

Quantification of semiconductor crystal surface features at the submicron level

Katie Ness
Electrical and Computer Engineering

Advisor: Dr. Rebecca Willett, Assistant Professor of Electrical and Computer Engineering

I. Introduction

Recent technological advances in imaging now allow surfaces to be examined more accurately at the submicron level. These developments particularly improve examination of semiconductors. Semiconductors vary in their surface textures, but until recently it has neither been possible to examine their texture directly, nor has any quantitative correlation between surface appearance and functional capabilities of a given crystal type been determined.

Semiconductor crystals grown under different conditions show significantly different surface features when imaged by an atomic force microscope (AFM). The AFM images represent surface topography with higher regions corresponding to brighter intensities. AFM images are acquired by measuring the return force experienced by a probe as it is moved along the surface. The higher the region, the more force the probe experiences as it taps the surface, thus the brighter it appears in the image. The AFM images used in this study are provided by Dr. April Brown. They are of GaN with AlGa_N Marker/GaN (referred to as D258) and Mg:Ga_N (T_g=750C)/GaN template (referred to as D259). For both D258 and D259 we have two images taken from different places in the sample. These images are shown in Figure 1.

A. Objectives

It is clear that these two sets of data have some distinct differences just by looking at them. It appears that D258 has fewer smooth regions than D259. Also, D259 has deep holes with a specific shape, while D258 has thin, linear canyons. D258 appears to have slopes that move up to a peak, then suddenly drop down while D259 is fairly flat except for the holes. Quantitative image analysis would lead to a method of distinguishing the surfaces mathematically, and the methods developed could be used to process future data sets. The purpose of this project is to explore these differences in the four study images in a quantitative way.

B. Literature Review

The wavelet transform is the most obvious choice for examining image smoothness. It is the standard transform for processing in many applications because it has a sparse representation of smooth regions in an image [1]. Thus, the magnitude and number of nonzero wavelet coefficients varies directly with the smoothness of an image. The wavelet transform naturally lends itself to smoothing using thresholding [5]. In this case, the wavelet coefficients will not be used to process the data, but instead as a quantitative measure of the smoothness of the image.

For many applications, ranging from identifying bioluminescent plankton to cancerous tumors [3], [5], it is imperative to pick out a specific region from an image and then examine its characteristics [3]. These features may include size, shape, frequency, duration, or any other number of useful distinctions. Multiple methods have been described to identify the region of an image for study. In some cases, simple thresholding will pull out the desired regions for processing [5]. In other cases, more sophisticated methods are required. These may include taking the derivative, gradient, or some other form of edge detection [5].

Image analysis may also require smoothing or removing undesired features of the image. One way of removing these unwanted regions is to use inpainting [5]. Inpainting replaces a given region with values that are similar to the surrounding region making. In most cases, this makes the studied region more representative of the rest of the image.

II. Preprocessing

Before analyzing the images quantitatively, all four images were preprocessed in several steps.

It is important to note that the D258 images have several particularly bright spots. These spots are actually reflections, and are not a part of the material. While they appear to be speckle-like noise, standard speckle removal methods such as a median filter did not effectively remove the spots. This is because the spots are larger than a few pixels across in most cases. An alternative method is to use some form of inpainting to set the values of those pixels to values that are similar to the rest of the image. Inpainting can be very complex, but one simple method is to find the boundaries of an unwanted region, and then use interpolation based on the surrounding pixels to fill in the undesirable area. The perimeter values provide the initial conditions to solve the system of equations. The Matlab command *roifill* uses such an algorithm based on using Laplace's equation:

$$\frac{\partial^2 I}{\partial x^2} + \frac{\partial^2 I}{\partial y^2} = 0,$$

where I is the region to be filled. The harmonic solutions to this linear system provide a smooth interpolation of the region.

For the D258 images, there are two types of unwanted regions. First, there is an unusual dark spot at the top of D258 Image 1. Second, there are very bright regions in both D258 images. These two images were thresholded between the values of -3 and 4.50; thus all values outside of this range were set to one. This mask describes the unwanted regions. The Matlab command *roifill* uses the boundary values of the region to be filled as initial conditions in solving Laplace's equation as stated above. Therefore it is important to ensure that these initial conditions are the desired values from the image. To achieve this, the mask was then blurred using a Gaussian low pass filter with sigma equal to 1.25 and thresholded above 0.01. To determine these values for the Gaussian filter, a nested loop was used to compare various values of sigma and the threshold value. The resulting image varied greatly. These particular values were chosen because they minimized the appearance of the unwanted regions.

In order to be able to directly compare the images, it is important for all of them to have the same intensity scale. Therefore, all of the images were renormalized using the equation

$$I_{norm} = \frac{I - \min(I)}{\max(I) - \min(I)},$$

where I is the original image. In this way all images have the same intensity scale from zero to one. Figure 1 shows the original data, and Figure 2 shows the preprocessed data.

