and second growth as well as determine what percentage of the total leaf area has been destroyed.

Alternatively, the algorithm could be focused more on detecting the damage causing insects instead of determining if damage warrants pesticides. At Virginia Tech's Kentland Farms, broccoli suffers from a variety of pests, but the most damaging is the cross-striped cabbage worm; also know as the imported cabbage worm. A computer vision system that could detect this insect on broccoli leaves and distinguish it from others would allow pesticides to be applied only were this insect was present.

Another area of research is the application of color to the co-occurrence matrix. Shearer and Holmes [11] built on Haralick's [5] original concept for a co-occurrence matrix by first separating the image into the three color planes, instead of converting it to grey-scale. They then created co-occurrence matrices for each of the color planes and applied Haralick's texture equations to each matrix. While this does increase the computation time, using this method of texture analysis may produce additional useful information about the maturity of the broccoli or about other aspects of the plant's health.

6.4 Additional Applications for Harvester

While the autonomous harvester is selectively harvesting the broccoli, it is also amassing a great deal of information, specifically the state of maturity of every broccoli plant in the field. The applications for this information are numerous. Most obviously, the data could be used to select the next time to harvest based on the maturities of the plants left in the field.

The information from the harvester could also be used to help determine if any parts of the field are suffering from a lack of water or nutrition based on the size of the mature broccoli heads (from the head sizing algorithm). A contour plot, displaying the productivity of the field

during different periods of the growing seasons, would be a valuable tool to farmers when planning irrigation or fertilization systems or even when deciding where to plant crops.

Over the years, farmers could compare the development of the broccoli plants to other variables, such as weather, thereby creating more accurate predictions of broccoli yield given certain growing conditions.

6.5 Further Plant Research

Aside from developing computer vision algorithms that can detect various ailments in broccoli, it would also be valuable to develop similar algorithms for other crops. Although the specific vision algorithms developed in this paper will probably not be directly applicable to other crops; distinguishing features of other crops can be identified and used to similar ends. Work has already been done to autonomously find and harvest bright radicchio and fennel plants in the field and cucumbers in a greenhouse [6, 7]. The algorithms used for these processes do not appear to take ripeness into account, so work could be done to further develop the current algorithms and machinery.

Research could also be done to determine if crops are stressed. It has long been known that crops give off unique ultra-violet indicators when stressed. These markers are sometimes enhanced (or new ones created) via genetic manipulation. Seeing as the technology for this manipulation has already been developed, the possibilities for linking ultra-violet indicators to stress produced proteins opens the doors to the possibility of diagnosing individual plants and fields of plants simply based on the type of ultra-violet light that they emit. Computer vision and image processing with ultra-violet light images is an interesting new field to explore.

Chapter 7. Conclusion

In this thesis, research done to design both the software and hardware of an autonomous broccoli harvester has been described. While only a few pieces of this problem have been studied in detail, the basic path of what needs to take place in order to develop such a harvester has been laid out. The many aspects of autonomously monitoring, harvesting and even planting crops encompass many fields, everything from mechanical engineering to horticulture to computer vision and electrical engineering. It is hoped that this work will help to motivate additional research in the field of autonomous horticulture.

References

1. Benson, Brinkley. Personal interview. 11 July 2006.

2. Benson, Brinkley. Personal interview. 20 October 2005.

3. Benson, E., Reid, J., and Zhang, Q. "Machine Vision Based Steering System for Agriculture Combines," *2001 ASAE Annual International Meeting Presentation,* Sacramento, CA, USA, July-August 2001, Paper Number: 01-1159.

4. Casada, J. H., L. R. Walton and M. J. Bader. 1988. Single pass harvesting of broccoli. Paper No. 88-1041, St. Joseph, MI: ASAE.

5. Haralick, R.M., K. S. Shanmugan and I. Dinstein. 1973. Texture features for image classification. *IEEE Transaction on Systems, Man, and Cybernetics*. SMC-3: 610-622.

6. Henten, E., Hemming, J., Van Tuijl, B., Kornet, J., Meuleman, J., Bontsema, J., and Van Os, E. "An Autonomous Robot for Harvesting Cucumbers in Greenhouses," *Autonomous Robots*, 13: pp. 241-258, 2002.

7. Milellat, A., Reinaz, G., Foglia, M., and Genteil, A., "Computer vision applications in agriculture robotics," *11th IEEE International Conference on Mechatronics and Machine Vision in Practice,* Macau, 2004.

