Digital Image Analysis of the Shroud of Turin: An Ongoing Investigation

Raymond J. Schneider
Bridgewater College, Department of Computer Science, Bridgewater, VA 22812, USA

Abstract

Since the Shroud of Turin is only available for viewing on long intervals at the pleasure of its custodians and is primarily an object of religious veneration, it is difficult to get new materials for study. However, since 1898 and the original photographic work of Secundo Pia a number of high quality images of the shroud have been collected by Enrie, Miller, Schwortz, Durante and others. More recently high resolution digital scanning of such images has made the application of digital algorithms feasible even with relatively modest equipment.

The present work has as its objective the development of a research program to apply comparative analysis to the several images of the shroud that are available. Among the objectives is a program to determine how much information can be extracted not only from the individual images but from comparing the same regions of different images. The work reported in this paper is the initial effort in this larger project.

1. Introduction

The history of scientific shroud research is intimately entwined with photography. While the Shroud of Turin was simply a relic enshrined in a cathedral in Turin Italy and seen only on rare occasions, it did not stir much controversy. The faithful probably believed it to be true while the skeptics didn't think about it at all. This changed in 1898 when the shroud was first photographed by Secondo Pia. [1]

What Pia discovered when he developed his photographic plates is that the shroud acted as a negative so that his plates contained positive images. Photography was then still a young science. Since Pia's historic photographs the shroud has been photographed with increasingly sophisticated technology numerous times, but most notably by Enrie in 1931, by Miller and Schwortz in 1978, by Durante in 2000 and in 2002, as well as others.

The most recent image of the shroud is that taken by HAL 9000, a company specializing in art photography that obtained permission to take an extraordinarily high definition composite image of the shroud on 22 January 2008. This image is composited from over 1600 credit card sized images for a total image of 12.8 billion pixels according to a story reported by discovery.com [2].

The present work uses images from the work of Barrie Schwortz, the documenting photographer of the 1978 STURP project and scans of images from Durante 2000 provided by Giulio Fanti.[12]

2. Levels of Analysis

Four levels of analysis are distinguished as 1) Data, 2) Recognition, 3) Aggregation, and 4) Meaning. The levels are based on the degree of interpretation required to express them.

2.1 Data The level which is entirely objective is the first, the Data level which simply consists of the quantitative measurements that defined the images. These are simply the two dimensional manifolds of 24-bit (8 bit RGB color planes) values.

2.2 Recognition Each point in the images can be categorized, generally by color and texture and immediately adjacent pixel values as one of a finite set of categories such as image, blood, cloth, burn/scorch, and others distinguishable by very small local regional features.

2.3 Aggregation This level recognizes integrative collections of small regions, generally contiguous or nearly contiguous elements such that continuity and smooth transition of metrics can be distinguished. These regions can be aggregated into recognizable patterns of an extended sort so as, for example, to be able to distinguish such things as facial features, hands, blood flows, and other spatially extended and connected features.

2.4 Meaning At the level of meaning one puts the aggregation features in context and goes beyond the aggregation data inferentially to generate interpretations of the meaning of the features. Some speculative examples would involve inferences about the flow of the blood, the pattern of wounds, the design of the scourge, the chemistry of the image, the shape of the enclosing cloth as explained by features present in preceding levels of analysis and other external factors. Meaning goes outside the domain of the three preceding levels to a framing context which itself has data, recognition and aggregation levels.

3. Initial Motivations

Three elements motivated the current work. The first was the presence in the shroud image of banding features of several kinds: 1) reflective, 2) transmitted, and 3) intrusive. The second element was JavaScript application written by Mario Latendresse [3] which emphasized the need for quantitative work on the shroud images. The third element was the availability of high resolution digital scans of color photographs of the shroud from the 1978 STURP expedition made available by the documenting photographer, Barrie Schwortz. Later other high resolution photographs used in this effort were provided by Giulio Fanti.

3.1 Banding Features The first kind is one that quite literally reflects the herring bone weave of the cloth giving a striped appearance to the cloth at the points where the chevrons change direction. This first kind of banding might be termed reflective banding caused by the interaction

of light with the weave pattern of the cloth. A second kind of reflective banding shows changes in the relative acceptance on the cloth of the image mechanism and is evidenced by a difference in image density. This effect is most easily seen along the sides of the face where there are areas that appear to have no image but on image enhancement it can be shown that image features extend into the apparently
image-less regions. One possibility advanced, but not definitively confirmed, is the idea that the image mechanism depends on some aspect of yarn lot treatment and so varies by the treatment the hanks of yarn received prior to being woven into the shroud. [4] A third kind of banding arises from yarn density and is particularly apparent in transmitted light photographs of the shroud.

3.2 Quantitative Analysis The application that Mario Latendresse developed is relatively simple. [3] It allows features on the shroud to be located using x,y pixel coordinates and converts pairs of these to Cartesian distances. This is a small step in the direction of
doing more quantitative analysis of shroud images individually and collectively (see three different images on the *Measure* sites). If high quality images of the shroud were available, clearly a wide variety of studies based on image processing could be explored.

3.3 Image Availability There have been a large number of high quality shroud images produced using photography. The initial image by Secundo Pia was not of particularly high quality both due to the relatively primitive state of photography at the time and limitations of lighting and time to take exposures as well as other difficulties. The images of Enrie in 1931 were very high quality but very high quality digital scans of
these images are not available at least to this author and the negatives have degraded over time. Moreover they are black and white images taken with very non-linear orthochromatic film. More recent images have been taken with color film and scanned at high resolution. This work takes advantage of the availability of high resolution color photographic scans provided by Barrie Schwortz and Giulio Fanti. It also takes
advantage of the availability of many high-quality image processing tools that have been developed for use on relatively inexpensive personal computers. [5] It is doubtful if this work, as limited as it is, could have been undertaken by an individual as recently as 20 years ago simply because of the state of the computing art.

4. Initial Result

The initial work done on the shroud images was done on images taken by Barrie Schwortz in 1978. These were cropped since the original scans were nominally 300 MB and this was too large to handle and still do significant processing even with 2 GB of main memory. The image was cropped to a ventral (front) head image to explore the banding alongside the face. Since the image is known to be a straw yellow color it is not too surprising that the pixels, with only a few isolated exceptions, all have the red plane of the three RGB planes as the one with the highest intensity.

One of the image characteristics commonly attributed to the shroud is that the chromophore is monochromatic and the apparent differences in color are due to changes in areal density like a half-tone image only composed of color patches of all the same color with more where the image is darker and fewer where the image is fainter. In view of this it seemed reasonable to suppose that a band intensity normalization of the red
band maintaining the ratios of the green and blue bands might tend to remove or reduce banding if it was due to differences in image take up. It would also reduce contrast eliminating it entirely in the red band which would go to a single intensity.

A simple algorithm was written in Python using the Python Imaging Library (PIL) to accomplish this transformation. The result is shown below in Figure 1.

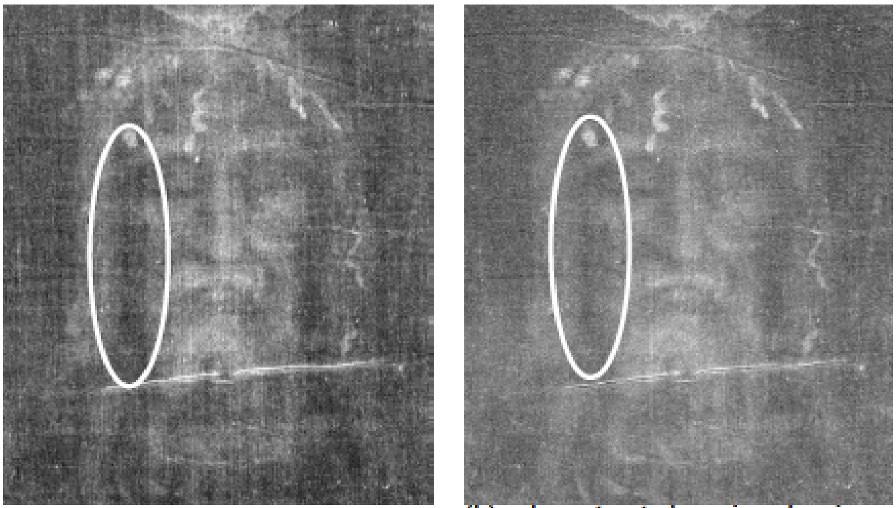

(a) negative showing band at side of face

(b) color saturated version showing smooth gradient in band area

Figure 1. Suppressed Face Bands with color normalization

This image (Figure 1) was produced by a simple normalization algorithm. The image was then converted to grayscale and inverted. Figure 1.a is the original simply converted to grayscale and inverted. Figure 1.b was saturated, converted to grayscale and inverted. The result of applying this simple algorithm was dramatic. The whole image was smoothed, the bands both near the face and elsewhere in the image all but disappeared, and the face filled out and appeared more contoured. This preliminary result was very encouraging.