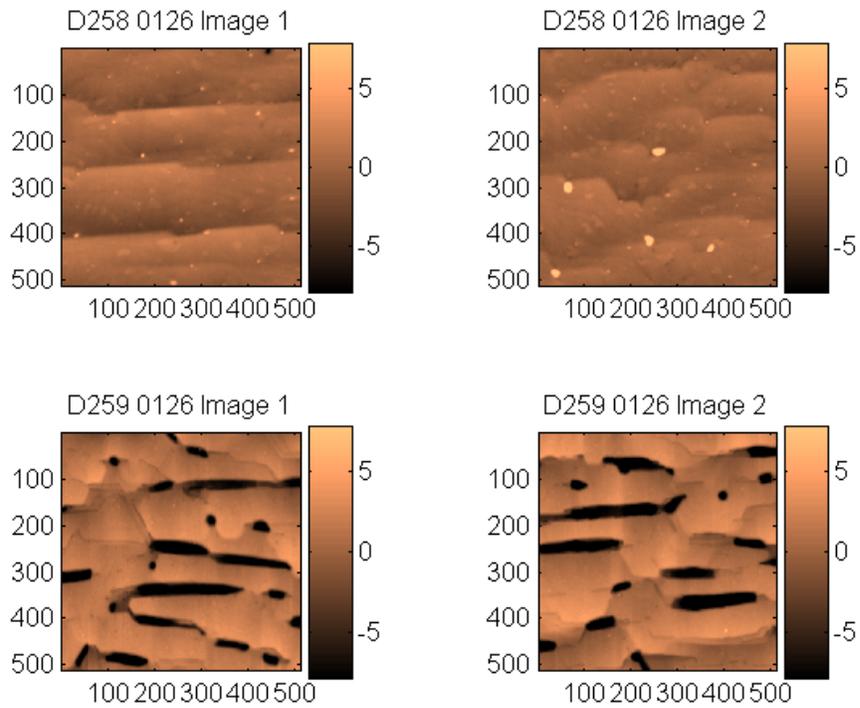


Figure 1: Original data with D258 in the top two images and D259 in the bottom two images.

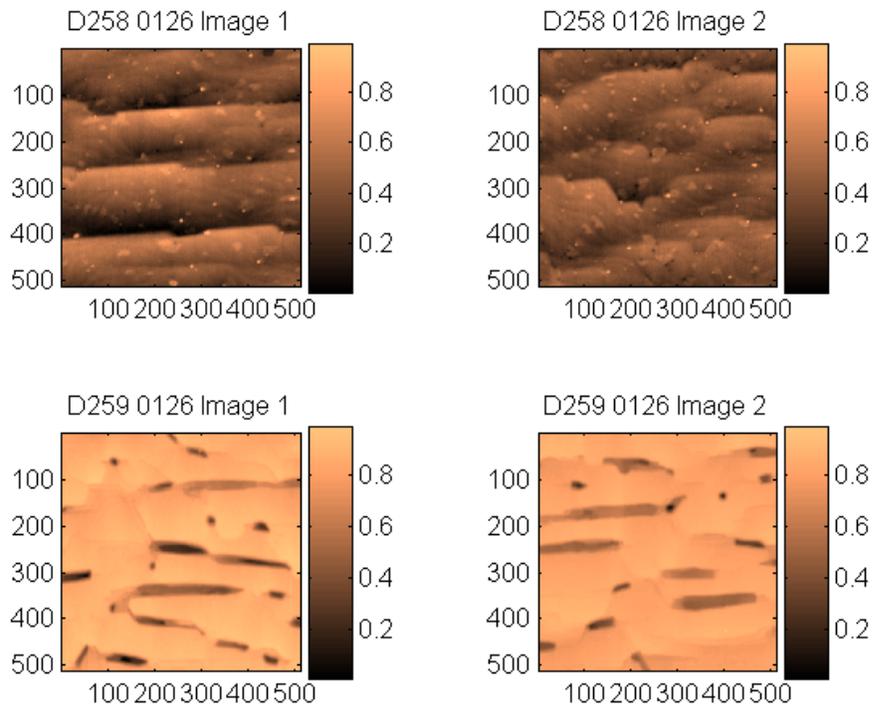


Figure 2: Preprocessed images in which bright spots have been removed and all images have been renormalized.

Comparing the originals to the preprocessed images, the bright spots have been removed from D258, and all images have intensities varying from zero to one. These images will be used as input for all later processing.

Since the images represent height at a given point, it is useful to plot them in three dimensions to visualize the surfaces. Figure 3 shows the three dimensional interpretation. Note that D258 has deep canyons and a generally rough surface while D259 has fairly flat regions that are interrupted by deep holes.

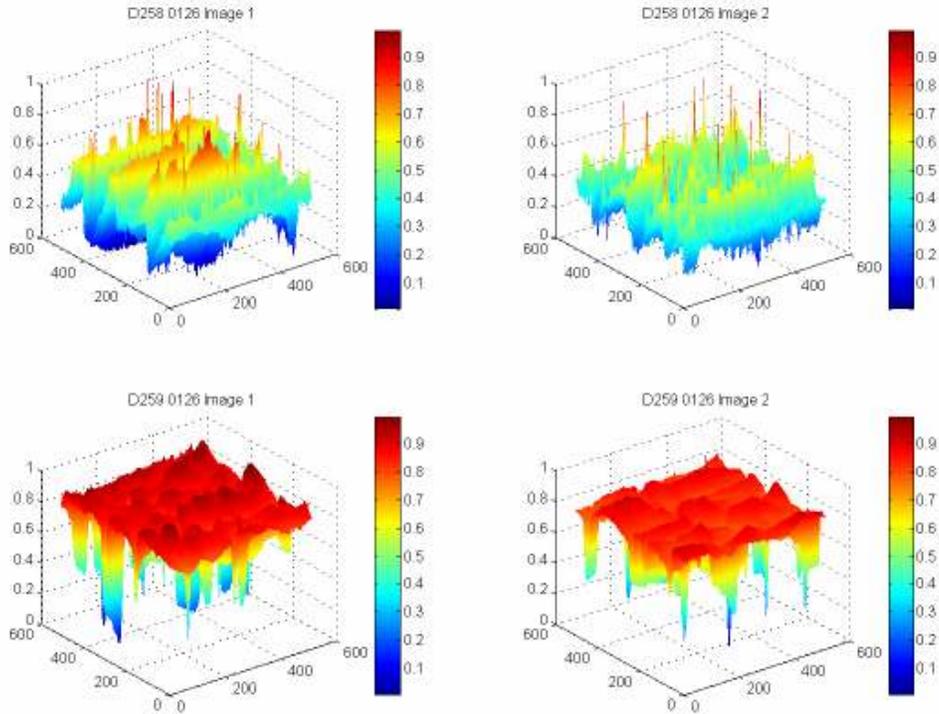


Figure 3: Three dimensional representation of the surface of the materials.

III. Smoothness

A. Expectations

For the given data sets, D258 is expected to have more nonzero coefficients and coefficients with a higher magnitude because it appears to be less smooth than D259.