8. National Agricultural Statistics Service. 2005. USDA-NASS Agricultural Statistics 2005. http://www.usda.gov/nass/pubs/agr05/05-ch4.pdf June 1, 2006. p. 8.

9. Qui, W. and S.A. Shearer. 1992. Maturity assessment of broccoli using the discrete Fourier transform. *Transactions of the ASAE*. 35(6):2057-2062

10. Shearer, S. A., P.T. Jones, J. H. Casada and L. D. Swetnam. 1990. A cut-off saw mechanism for selective harvest of broccoli. ASAE Paper No. 902151. St. Joseph, MI: ASAE.

11. Shearer, S.A. and R.G. Holmes. 1990. Plant identification using the color co-occurrence matrices. Transactions of the ASAE. 33(6): 2037-2044.

12. Shearer, S.A., P.T. Jones, J.H. Casada and L.D. Swetnam. "Multiple-pass selective broccoli harvester field trial." Presented at the 1990 International Winter Meeting, ASAE Paper No. 91-1018. ASAE, 2950 Niles Rd., St. Joseph, MI 49085-9659 USA.

13. Shearer, S.A., P.T. Jones and J.H. Casada. "Development of a mechanized selective harvester for cole crops." Presented at the 1991 International Summer Meeting, ASAE Paper No. 90-1611. ASAE, 2950 Niles Rd., St. Joseph, MI 49085-9659 USA.

14. Shearer, S.A., T.F. Burks, P.T. Jones and W. Qui. “One-dimensional image texture analysis for maturity assessment of broccoli.” Presented at the 1994 International Summer Meeting, ASAE Paper No. 94-3017. ASAE, 2950 Niles Rd., St. Joseph, MI 49085-9659 USA.

15. Smith, R., W. Chaney, K. Klonsky and R. DeMoure. “Sample cost to produce fresh market broccoli.” Central Cost Region, Montoy County. 2004.

16. Soule, H. M. and S. E. Sides. 1988. Engineering aspects of the Maine broccoli industry. ASAE Paper No. 88-1577. St. Joseph, MI:ASAE.

17. Walton, L. R. and J. H. Casada. 1988. Evaluation of broccoli varieties for mechanical harvest. *Applied Engineering in Agriculture* 4(1):5-7.

18. Wilhoit, J. H. and D. H. Vaughan. 1987. Measuring broccoli stalk cutting force. ASAE Paper No. 87-6523.. St. Joseph, MI: ASAE.

19. Wilhoit, J. H. and D. H. Vaughan. 1988. A powered cutting device for selectively harvesting broccoli. ASAE Paper No. 88-1042. St. Joseph, MI: ASAE.

20. Wilhoit, J. H., M. B. Koslav, R. K. Byler and D. H. Vaughan. 1990. Broccoli head sizing using image texture analysis. *Transactions of the ASAE* 33(5):1736-1740.

Appendix A: Results from Dynamic Standard Deviation

The dynamic means and standard deviations were found for the broccoli heads using pixel run lengths between one and fifty. The results of this test are tabulated in Table A1.