5. Long and Short Range Objectives
The initial result was sufficiently promising to suggest that significant work could be done on the shroud using high resolution scans and it appeared likely that a broad image processing research program could yield meaningful results. It also seemed likely that students even at the undergraduate level could contribute to such research. Another paper given at this conference discusses possible content for such a research program. To begin however, it appeared that one would have to conceptualize a comprehensive image study research program. Initial thoughts used the four level model discussed briefly in section two and sketched in the following elements (given in no particular order):

5.1 Some Thoughts on a Comprehensive Image Study Research Program
Since there are many images of the shroud that have been taken one of the first thoughts was that any meaningful program should consider as large a range of images as could be obtained.

- Address a Wide Range of Available Images

What ought to be done with the images? Images can be thought of as a set of features and images of the same thing ought to be composed of the same set of features. Differences in the images are likely to reveal more or less about the features as the nature of the images reveal them.

- Fully Characterize the Features in the Images

Among the earliest objectives is one to characterize the kinds of features. Some of the relatively straightforward categories are those commonly discussed when the shroud image is assessed.

- Categorize the Features, ex. Cloth (Banding), Image, Blood, Scorch, others

Among the tasks that may be especially important in considering multiple images is the question of how are differences and commonalities among the images to be distinguished? The term generally used to align common features is *normalization*. Kinds of normalization include spatial, temporal, and spectral. An example might be:

- Color Normalization to Reduce Banding and Enhance Image Features

A narrower study would be one that focuses on a particular kind of feature with the objective of understanding it more completely in the context of its creation:

- Blood Image Enhancement, especially of Scourge Markings

As projects are undertaken it is essential to be alert for new ideas and new information that leads to discoveries and fruitful directions of research previously not envisioned.

- Definition of New Projects as Intermediate Research Reveals Fruitful Directions

Narrowly held research fields are always only one generation away from disappearing if new researchers do not take up the challenge. It should be a goal of sindonology to interest younger researchers in the field to constantly revitalize it.

- Develop a Research Program that Brings New Researchers into the Field

5.2 The Goals of the Present Research

The initial goals which the author adopted for this research were to restrict the research to the first two levels discussed in section 2, Data and Recognition. A general problem in image processing is that of *image segmentation*. Image segmentation involves partitioning the image into regions, distinguishing distinct elements in the image from one another. There are many techniques in the literature but no general solution. Different domains often call for domain-specific techniques. [6] The problem addressed by the author is to use color, luminance, and texture to develop classification metrics to divide the shroud image initially into major categories such as blood, cloth, image, scorch, and such other categories as can be distinguished. The work was divided into two steps: 1) development of metrics, and 2) false color substitution to evaluate the success of the metrics.

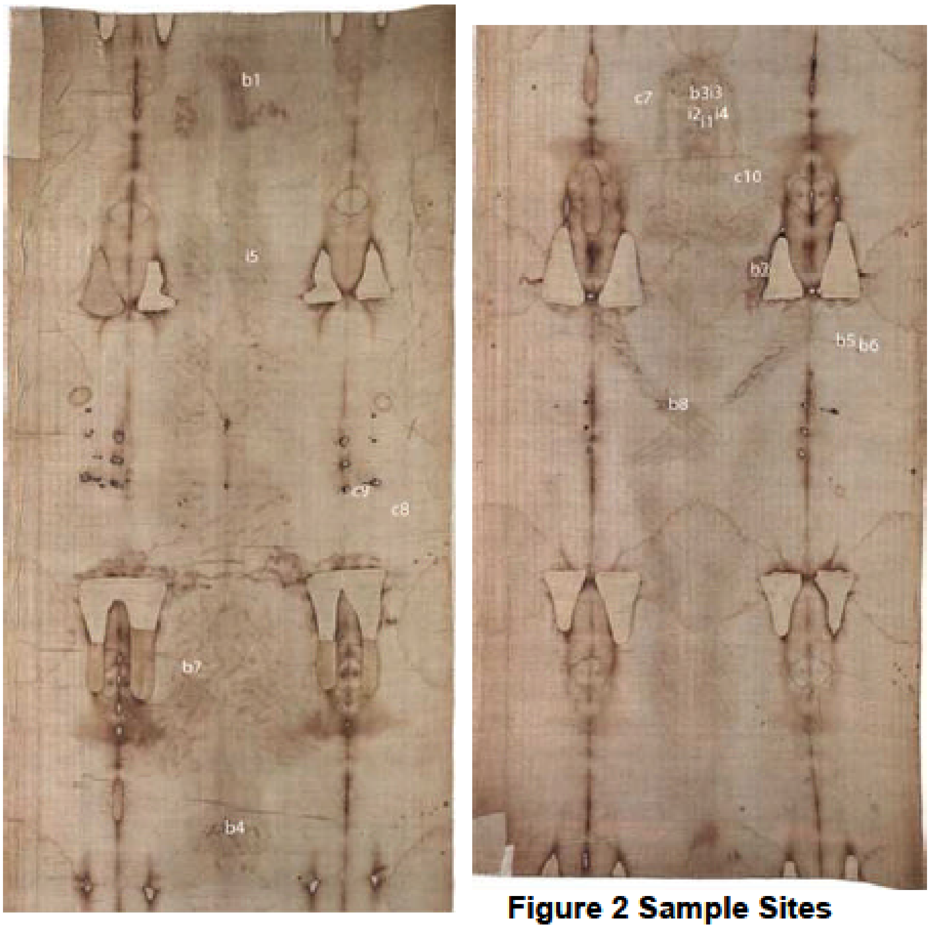

Figure 2 Sample Sites

6. Initial Steps: Sampling, Tools, and Coordinate Systems

The initial step in the study was to select representative samples of various categories for characterization. There are many possible categories that could be selected including: 1) blood, 2) image, 3) cloth, 4) scorch, 5) waterstains, 6) burns, 7) droppings, 8) dirt, 9) patches, and others. Within each category it would be possible to select various degrees such as the density or relative degree of admixture of other factors as judged by the observer. To begin it was decided to study only the first few categories. Further study could add additional categories if this preliminary study proved fruitful.

6.1 *Initial Sampling*

Samples were extracted at arbitrary locations from the Durante 2000 image. [7] The

Figure 3 Nineteen samples of cloth, image, blood and two scorch images

selection criteria were informal, primarily that the sample be representative of a single phenomenology and uniform in herringbone direction. The coordinates given on all but one of the samples are the horizontal and vertical pixel location coordinates at the approximate center of the sample. The prefixes stand for (c)loth, (i)mage, (b)lood, and (s)corch.

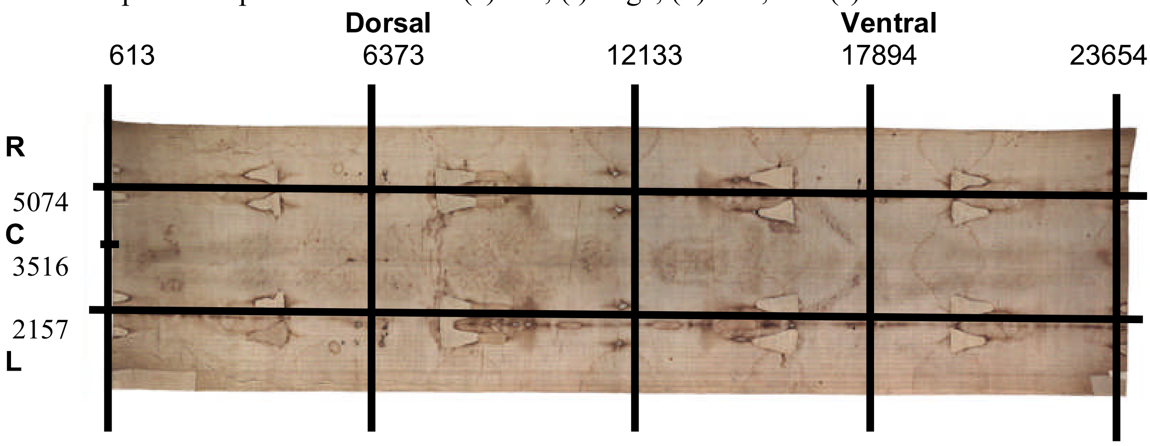

Figure 4 Approximate coordinate system roughly gridded on the longitudinal burn lines. Durante 2000 is 7632 pixels x 24197 pixels, including an off- image white border.