B. Processing

The preprocessed data was used as input for this test. The wavelet transform was computed using a D4 Daubechies wavelet. The resulting matrices were then sorted and renormalized for easy comparison. Next, the number of coefficients greater than the threshold was calculated as the threshold was varied from zero to one. Finally, the number of coefficients versus the threshold was plotted. For easier visualization, the results are shown on a log scale.

C. Results

Figure 4 shows the plot of the number of wavelet coefficients greater than a minimum threshold. As expected, for a given threshold D258 has more coefficients than D259.

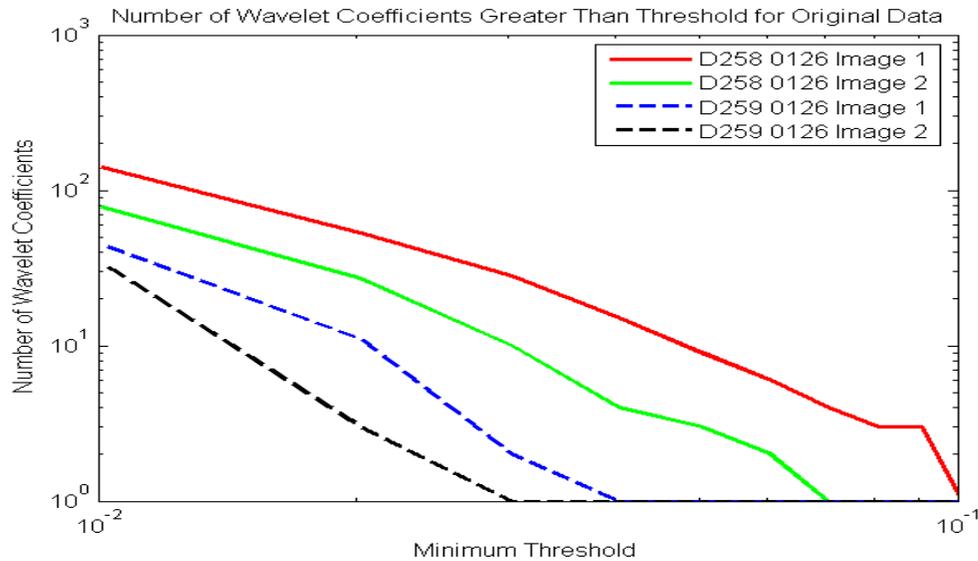


Figure 4: Number of wavelet coefficients greater than the threshold for the original data.

D. Analysis

The expected correlation is shown in Figure 4 where D258 is shown to be less smooth than D259. However, in this case it is a fairly loose correlation. For D258 there is a large difference between the two images, and the same is the case for D259. Thus it is difficult to say with certainty that there is strong evidence supporting that D258 is rougher than D259. The smoothness test has potential but does not provide a definitive quantitative difference between D258 and D259.

To explore why D258 and D259 do not have a greater difference in smoothness, several methods were used. Histogram equalization can be used to visualize the structure of the images more clearly. This process involves looking at the distribution of intensities. In this case, the intensities are divided into 1000 equally spaced bins. Figure 5(a) shows the histogram of D258 Image 1. For D258, most of the intensities fall in the middle range with a few higher and a few lower. During histogram equalization, the intensities are redistributed to make the distribution of intensities more even. The equalized histogram of D258 Image 1 is shown in Figure 5(b).

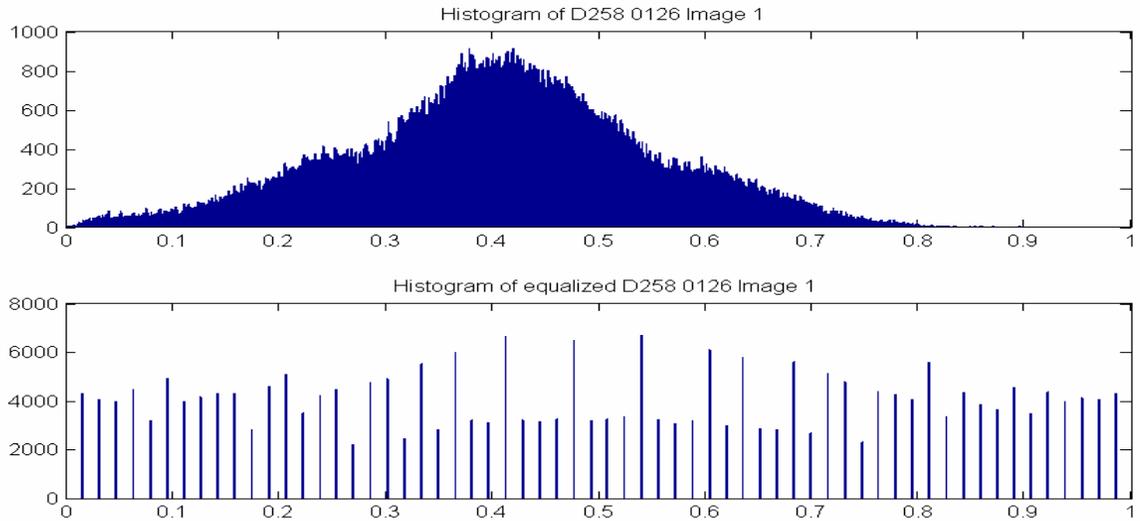


Figure 5: Histogram plots. (a)The histogram of D258 Image 1 is shown in the top plot. **(b)**The second plot shows the equalized histogram for D258 Image 1.

Figure 6 shows the histogram equalized images. During histogram equalization, the unwanted spots in D258 reappear, although not with the same intensity as in the original images. From this equalization, it is clear that D258 has many surface features. Interestingly, the flat plateaus of D259 no longer appear as flat. However, the observation that there are deep holes of a definable shape remains true for D259 even after histogram equalization.