Table A1. Dynamic Means and Dynamic Standard Deviations of Broccoli Heads

	Dynamic Mean Length	1	5	10	15	20	25	30	35	40	45	50
Ripe 1	Dynamic Mean	15.776	17.802	28.104	32.431	33.967	34.751	35.093	35.280	35.182	35.032	34.843
	Standard Deviation	1.017	1.088	1.990	2.424	2.558	2.508	2.538	2.533	2.555	2.533	2.495
Ripe 2	Dynamic Mean	15.385	17.349	26.281	30.042	31.952	33.097	33.799	34.190	34.482	34.640	34.730
	Standard Deviation	1.322	1.456	2.327	2.443	2.442	2.423	2.390	2.396	2.425	2.467	2.473
Ripe 3	Dynamic Mean	15.423	17.096	24.441	27.239	28.427	29.055	29.421	29.622	29.676	29.699	29.685
	Standard Deviation	1.153	1.184	1.805	2.095	2.218	2.258	2.298	2.316	2.322	2.360	2.362
Ripe 4	Dynamic Mean	17.295	19.480	30.119	35.749	38.873	40.911	42.153	42.864	43.369	43.748	43.990
	Standard Deviation	1.122	1.219	2.003	2.491	2.731	2.804	2.809	2.733	2.611	2.506	2.410
Ripe 5-a	Dynamic Mean	16.151	17.956	24.926	26.757	27.503	27.760	27.859	27.806	27.748	27.671	27.574
	Standard Deviation	1.393	1.557	2.422	2.579	2.635	2.649	2.665	2.686	2.686	2.710	2.733
Ripe 5-b	Dynamic Mean	10.368	11.721	18.560	22.293	24.265	25.419	26.126	26.558	26.863	27.061	27.197
	Standard Deviation	0.835	0.908	1.584	2.095	2.429	2.630	2.737	2.774	2.807	2.827	2.851
Ripe 6-b	Dynamic Mean	11.644	12.703	18.460	20.887	21.919	22.492	22.846	23.062	23.197	23.287	23.318
	Standard Deviation	0.915	0.907	1.489	1.736	1.784	1.802	1.839	1.871	1.883	1.875	1.861
Average	Dynamic Mean	14.577	16.301	24.413	27.914	29.558	30.498	31.042	31.340	31.502	31.591	31.620
	Standard Deviation	2.549	2.908	4.465	5.309	5.828	6.228	6.475	6.624	6.723	6.799	6.852
Unripe 1	Dynamic Mean	9.758	10.857	15.819	17.757	18.558	18.946	19.164	19.308	19.384	19.417	19.418
	Standard Deviation	1.035	1.107	1.646	1.903	1.992	2.006	2.008	2.029	2.100	2.155	2.190
Unripe 4-b	Dynamic Mean	4.982	5.523	8.740	11.514	13.038	13.895	14.446	14.812	15.091	15.268	15.430
	Standard Deviation	0.494	0.485	0.906	1.298	1.542	1.718	1.851	1.950	2.020	2.072	2.102
Unripe 5-a	Dynamic Mean	13.993	15.354	20.393	21.570	21.985	22.202	22.274	22.296	22.283	22.199	22.092
	Standard Deviation	1.683	1.828	2.670	2.691	2.717	2.737	2.740	2.823	2.872	2.823	2.759
Unripe 5-b	Dynamic Mean	7.701	8.766	14.384	17.486	18.863	19.586	20.042	20.329	20.496	20.581	20.621
	Standard Deviation	0.739	0.790	1.403	1.819	2.025	2.119	2.178	2.234	2.271	2.289	2.299
Unripe 7-b	Dynamic Mean	8.417	8.961	12.402	13.742	14.368	14.756	15.022	15.218	15.372	15.485	15.554
	Standard Deviation	0.805	0.708	0.929	0.980	0.988	0.996	1.014	1.032	1.077	1.122	1.143
Unripe 8-a	Dynamic Mean	9.025	10.097	14.011	15.142	15.577	15.889	15.968	15.979	15.882	15.760	15.545
	Standard Deviation	1.063	1.185	1.833	1.918	1.956	2.038	2.114	2.128	2.135	2.138	2.146
Average	Dynamic Mean	8.979	9.926	14.292	16.202	17.065	17.546	17.819	17.990	18.085	18.118	18.110
	Standard Deviation	2.955	3.226	3.848	3.522	3.328	3.214	3.138	3.085	3.043	3.001	2.972

Appendix B: Texture Analysis Equations from Haralick *et al* [4]

Haralick *et al* presents thirteen different texture quantifying equations for use with a co-occurrence matrix.

1. Angular Second Moment

$$f_1 = \sum_i \sum_j \{p(i,j)\}^2$$

2. Contrast:

$$f_2 = \sum_{n=0}^{N_g-1} n^2 \left\{ \sum_{i=1}^{Ng} \sum_{\substack{j=1 \\ |i-j|=n}}^{Ng} p(i,j) \right\}$$