Table 1 Samples cropped from Durante 2000 Image for Study

Durante 2000 Shroud Samples Used in Study						
	Small Samples					
	X	Y	W	H		Desc
1	3545	1609	34	30	b1	blood foot
2	3989	15828	28	29	b2	blood chest wound
3	3231	13352	27	25	b3	blood forehead base of E
4	3333	11408	23	23	b4	blood back of head
5	5148	16635	33	36	b5	blood elbow drip
6	5392	16734	17	22	b6	blood elbow drip smaller further out
7	2768	9341	25	32	b7	blood scourge back
8	2931	17472	31	37	b8	blood wrist wound
1	5100	10065	40	39	s1	scorch
2	TBD		43	41	s2	scorch dark
1	2453	13512	38	39	c7	clear cloth
2	5474	7284	25	30	c8	clean cloth
3	4928	7018	29	29	c9	clean cloth
4	3646	14588	46	64	c10	clear cloth right head strip below face and neck
1	3304	13858	43	42	i1	image tip of nose, possibly dirt too
2	3293	13666	45	60	i2	image left cheek looking at positive image
3	3424	13312	50	56	i3	image above right eye
4	3541	13625	38	28	i4	image right cheek looking at positive image
5	3693	3937	34	37	i5	image right calf
	Larger Samples					
1	3205	13174	73	172	bf1	blood stain on forehead
2	2736	17367	297	189	bw1	blood wrist wound
3	2949	13299	67	92	bf2	blood forehead left large spot
4	3876	15461	393	817	bc1	blood chest lance wound
5	2118	12929	579	1030	c1	cloth left of face
6	956	13374	728	533	c2	cloth on left of face beyond vertical burn line
7	3892	13022	640	821	c3	cloth on right of face
8	5032	12986	812	793	c4	cloth on right beyond burn line
9	910	16583	858	813	c5	cloth left of forearms above wriste and left of burn line
10	4887	16784	1203	676	c6	cloth right of forearms right of burn line
11	3157	22854	461	387	bvf1	blood ventral feet (large stain)
12	3670	23177	229	242	bvf2	blood ventral feet (smaller stain)
13	2928	13459	90	581	fbl1	faceband left 1
14	3596	13423	104	645	fbr1	faceband right 1
15	2624	1657	456	731	bdf1	blood dorsal feet 1
16	3154	1269	704	1067	bdf2	blood dorsal feet 2
17	3864	1860	281	272	bdf3	blood dorsal feet 3

Table 1 lists samples cropped from the Durante 2000 image. The samples are listed by arbitrary sample number, approximate center coordinates (X,Y), the height, H, and width, W of the samples in pixels, a short designation using in naming the image file for processing, and a brief description.

6.2 *Tools*

Software tools used in the course of this work included:
1. Adobe Photoshop Elements http://www.adobe.com/
2. Python PIL http://www.pythonware.com/products/pil/
3. Matlab Image Processing Toolbox http://www.mathworks.com/products/image/
4. CVIPtools http://www.ee.siue.edu/CVIPtools/
5. ImageJ http://rsbweb.nih.gov/ij/

The size of the Durante 2000 image was so large that it was not realistic to use the entire image in the computers which the author had available for the work. In view of the exploratory nature of this initial work it was decided to limit the effort to a series of samples of various sizes. Small samples were used to isolate features so that statistics could be collected on single classification characteristics. Larger samples were used to test processing concepts. The largest image used repeatedly was termed *facecrop* and included the face and upper chest of the shroud image.

6.3 *Coordinate Systems*

Several coordinate systems were used in processing the data. The initial images used are all in 24 bit, RGB coordinates. The RGB coordinate system is not particularly suited to image processing. One of the processing characteristics desired was to separate color and luminance information.

This was done by retaining the basic RGB color space but doing a coordinate transformation to spherical coordinates. The color information is then the direction in the RGB space normalized to a unit vector and the luminance is the magnitude of the resulting vector. This is not one of the more customary definitions of these entities, so the reader should be cautioned that there are many color spaces in the literature and that this particular one is not intended to have human color perception significance, but only a convenient way to separate color and intensity for analysis.

RGB Color Space and transformed versions of it.
The original color, (R,G,B), can be represented as a color vector:

$\vec{C} = R\hat{r} + G\hat{g} + B\hat{b}$ where the values (R, G, B) are the one byte encodings [0 .. 255] of the color vect components and the elements $\hat{r}, \hat{g}, \hat{b}$ are respectively the directions in an abstract orthogonal color three-space of the three colors, red, green, and blue.

We can create a color unit vector by defining $L = \sqrt{R^2 + G^2 + B^2}$ the magnitude of the color and dividing through the color vector by L giving us a color unit vector:

$\hat{C} = \dfrac{R\hat{r} + G\hat{g} + B\hat{b}}{\sqrt{R^2 + G^2 + B^2}}$ and this color unit vector is a two dimensional entity which can be remapped to a two dimensional (φ,θ) space where $\varphi = \arctan(G/R)$ and $\theta = \arctan(\dfrac{\sqrt{R^2 + G^2}}{B})$ here the color is encoded as a point in (φ,θ) space where both φ and θ have a range of $0 .. \pi/2$.

These simple transformations of the RGB color space allow representation of color as a direction in the space and intensity as the magnitude of the color vector. For more on this particular color transformation and a related color segmentation algorithm (SCT/Center) which uses it see [8].

6.4 *Narrow Color Space*
Using the SCT coordinate system it is relatively easy to map all the pixels in a sample to the (φ,θ) space and see the range of color in the image. Figure 5.1 below gives the color range by plotting all the pixels in the Face Crop as black, cloth sample c1 as blue, an image sample taken across cheeks and nose as green, and blood sample b1 as red. Two immediate observations are appropriate: 1) A very small part of the overall color space is actually occupied by the full range of shroud pixels, and 2) because of this small range there is a good deal of color overlap between ostensibly different areas of the shroud.

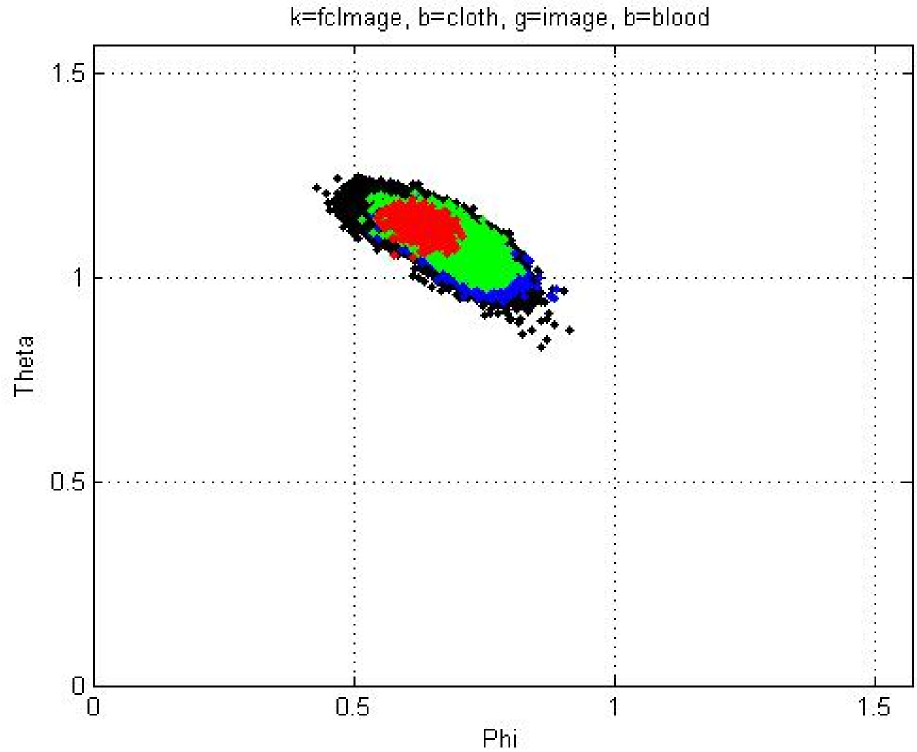

Figure 5.1 Shroud pixels plotted in Phi,Theta unit vector space where black pixels are all the pixels in the FC (Face Crop) image, blue pixels are cloth pixels from the c1 sample, green are image pixels from the cheeks and nose in the face area, and red are pixels from blood sample b1.