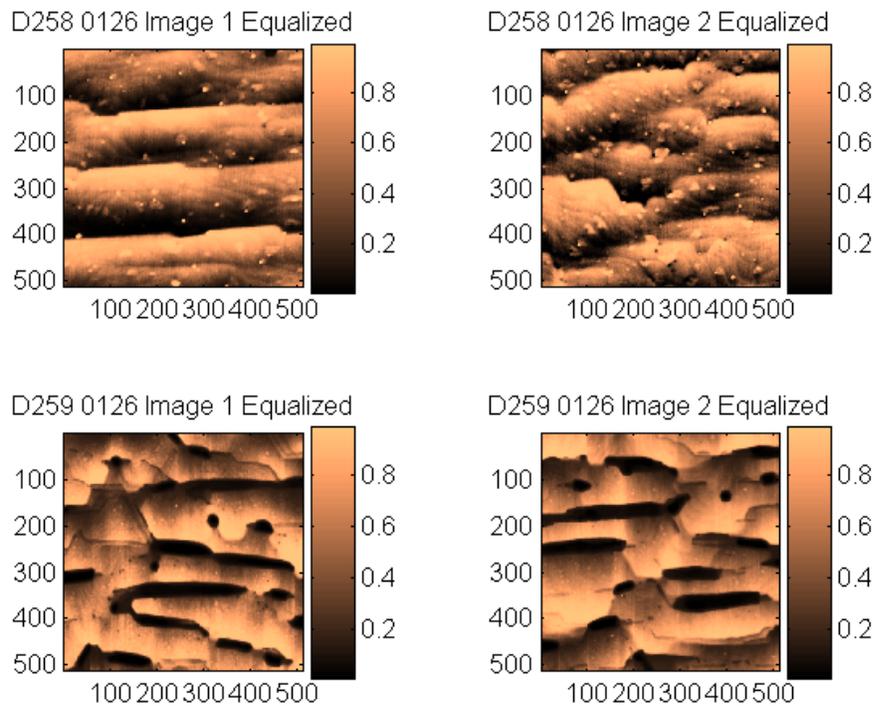


Figure 6: Histogram equalized images.

After looking at the histogram equalized images, it appears that the flat areas, particularly those in D259 may not be as smooth as they appear. Thus, it would be useful to look at the derivative of the image. From the derivative, any roughness in areas that seem to be smooth would become obvious since the derivative measures changes in intensity. In this case, the derivative was taken in the Fourier domain using the following equation:

$$G(u, v) = j(u + v)F(u, v)$$

$F(u, v)$ is the Fourier transform of the image and $G(u, v)$ is the Fourier transform of the derivative. After taking the inverse Fourier transform, the absolute value of the derivatives are shown below. By taking the absolute value, areas with a high magnitude of change will appear to be very bright while areas with small magnitudes of change will appear dark.

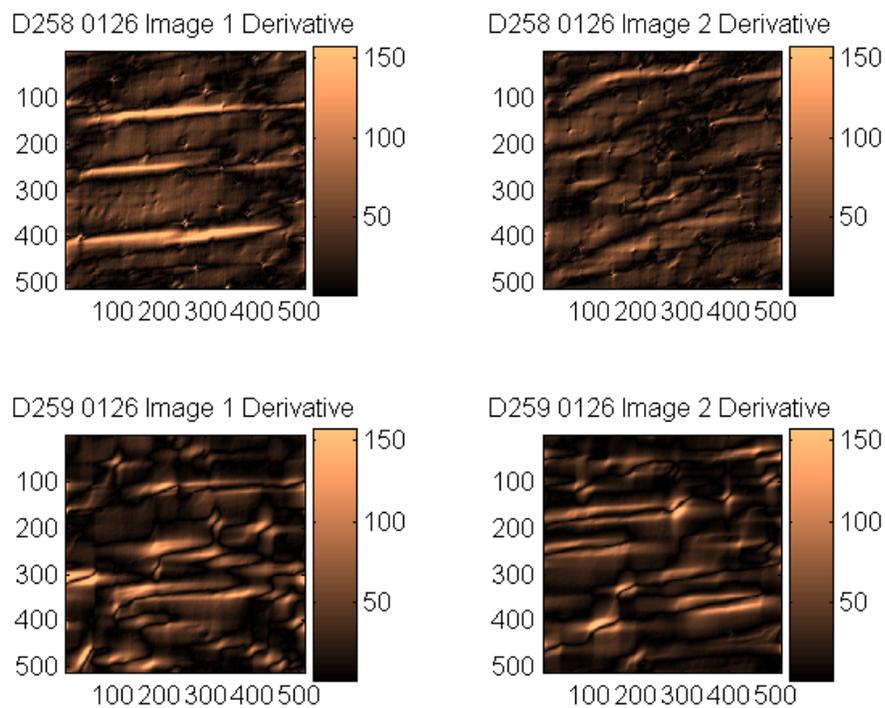


Figure 7: Absolute value of the derivative of the images calculated in the Fourier domain.

The derivatives show that D259 still has fairly flat regions, but is by no means completely smooth.

IV. Ridge Height

A. Expectations

One of the distinguishing features of D259 is its smoothness with deep holes. These dark holes have very bright regions immediately adjacent to them suggesting some sort of “ridge.” D258 has linear canyons which could also be considered ridges. Depending on how the ridge is defined, the material with the higher ridge height may vary. Based on a first-order differential, D258 should have a higher average height. For the second-order differential, D259 will have a

greater ridge height. This is because D259 has very sharp contrasts, so the magnitude of the rate of the rate of change will be very high.