3. Correlation

$$f_3 = \frac{\sum_i \sum_j (ij) p(i,j) - \mu_x \mu_y}{\sigma_x \sigma_y}$$

4. Sum of Square: Variance

$$f_4 = \sum_i \sum_j (i-\mu)^2 \, p(i,j)$$

5. Inverse Difference Moment

$$f_5 = \sum_i \sum_j \frac{p(i,j)}{1+(i-j)^2}$$

6. Sum Average

$$f_6 = \sum_{i=2}^{2N_g} i p_{x+y}(i)$$

7. Sum Variance

$$f_7 = \sum_{i=2}^{2N_g} (i - f_8)^2 \, p_{x+y}(i)$$

8. Sum Entropy

$$f_8 = -\sum_{i=2}^{2N_g} p_{x+y}(i) \log\{p_{x+y}(i)\}$$

9. Entropy

$$f_9 = -\sum_i \sum_j p(i,j) \log(p(i,j))$$

10. Difference Variance

$$f_{10} = \text{var}\,iance\_of\_p_{x-y}$$

11. Difference Entropy

$$f_{11} = -\sum_{i=0}^{N_g-1} p_{x-y}(i) \log\{p_{x-y}(i)\}$$

12. Information Measure of Correlation 1

$$f_{12} = \frac{HXY - HXY1}{\max\{HX, HY\}}$$

$$HXY = -\sum_i \sum_j p(i,j) \log(p(i,j))$$

$$HXY1 = -\sum_i \sum_j p(i,j) \log\{p_x(i) p_y(j)\}$$

13. Information Measure of Correlation 2

$$f_{13} = (1 - \exp[-2.0(HXY2 - HXY)])^{1/2}$$

$$HXY2 = -\sum_i \sum_j p_x(i) p_y(j) \log\{p_x(i) p_y(j)\}$$

Appendix C: Results from all Texture Analysis Equations

The results from applying the thirteen texture analysis equations, listed in Appendix B, to the co-occurrence matrix of each broccoli head are tabulated in Table C1. These results are the averaged values of the equations applied to the heads classified as mature and those classified as immature. The standard deviations of those averages are also included for each displacement and texture equation.

Table C1. Results from application of texture equations to co-occurrence matrix of broccoli head for displacements between one and thirty.

Displacement		f1	f2	f3	f4	f5	f6	f7	f8	f9	f10	f11	f12	f13
1	Mature Average	13.86	485.845	0.0556	25895	0.069	311.74	99635	5.653	9.294	9.855	3.772	3.850	0.00186
	Mature Standev	4.95	139.488	0.0308	2326	0.011	15.91	9166	0.202	0.321	2.120	0.164	0.026	0.00121
	Immature Average	37.70	232.004	0.0193	25596	0.107	314.18	98938	5.189	8.425	16.168	3.322	3.847	0
	Immature Standev	17.36	123.484	0.0048	6652	0.029	40.97	26166	0.118	0.363	4.491	0.312	0.031	0
3	Mature Average	7.99	1933.88	0.0259	25966	0.035	312.20	98594	5.505	9.811	4.028	4.473	3.957	0.00271
	Mature Standev	2.93	584.328	0.0213	2324	0.005	15.81	9204	0.229	0.343	0.961	0.131	0.012	0.00206
	Immature Average	21.60	771.801	0.0083	25644	0.057	314.50	98688	5.059	8.913	7.560	3.981	3.960	0.0005
	Immature Standev	8.54	308.334	0.0026	6656	0.015	40.94	26158	0.079	0.288	1.895	0.282	0.011	0.00055
5	Mature Average	7.50	2389.67	0.0163	26021	0.030	312.57	98388	5.463	9.865	3.395	4.624	3.969	0.00329
	Mature Standev	2.90	827.504	0.0139	2320	0.004	15.72	9213	0.225	0.363	0.889	0.141	0.007	0.00214
	Immature Average	20.07	880.29	0.0062	25691	0.051	314.82	98791	5.022	8.957	6.655	4.090	3.971	0.0005
	Immature Standev	7.00	261.8	0.0028	6659	0.010	40.92	26142	0.102	0.254	1.255	0.226	0.013	0.00055
7	Mature Average	7.38	2530.27	0.0133	26069	0.029	312.89	98443	5.453	9.877	3.244	4.658	3.972	0.00329
	Mature Standev	2.94	949.703	0.0108	2312	0.005	15.64	9196	0.223	0.373	0.934	0.166	0.006	0.00214
	Immature Average	19.76	911.338	0.0053	25732	0.049	315.09	98931	5.009	8.964	6.436	4.140	3.974	0.0005
	Immature Standev	6.74	248.821	0.0025	6661	0.009	40.89	26148	0.113	0.247	1.129	0.210	0.013	0.00055
10	Mature Average	7.33	2639.41	0.0107	26140	0.028	313.38	98622	5.441	9.881	3.150	4.694	3.973	0.00329
	Mature Standev	2.96	1056.27	0.0087	2309	0.005	15.57	9198	0.220	0.380	0.973	0.192	0.006	0.00256
	Immature Average	19.57	930.611	0.0047	25798	0.048	315.53	99178	5.000	8.965	6.307	4.154	3.975	0.0005
	Immature Standev	6.52	246.44	0.0026	6669	0.009	40.89	26176	0.114	0.243	1.050	0.182	0.012	0.00055
12	Mature Average	7.33	2681.82	0.0096	26184	0.028	313.68	98755	5.436	9.882	3.115	4.716	3.974	0.00329
	Mature Standev	2.96	1093.95	0.0079	2303	0.006	15.50	9189	0.219	0.381	0.979	0.199	0.006	0.00256