6.5 *Adding Luminance*

In order to enhance the separation of the various elements of the shroud the luminance information is important. We can cut the three space formed by (φ, θ, L) and show how the addition of luminance information enhances the separation of the elements. However, even with the addition of luminance information the color space remains very small and confounded.

Figure 5.2 uses the same subimages (face crop, cloth, image, blood) to illustrate that the addition of luminance information does produce substantial discrimination. The cloth for example has the highest luminance values while blood stains are among the darkest pixels. Pixels related to the image area across the nose and face fall in an intermediate range. Even with the addition of luminance information is it obvious that the shroud pixels do not divide nicely into disjoint sets and that fact substantially complicates the general problem of image segmentation applied to the shroud.

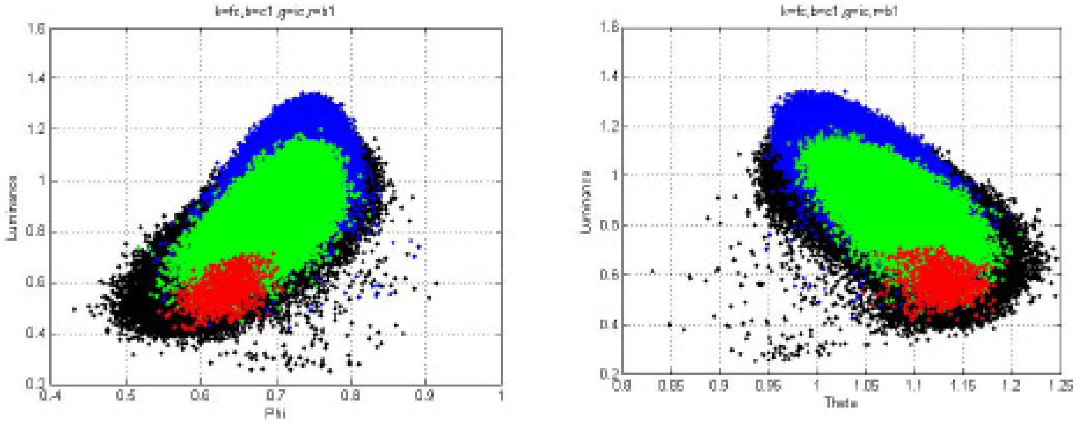

Figure 5.2 Two slices through the (φ, θ, *L*) space showing some enhanced separation provided by the addition of Luminance information. Nevertheless substantial overlap among the various classifications of pixels remains.

7. Methodology

Traditional methods of image segmentation are challenged by the narrow and confounded color range provided by the shroud. The initial methodology restricts the analysis to the Data and Recognitions levels presented in section 2 used to explore pixel categorization based as much as possible on single pixel metrics. This goal was pursued in eleven steps:
1) false color substitution, 2) image color indexing using minimum variance, 3) color separation using the SCT described in section 6, 4) contrast stretching of the separated color unit vectors, 5) combination of unit vector separation, indexing and false color using black, 6) black and white false color substitution in a 24 color indexed image, 7) combining color directionality and luminance intervals to create categories, 8) using the sample median to create stripe and interstitial subcategories, 9) using wrist stripe and interstitial measures and transferring them to the chest wound, 10) statistics of stripe and interstitial subcategories, 11) application of subcategories and using the nearest subcategory to false color the face crop image.

7.1 *False Color Substitution*

The first method employed to explore pixel classification was simply to create a threshold category to make the color substitution decisions going over the image. The earliest experiments used the PIL (Python Imaging Library) [9] Figure 6 shows the result of a slightly later experiment injecting red (1, 0, 0) into a face image that has been converted to an indexed image of 8 different colors. The pixels to convert were selected empirically by cycling black, (0, 0, 0), through the eight colors to see which would correlate most with the blood images. Then the strongest correlates were changed to red resulting in the image of Figure 6.

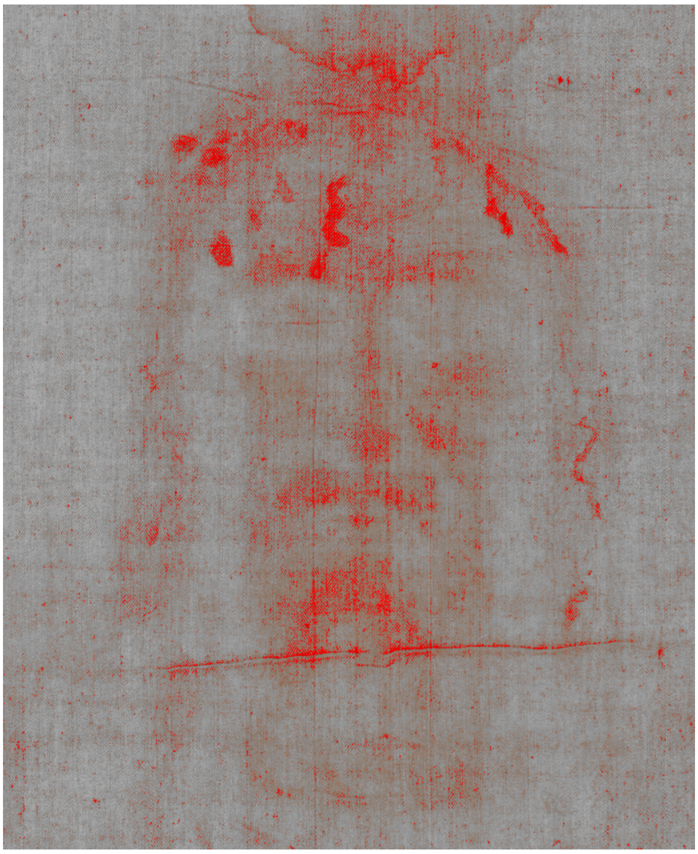

Figure 6 False color substitution into an image converted to an indexed image with eight colors.

7.2 *Minimum Variance Color Indexing*
An RGB image encodes colors with red, green, and blue pixels each encoded with a value in the interval 0..255. This coding results in a possible 16,777,216 different colors. An indexed color image codes the image with a number that identifies each unique color used in the image with a number that is associated via a map with an RGB code. Matlab contains several interesting functions to convert images to indexed images. Two functions that are particularly useful are: `cmunique` and `rgb2ind`.

`cmunique` — `[Y,newmap] = cmunique(RGB)` converts the true-color image `RGB` to the indexed image `Y` and its associated colormap `newmap`. The return value `newmap` is the smallest possible colormap for the image, containing one entry for each unique color in `RGB`.

The maximum possible unique colors in an image is the number of pixels. The Face Crop (FC) image used here is 1554x1867 = 2,901,318 pixels. Using `cmunique` we can determine that the image has 116,619 unique colors, only 4% of the possible numbers. `cmunique` is a simple way to determine how many colors are present in the image and also an easy way to generate a data set for color analysis.