B. Processing

There are several methods of determining the location of a ridge. They are all forms of edge detection which will emphasize ridges. First, a ridge could be defined as a region where the derivative is above a given threshold. In other words, a ridge is a region in which there is a high rate of change. For this test, the derivative was calculated as described previously. Second, a ridge could be defined as a region where the gradient is above a given threshold. The gradient is calculated based on differences from pixel to pixel. It is similar to the derivative, but calculated in the spatial domain. Finally, a ridge could be defined as a region where the Laplacian is above a given threshold. The Laplacian is calculated in the Fourier domain using the following equation:

$$G(u, v) = ((u - M / 2)^2 - (v - N / 2)^2)F(u, v)$$

$G(u, v)$ is the Fourier transform of the Laplacian. $F(u, v)$ is the Fourier transform of the image. M and N are the number of the rows and columns respectively of $F(u, v)$. Height is proportional to the absolute value of the derivative, gradient or Laplacian. Figures 8-10 show the absolute value of the derivative, gradient, and Laplacian of each image respectively.

A set of masks is created by varying the threshold from zero to the maximum value of each image. The mask is multiplied by the appropriate input (the derivative, gradient, or Laplacian image) which gives the magnitude of the heights and the average of these magnitudes is stored. Then the average height is plotted versus the threshold. In the graphs, the threshold is the percentage of the maximum of the image.

C. Results

The average height versus threshold is shown in Figures 8-10.

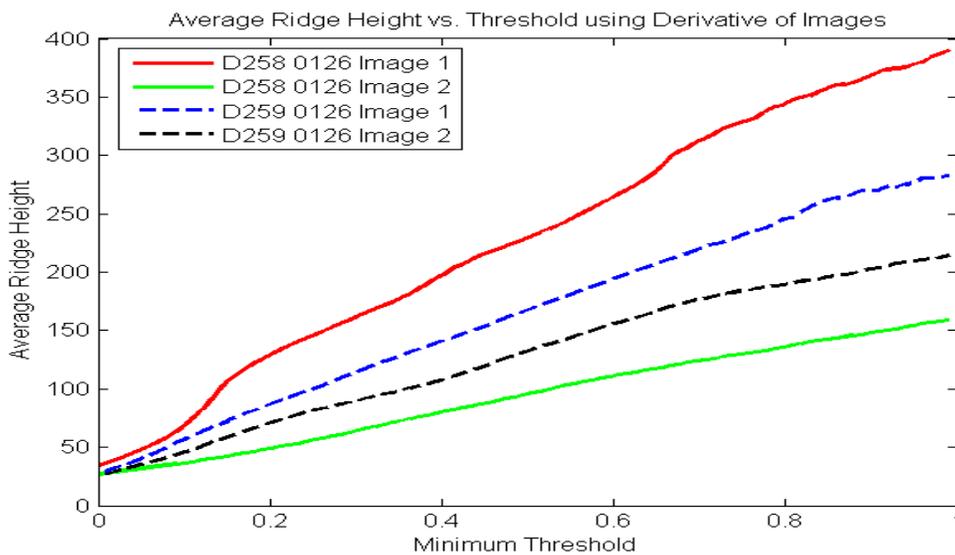


Figure 8: The average ridge height determined using the derivative of the images calculated in the Fourier domain.

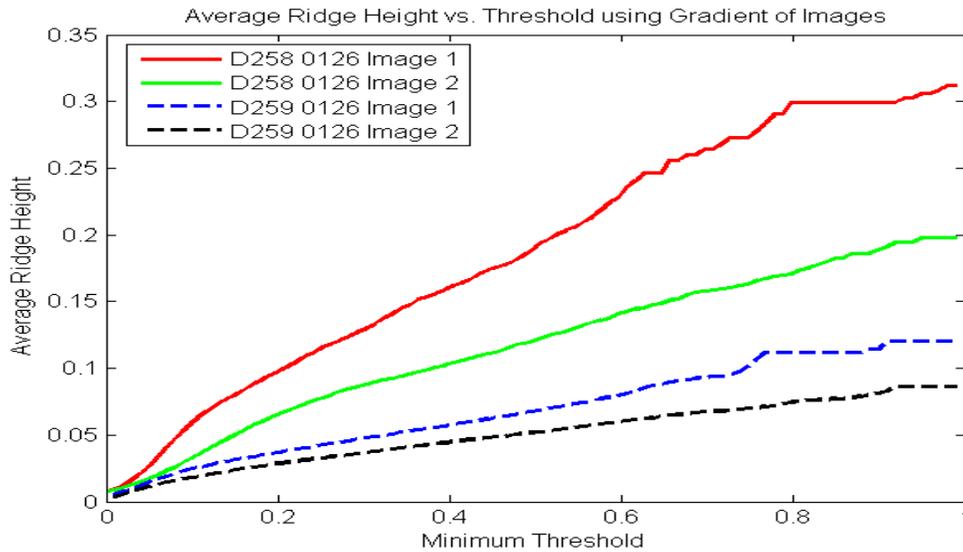


Figure 9: Average ridge height based on the gradient of the images.

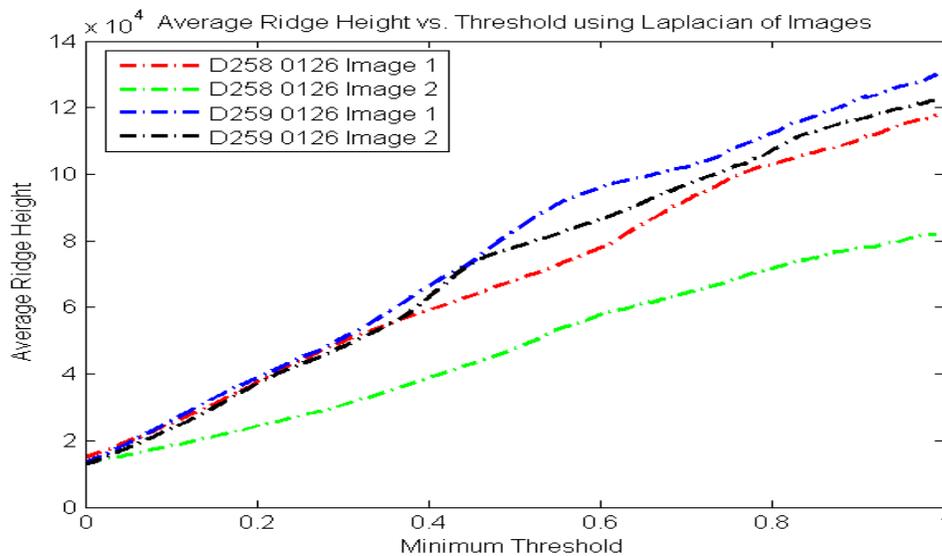


Figure 10: Average ridge height based on the Laplacian of the images.