	Immature Average	19.54	941.424	0.0047	25839	0.047	315.81	99333	4.993	8.964	6.248	4.176	3.976	0.0005
	Immature Standev	6.43	245.14	0.0026	6675	0.008	40.89	26200	0.112	0.242	1.007	0.190	0.011	0.00055
15	Mature Average	7.33	2714.69	0.0089	26244	0.028	314.09	98961	5.430	9.881	3.083	4.698	3.974	0.00329
	Mature Standev	2.96	1114.07	0.0074	2298	0.006	15.43	9169	0.218	0.382	0.982	0.213	0.006	0.00256
	Immature Average	19.53	956.14	0.0042	25898	0.047	316.20	99554	4.982	8.961	6.172	4.182	3.977	0.0005
	Immature Standev	6.31	246.714	0.0021	6681	0.008	40.88	26223	0.111	0.239	0.969	0.171	0.011	0.00055
17	Mature Average	7.33	2724.52	0.0084	26285	0.028	314.36	99114	5.428	9.881	3.075	4.714	3.974	0.00329
	Mature Standev	2.96	1120.16	0.0073	2299	0.006	15.40	9163	0.218	0.382	0.980	0.205	0.006	0.00256
	Immature Average	19.52	962.765	0.0038	25938	0.047	316.47	99709	4.977	8.959	6.127	4.181	3.977	0.0005
	Immature Standev	6.27	250.203	0.0024	6685	0.007	40.88	26241	0.110	0.237	0.960	0.166	0.011	0.00055
20	Mature Average	7.33	2739.1	0.0080	26342	0.028	314.76	99325	5.424	9.878	3.057	4.715	3.974	0.00314
	Mature Standev	2.96	1135.92	0.0070	2299	0.005	15.36	9162	0.218	0.382	0.974	0.201	0.005	0.00254
	Immature Average	19.50	970.604	0.0035	25994	0.046	316.84	99926	4.970	8.957	6.078	4.203	3.978	0.00033
	Immature Standev	6.16	244.253	0.0021	6692	0.007	40.87	26268	0.112	0.235	0.911	0.178	0.011	0.00052
25	Mature Average	7.36	2756.02	0.0071	26436	0.028	315.41	99683	5.417	9.875	3.044	4.719	3.974	0.003
	Mature Standev	2.95	1146.32	0.0068	2304	0.006	15.31	9169	0.217	0.381	0.974	0.208	0.006	0.00224
	Immature Average	19.62	985.661	0.0030	26094	0.046	317.50	100315	4.958	8.951	6.005	4.207	3.978	0.00033
	Immature Standev	6.12	243.097	0.0018	6702	0.007	40.87	26312	0.110	0.233	0.879	0.165	0.011	0.00052
30	Mature Average	7.36	2783.33	0.0063	26521	0.027	315.99	99992	5.411	9.872	3.011	4.735	3.974	0.003
	Mature Standev	2.93	1152.59	0.0062	2308	0.005	15.28	9179	0.216	0.380	0.960	0.214	0.006	0.00224
	Immature Average	19.77	992.768	0.0028	26188	0.045	318.13	100684	4.947	8.942	5.966	4.219	3.979	0.00033
	Immature Standev	6.09	241.315	0.0018	6711	0.007	40.87	26346	0.107	0.231	0.873	0.159	0.011	0.00052

Appendix D: Basic Texture Images and Results

Figure D1 through D12 are the various basic texture images that were created in order to better define how the texture analysis equations quantify texture. Table D1 shows the normalized angular second moment, inverse difference moment, difference variance and contrast for each of the basic texture images.