`rgb2ind` — Converts a truecolor (RGB) image to an indexed image. Matlab provides two quantization methods to create the index colors of an RGB image, *uniform* and *minimum variance*. . The syntax for the minimum variance quantization is `[X,map] = rgb2ind(RGB,N);` where N is the number of colors in the indexed image. [10]

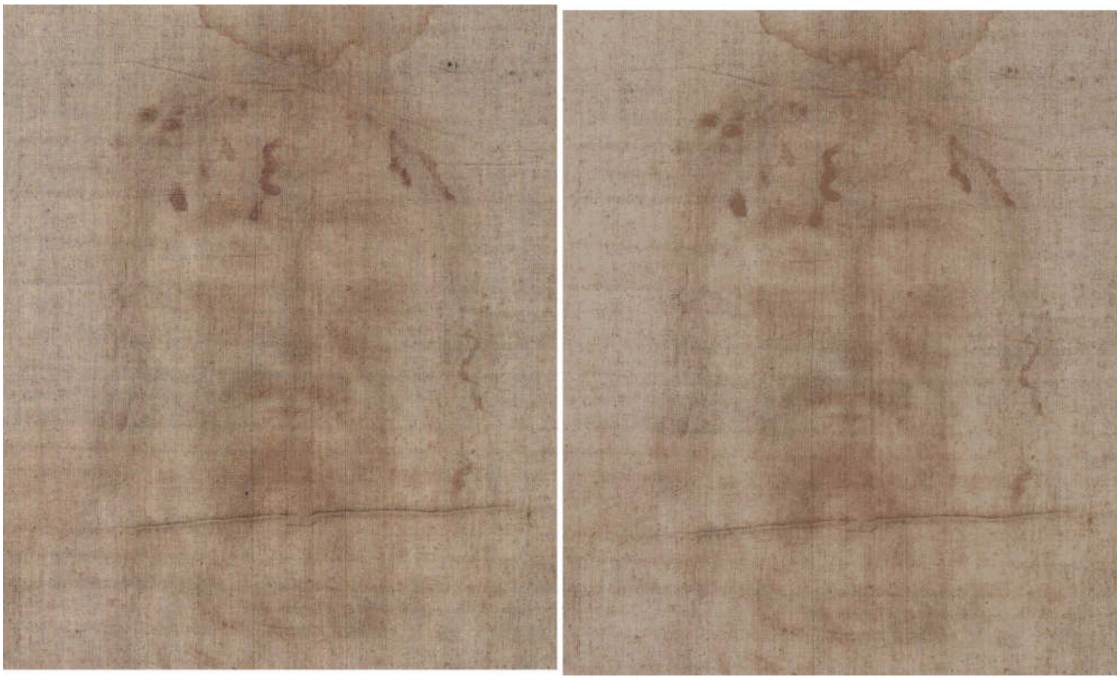

Figure 8 The image on the left is a Face Crop image with 116,629 colors while the image on the right is an indexed version of the same image with only 8 colors. They look identical on casual inspection although there is some loss in luminance information, the narrow range of the shroud image reduces the degradation due to indexed imaging.

Figure 8 shows how little subjective information is lost in converting a full range shroud image even rather aggressively to an indexed image with far fewer colors. This is particularly significant for exploring image segmentation since it facilitates false color injection.

7.3 *SCT Color Separation*

The previous subsection illustrated the effectiveness of minimum variance quantization in reducing the 116,629 color original image to an 8 color image with a minimal loss of image quality. Separating color and luminance information allows the investigation of how much information lies in the color encoding versus the luminance. Using the transformation described in section 6.3 the chest wound image in Figure 9 is transformed into a luminance and normalized color image. The original is then the product of the two images.

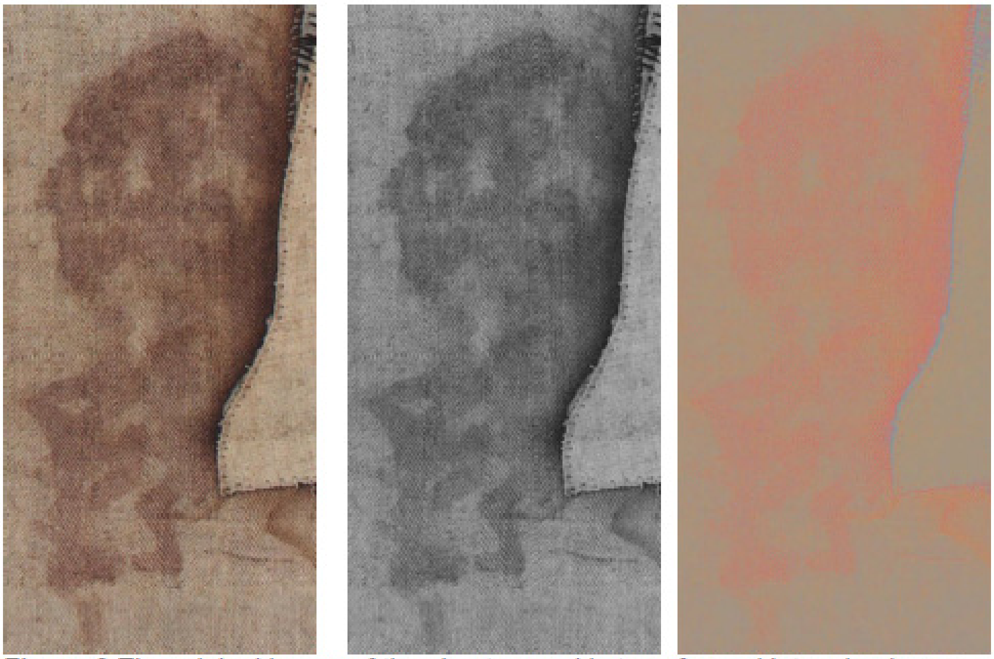

Figure 9 The original image of the chest wound is transformed into a luminance image and a normalized unit vector color image.

To recover the original image just multiply the luminance image by the normalized color image. Notice in particular how uniform the colors in the normalized color image become. In order to study the actual range of colors it would be helpful to do a contrast stretch.

7.4 *Contrast Stretching SCT Color Separations*

The color unit vector image can be enhance by using histogram equalization which stretches the contrast by distributing the intensity of the unit vector components more uniformly across the available intensity space. In Matlab this is rather easily accomplished by using the `histeq` function three times on the components and concatenating the result. If `RGB` is a unit vector color image the conversion to a stretched image is:

```
RGBs = cat(3,histeq(RGB(:,:,1)),
             histeq(RGB(:,:,2)),
             histeq(RGB(:,:,3)));
```

The resulting image is a color contrast enhanced version of the unit vector color image as shown in Figure 10. The default behavior of `histeq` is to convert the existing histogram of intensities to a uniform histogram of 64 bins. Other histograms can be used.

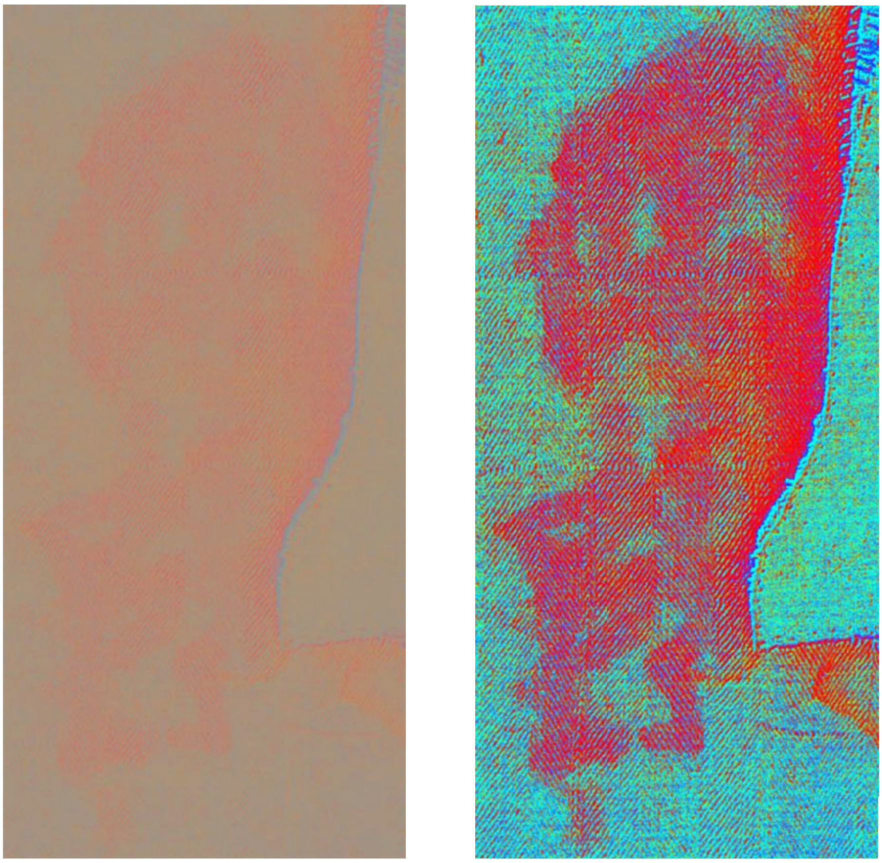

Figure 10 Color Contrast Enhanced version of the color unit vector image using the histogram equalization function in Matlab on the three RGB color components. Color Unit Vector image on left and triply histogram equalized image on right.