D. Analysis

As expected, depending on the definition of the average ridge height, the material with the greater ridge height varies. Based on the derivative, there is no conclusive difference between the materials. Based on the gradient, the ridge height is greater for D258 than for D259. Again, this makes sense because D258 has more regions where there will be a nonzero first-order rate of change. Based on the Laplacian, D259 has higher average ridge height than D258. As the threshold increases, the ridges for D259 stay the same until the threshold is very high because D259 has high plateaus with deep holes.

V. Holes/Deep Regions

A. Expectations

The most obvious feature of D259 is its distinctive holes. D258 does not have holes per se, but it certainly has deep regions in the form of long canyons. Several properties would be useful to examine, including total area of deep regions, average area of each deep region, and average major and minor axis lengths.

Upon observation of the original images, for a given threshold D259 should have shorter average major and minor axis lengths. Also, D258 should have a greater area (and thus average area per region) for a given threshold. As the threshold varies from zero to one, the total area will jump at a lower value for D258 than for D259. This is because D258 has lower contrast between low and high regions. D259 has distinctive holes which are much deeper than the surrounding area. Thus, as the threshold increases, only those holes will be considered deep until a much higher threshold is reached.

B. Processing

Each image is thresholded such that all pixels less than or equal to a given threshold are set to one. This mask is then passed through a wavelet low pass filter to smooth regions (this is especially important for D258) as well as remove very small regions that are not useful. Using this filtered mask, each pixel in a contiguous region is given a number where the highest number corresponds to the number of separate regions in the image. Two examples are shown in Figures 11 and 12 where the threshold is 0.4 and 0.6 respectively. In these images, very small regions with area less than 25 pixels have been removed. In Figure 12 the very small areas with a value of zero are not holes, but rather the very small areas that are still above the threshold. The bar on the side shows a scale of 0 to 1. However, there are no values in between 0 and 1 in the image because the regions have been labeled using integers. Therefore, the bar is showing that there is only one region that is considered to be deep.

The Matlab command *regionprops* is used with the labeled mask as input. The output of *regionprops* is a cell array that includes the total area, major axis length, and minor axis length for each region in the image. Next, for each image, the total area, the median area for each region, as well as the median major and minor axis lengths is calculated. The median is used instead of the average so that the answer does not depend on the number of regions. The number of regions does not increase and decrease smoothly, but appears noisy. Thus, while the totals increase smoothly, the averages may also appear noisy. Only regions with an area greater than 25 pixels are counted. This is to prevent very small areas from affecting the results. Finally, each property (total area, median area, major axis length and minor axis length) is plotted against the threshold.

C. Results

First, Figures 11 and 12 show the labeling process and also demonstrate that D258 has larger deep regions than D259 at a given threshold.

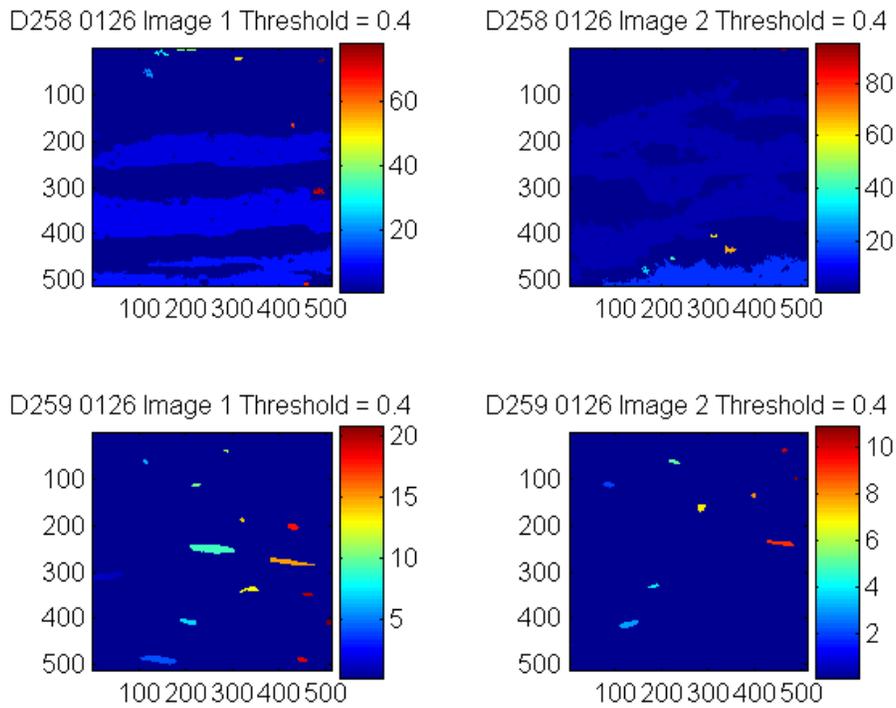


Figure 11: Deep region labeling with threshold equal to 0.40.