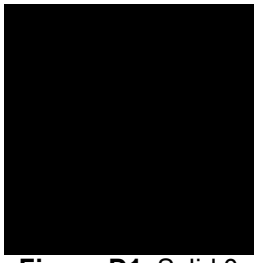

Figure D1. Solid 0

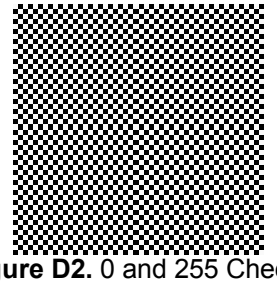

Figure D2. 0 and 255 Checks

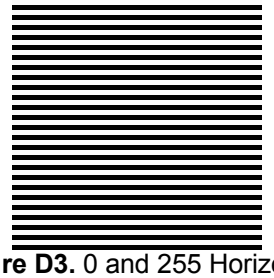

Figure D3. 0 and 255 Horizontal Stripes

Figure D4. 0 and 255 Random

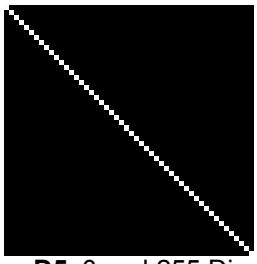

Figure D5. 0 and 255 Diagonal Stripe

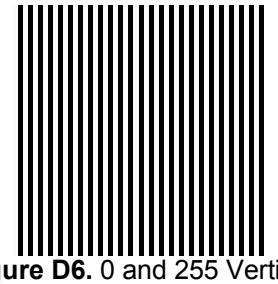

Figure D6. 0 and 255 Vertical Stripes

Figure D7. Solid 125

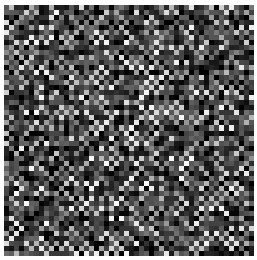

Figure D8. 0-255 Checks

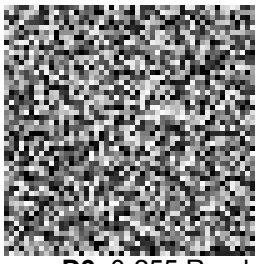

Figure D9. 0-255 Random

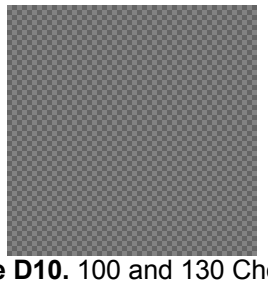

Figure D10. 100 and 130 Checks

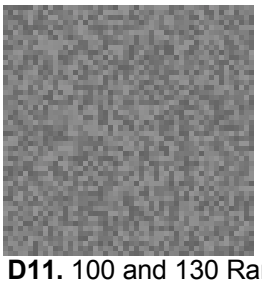

Figure D11. 100 and 130 Random

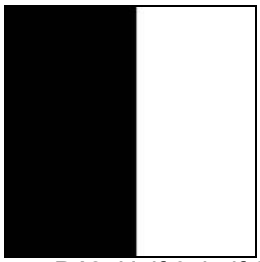

Figure D12. Half 0, half 255

Table D1. Normalized texture analysis results for basic texture images

Figure	Angular Second Moment	Inverse Difference Moment	Difference Variance	Contrast
D1. Solid 0	100.000	100.000	100.000	0.000
D2. 0 and 255 Checks	50.000	0.000	100.000	100.000
D3. 0 and 255 Horizontal Stripes	25.000	50.000	49.802	50.000
D4.0 and 255 Random	25.043	49.600	49.805	50.391
D5. 0 and 255 Diagonal Stripe	92.008	95.900	92.062	4.123
D6. 0-255 Vertical Stripes	25.000	50.000	49.802	50.000
D7. Solid 125	100.000	100.000	100.000	0.000
D8. 125-255 Checks	0.014	1.300	0.119	20.347
D9. 0-255 Random	0.013	1.200	0.163	17.158
D10. 100 and 130 Checks	50.000	0.109	99.608	1.499
D11.100 and 130 Random	25.011	49.500	49.808	0.757
D12. Half 0, half 255	94.906	98.892	97.016	0.551