Applying the same technique of color unit vector conversion and stretching the contrast using component by component histogram equalization to the face crop image gives the image shown below in Figure 11. Perceptually at least this image seems to give a good segmentation of the main categories of interest. The blood on the forehead is rather magenta. What appears to be blood on the moustache and beard is slightly redder than the orange in the face, nose and cheek image areas. The scorch renders a fiery red and the relatively unmarked cloth renders in the cyan-blue range. This seems like a fairly complete segmentation based only on color after being histogram equalized.

This seems promising enough to suggest further work to improve the behavior of the color manipulation.

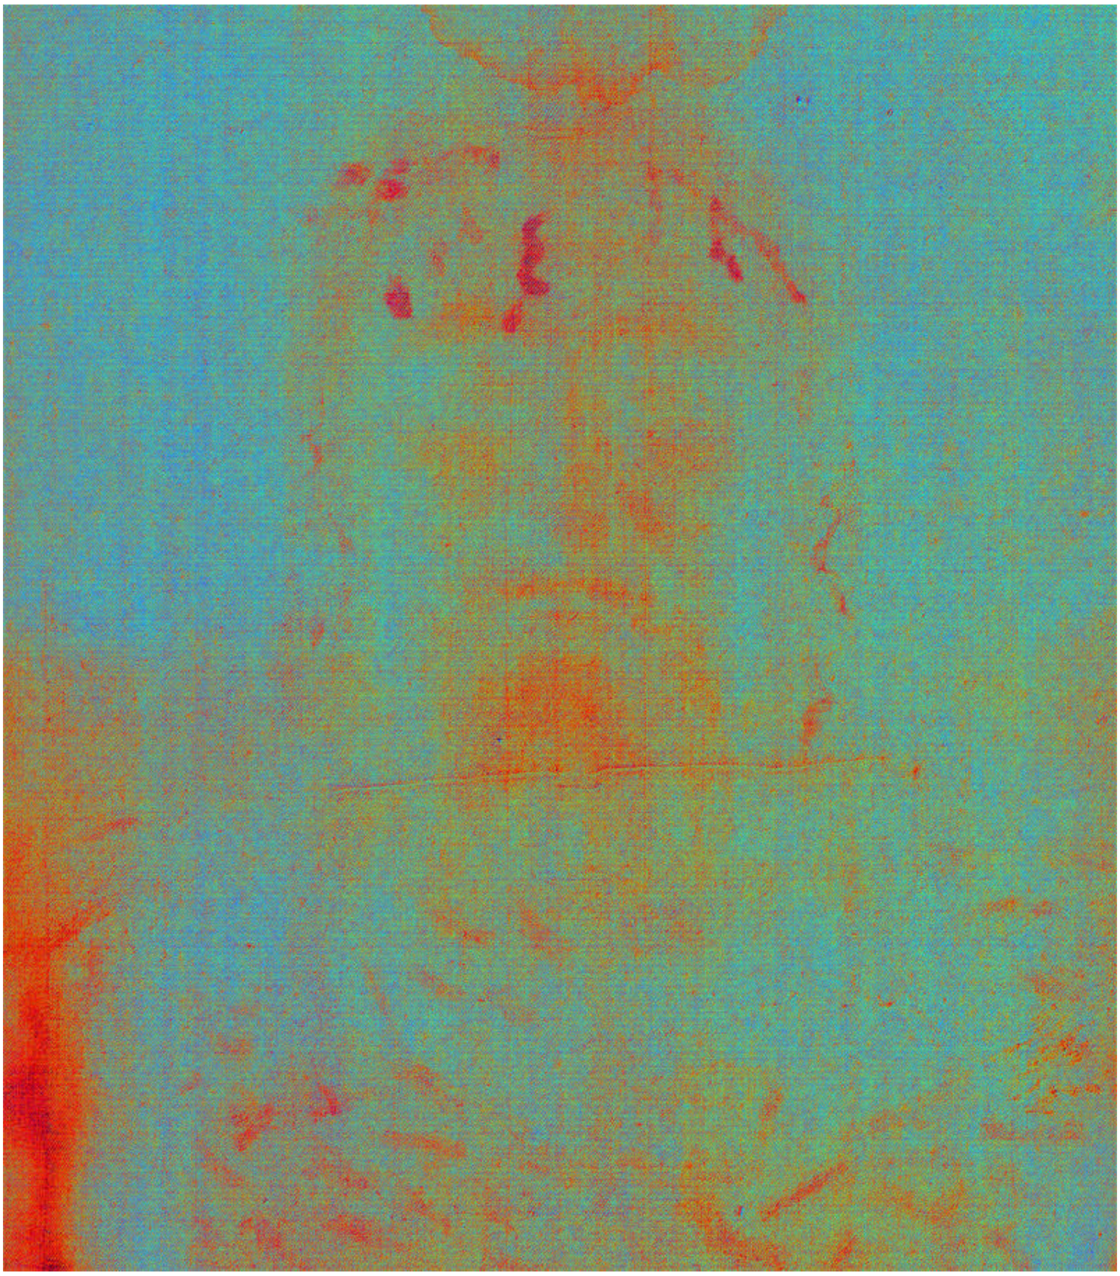

Figure 11 Color Unit Vector of Face Crop Image that as been Contrast Stretched using the method illustrated in Figure 10 for the chest wound.

7.5 *Combination Methods Applied to Isolated Color*

Using techniques in combination can produce interesting and stimulating effects. In Figure 12 below the method of converting to color unit vectors and contrast stretching is further extended by using color indexing to reduce the image to 16 colors and then using black color substitution on those colors that have Red, Green, or Blue as the largest components. The final image is one in which black is false color substituted for the colors that have either Green or Blue as their largest components. Notice that there is some enhancement of the vertical bands caused by the shifting of the herring bone weave of the shroud. Also that blacking out the Green and Blue leaves an image in Red which is very much like a positive image. The reader is encouraged to compare this last image with the red color normalized image in Figure 1b.