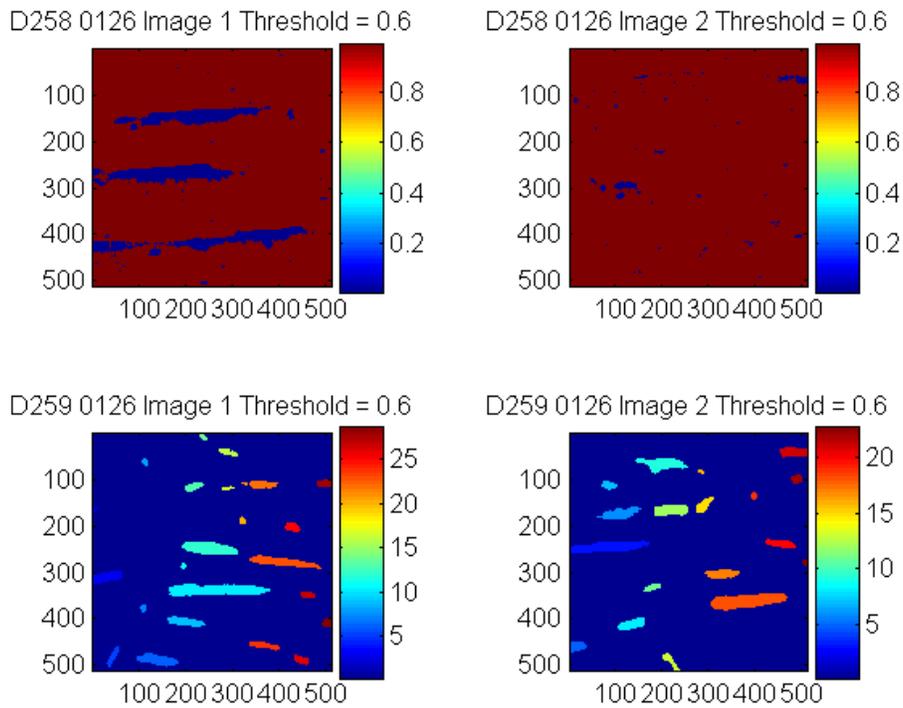


Figure 12: Deep region labeling with threshold equal to 0.60.

The total area and average area versus threshold are shown in Figures 13 and 14 respectively. The total area plateaus at 262144 pixels, which makes sense because when the threshold equals 1, all pixels are less than or equal to 1 and should be counted.

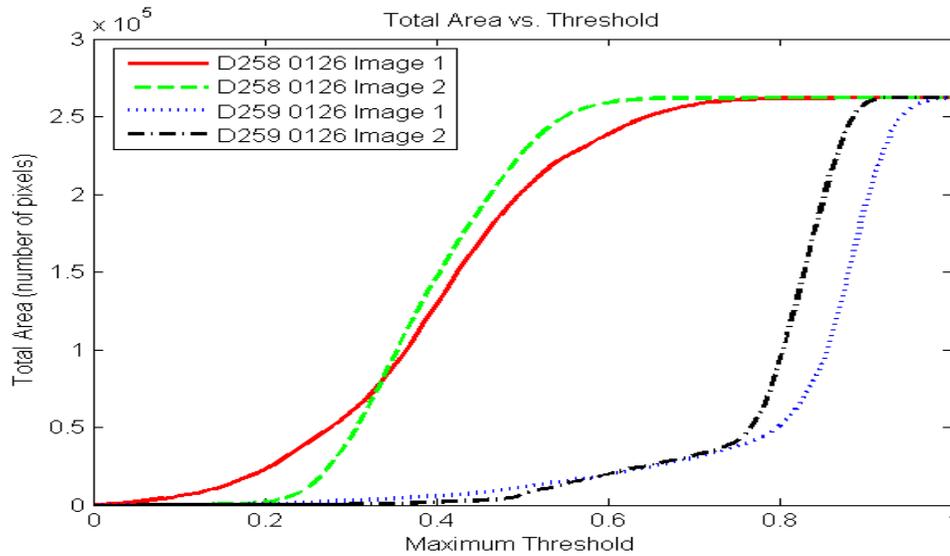


Figure 13: Total area below threshold.

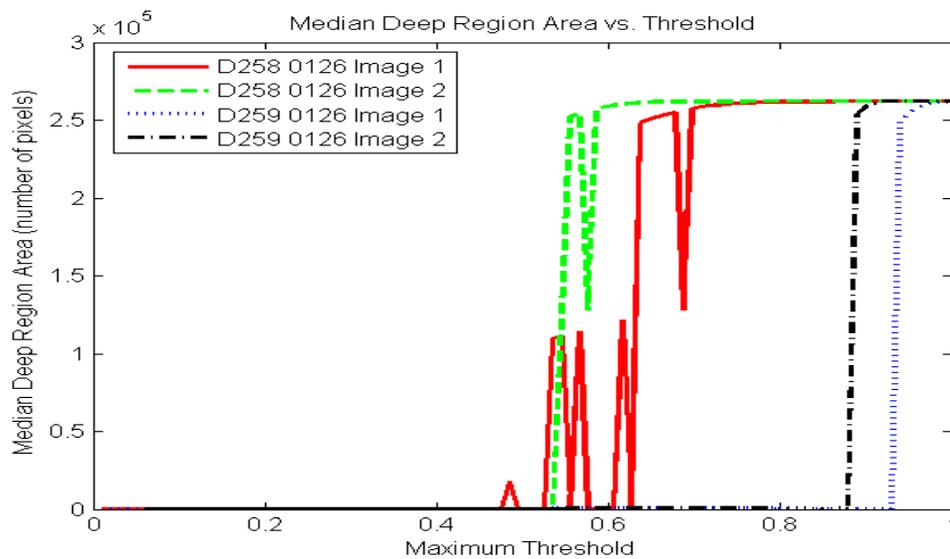


Figure 14: Median area of individual holes based on given threshold.

Figures 15 and 16 show the average major and minor axis lengths versus threshold respectively. Both reach a plateau of 256 which is to be expected because at threshold equal to 1, all pixels should be counted.