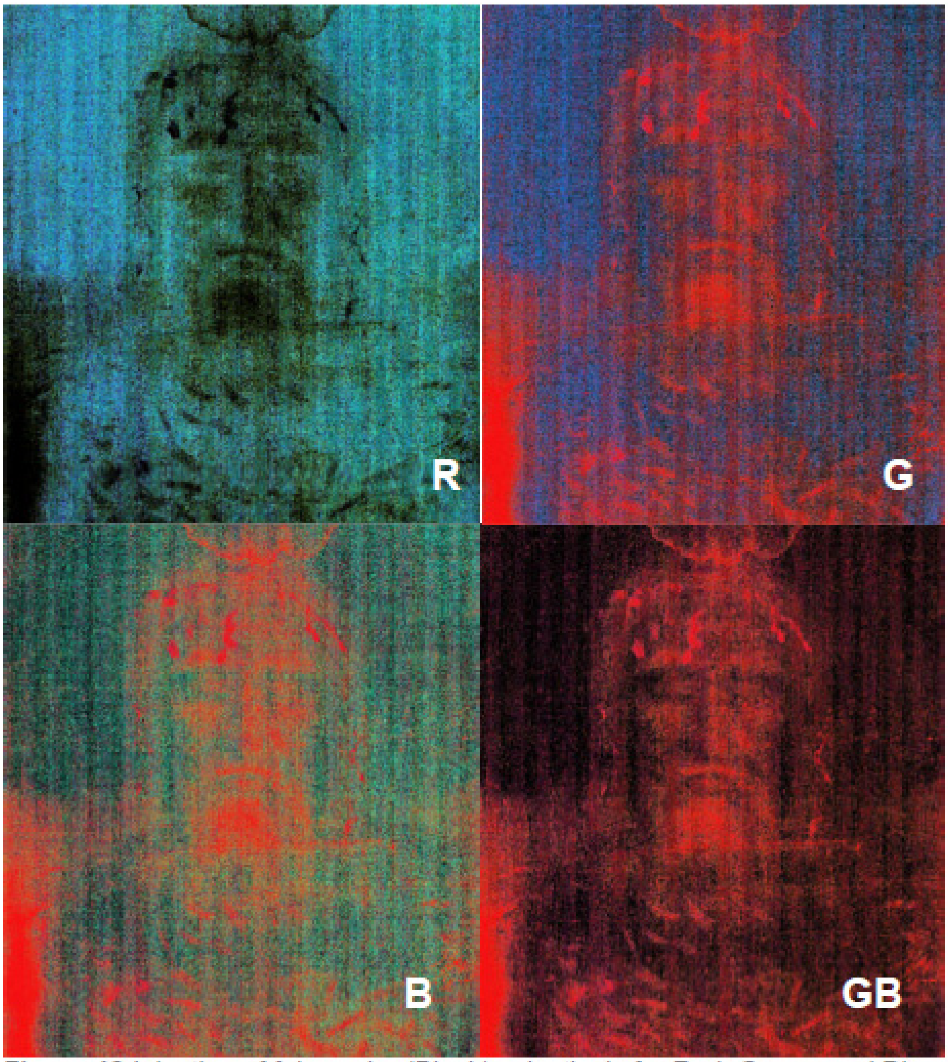

Figure 12 Injection of false color (Black) selectively for Red, Green, and Blue largest components in a 16 color indexed image conversion of the color unit vector image of Figure 11. Final image substituted black for both Green and Blue resulting in an image similar to Figure 1b.

7.6 *Combination Methods Applied to Luminance*

Applying the same two methods, indexed images and false color substitution, to the luminance variable instead of the largest color in the color normalized image gives some insight into the behavior of the image along the luminance dimension. The sub-images of Figure 13 were created by first converting the Face Crop image to a 24 color minimum variance indexed color image. Then false colors of black and white were substituted for pixels in stages. The top left image has black substituted for the darkest three colors. The next image has white substituted for the brightest five, and so on through the six images, white substituted for the brightest pixels and black for the darkest.

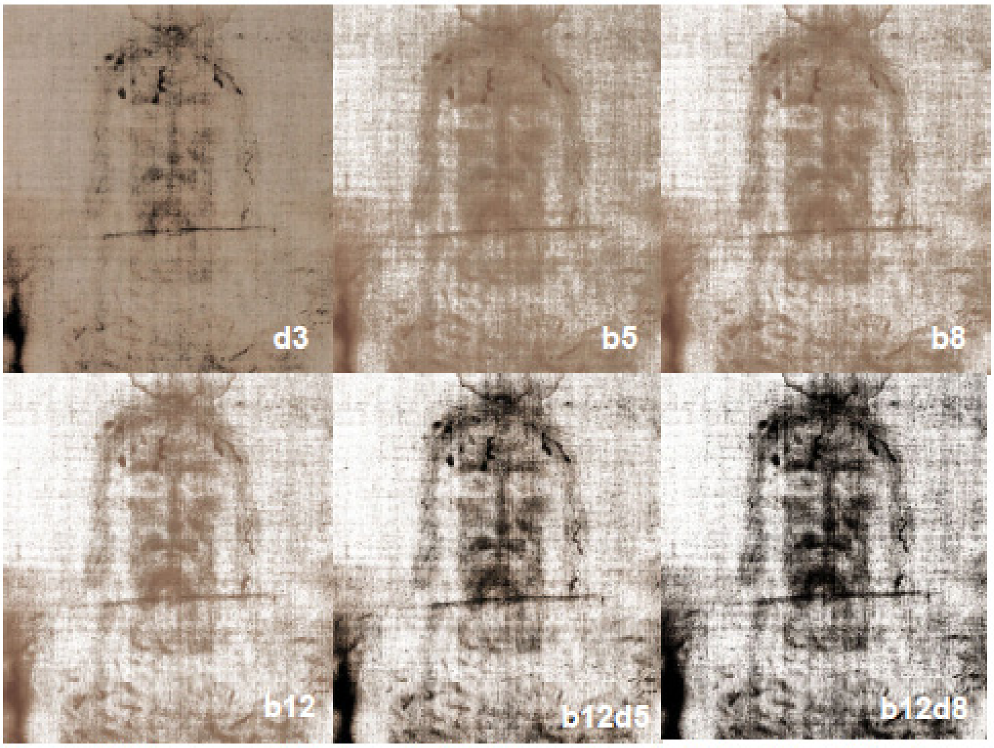

Figure 13 Sequence of black and white false color substitution into a 24 color indexed Face Crop image. Black is substituted for pixels with darkest luminance and white for pixels with brightest luminance.

Note that the darkest three pick up the darkest scorch and blood marks. The brightest five show the cleanest cloth and as more are turned white (see b12) the darkest indexed colors show the image, blood, scorch and miscellaneous dirt and off image marks progressively darker as more have black substituted for them. False color substitution is a nice technique for probing an image to see what features a particular indexed color is associated with.

7.7 *Combined Color/Luminance Intervals*

At this point it is clear from the previous work illustrated in Figures 12 and 13 that both color and luminance contain important information. In developing feature sets with which to characterize decision logics to classify pixels it could be important to combine measures on the unit vector color image and the luminance information. Color as characterized in the SCT is just a direction in the color space. So characterizing a color involves picking a direction and setting an angular interval. The luminance is a single metric on the image and one can set that with a single interval. So two metrics and two intervals, an angular one and a amplitude one form a simple specification. Figure 14

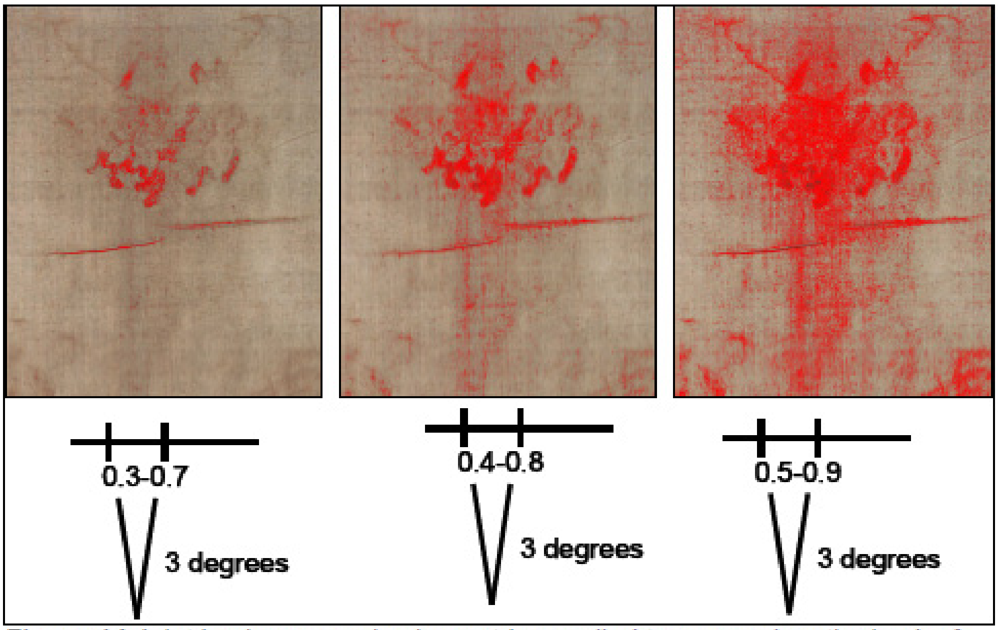

Figure 14 Joint luminance and color metrics applied to a wound on the back of the head substituting red as a false color to mark points that satisfy the criteria show how opening up the luminance interval around a fairly narrow blood color vector admits more and more pixels.

illustrates applying a single angular interval to set the color and then a second interval on luminance that allows an increasing number of the pixels that satisfy the color specification to be marked with a false color, here red. Each luminance interval is the same width but lighter. With more computing power this could be done dynamically.

7.8 *Stripe and Interstitial Isolation By Category*

As more and more exploration of pixel characterization continued one of the aspects that became clear is that extending the classification measures into pixel clusters would pose some problems. This is because of a feature that was apparent in the small samples shown in figure three. Each sample divides rather clearly into parallel stripes dictated by the weave of the cloth and about half are darker and half lighter. For the purposes of discussion the darker half is called here *stripes* and the lighter half *interstitials*.

It is straightforward to isolate the two by sorting the pixels by intensity, identifying the median and false coloring the darker half black and the lighter half white. Figure 15, below illustrates this strategy.