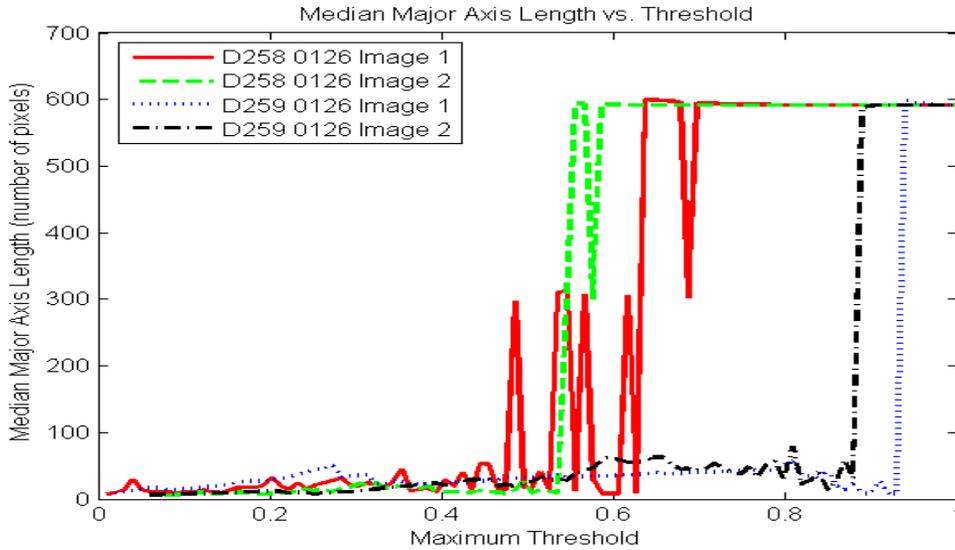


Figure 15: Median major axis length for individual holes.

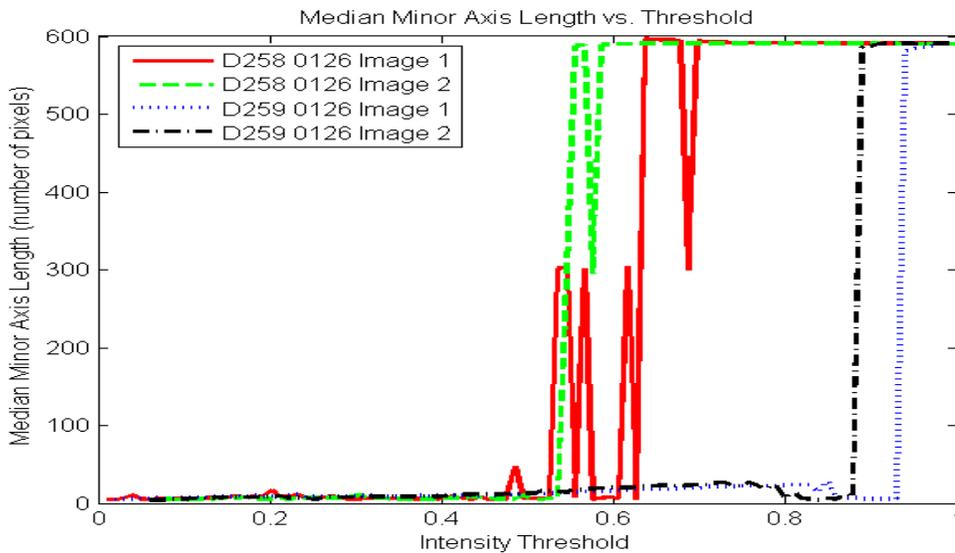


Figure 16: Median minor axis length for individual deep regions.

D. Analysis

The hole/deep region test shows the most definitive differences between D258 and D259. D258 has relatively shallower deep areas. Therefore, as the threshold is increased, it reaches a point where almost all pixels are below the threshold much earlier than in D259. It makes sense for this pattern to be repeated for each of the properties examined because they are all related to the size of the regions.

VI. Conclusions

It is obvious from visual observation that the two materials D258 and D259 are different based on the AFM images. Several visually distinct features include smoothness, average ridge height and deep region area and dimensions. Upon developing methods to quantify these differences, the most conclusive distinguishing features are the deep region area and dimensions. In this

case, D258 had deep areas, but nothing that could be described as a hole. This means that instead of polygon shaped regions, most of the regions are long with jagged edges, which may not allow an accurate estimate of the major and minor axis lengths. While this test was useful in comparing D258 and D259, it would be interesting to see the results after comparing two materials that had differently shaped holes.

Smoothness showed some of the expected relationships. The difference is not distinctive enough to say definitively that one image is smoother than the other. However, depending on the materials used in future comparisons; smoothness may be very useful in quantifying the differences between AFM images. Both of these materials are probably of medium smoothness, so comparing to a very rough image or a very smooth image would give some interesting results.

Finally, for ridge height, the definition of a ridge has a large impact on the ability to differentiate between D258 and D259. D258 has a rapidly changing surface, and therefore has a greater first order difference than D259. D259 has a greater second order difference because of its high contrast between the bottom of the holds and the high plateaus. This test would also be useful in comparing two materials that have features similar to those in D258 or in comparing two materials that have features similar to those in D259.

References

- [1]Carmona, René and Zhong, Sifen. “Adaptive smoothing respecting feature directions.” *IEEE Transactions on Image Processing*. Vol.7, No.3, March 1998.
- [2]Kawata, Y., Ohmatsu, H., Kakinuma, R., Eguchi, K., Kaneko, M., Moriyama, N. “Quantitative surface characterization of pulmonary nodules based on thin-section CT images.” *IEEE Transactions on Nuclear Science*. Vol. 45, No. 4, August 1998.
- [3]Kocak, Donna M. da Vitoria Lobo, Niels. Widder, Edith A. “Computer vision techniques for quantifying, tracking, and identifying bioluminescent plankton.” *IEEE Journal of Oceanic Engineering*. Vol. 24, No.1, January 1999.
- [4]Yuan, Xin and Shi, Pengsheng. “Microcalcification detection based on localized texture comparison.” *2004 IEEE International Conference on Image Processing*. 2004.
- [5]Yue, Yong. Croitoru, Mihai M., Bidani. Akhil, Swichenberger. Clark, John. “Nonlinear multiscale wavelet diffusion for speckle suppression and edge enhancement in ultrasound images.” *IEEE Transactions on Medical Imaging*. Vol 25, No. 3, March 2006.