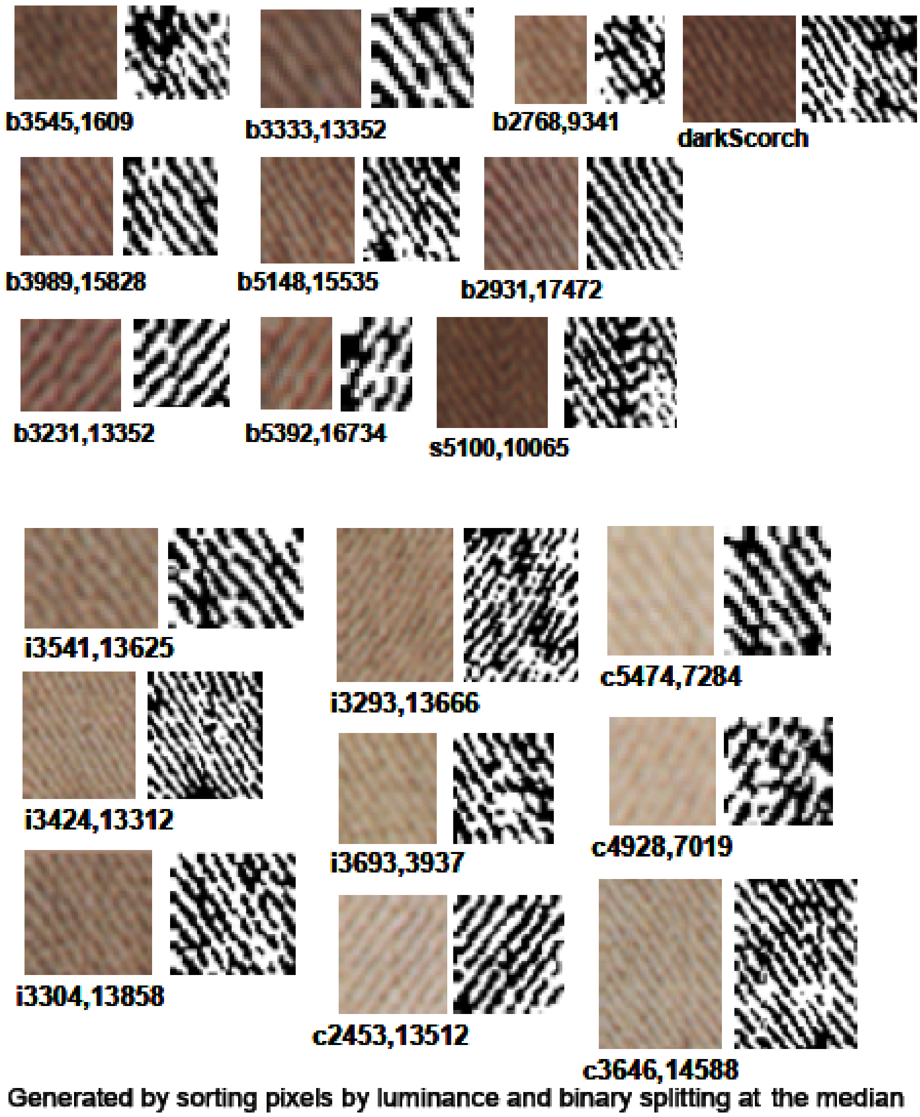

Figure 15 Separating stripe and interstitial weave markings in a sample by doing a false color binary split at the median. This is then used to generate separate statistics for the stripe and interstitial markings. Notation is the approximate image coordinates of the sample and a letter prefix to designate category b=blood, c=cloth, i=image, and s=scorch.

Isolating stripe and interstitials allows separate statistics to be calculated for each subcategory. Statistics calculated for each sample were the mean, standard deviation, maximum and minimum of the (φ, θ, L) phi, theta, luminance coordinates of each pixel in the sample. For example the statistics for b3545,1609 the first blood sample in Table 1 were:

Table 2 Statistics for sample b1 (3545,1609) Stripe and Interstitials

b3545,1609	s phi	s theta	s lum	i phi	i theta	i lum
mean	0.6234	1.1309	0.5226	0.6422	1.1252	0.6158
std-dev	0.0273	0.0228	0.0330	0.0235	0.0214	0.0357
maximum	0.6892	1.1813	0.5677	0.7076	1.1941	0.7123
minimum	0.5431	1.0476	0.4131	0.5727	1.0598	0.5681

7.9 *Transference of Metrics from One Site to Another*

Using statistics for stripe and interstitial one can false color a sample image such as the wrist wound shown in Figure 16.

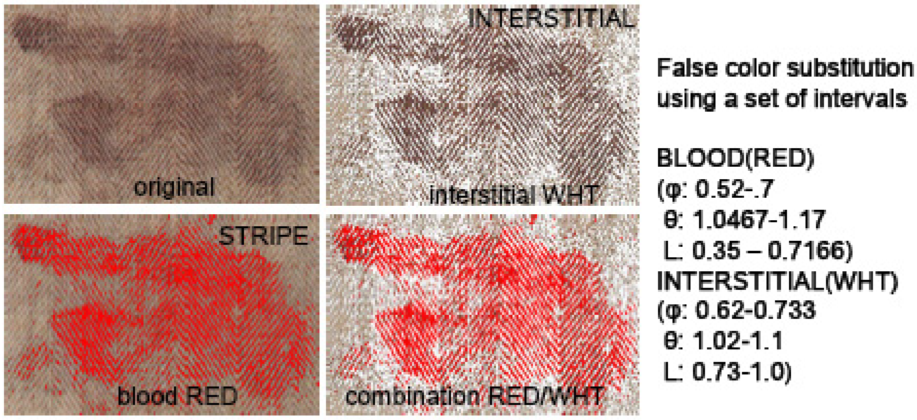

Figure 16 Applying false color to delineate stripe and interstitial markings in the wrist wound area. The color ranges overlap but the luminance ranges are disjoint. The subimages are the original, false color interstitial, false color stripe, and the combination. Pixels which did not meet the stripe or interstitial criteria were left their original values.

A test of the universality of a set of criteria for a category is the degree that it can be transferred without change to another part of the image. This was tested in a preliminary way by applying the same wrist wound criteria to the chest wound. The result is shown in Figure 17. The transference appears excellent. Blood is the darkest category except for scorch and so is largely unconfounded compared with image and cloth. Among the obvious next steps was to explore the application of statistics taken from different categories and seeking to create a pixel classification algorithm that would assign pixels to categories (initially blood, cloth, image) by their relation to the criteria.

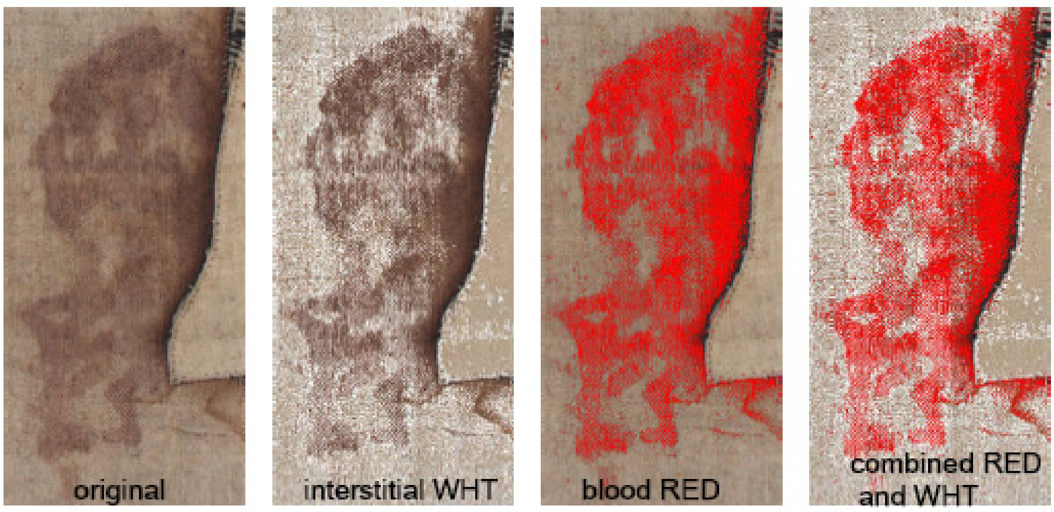

Figure 17 Application of Wrist Wound Criteria to Chest Wound to test Transference

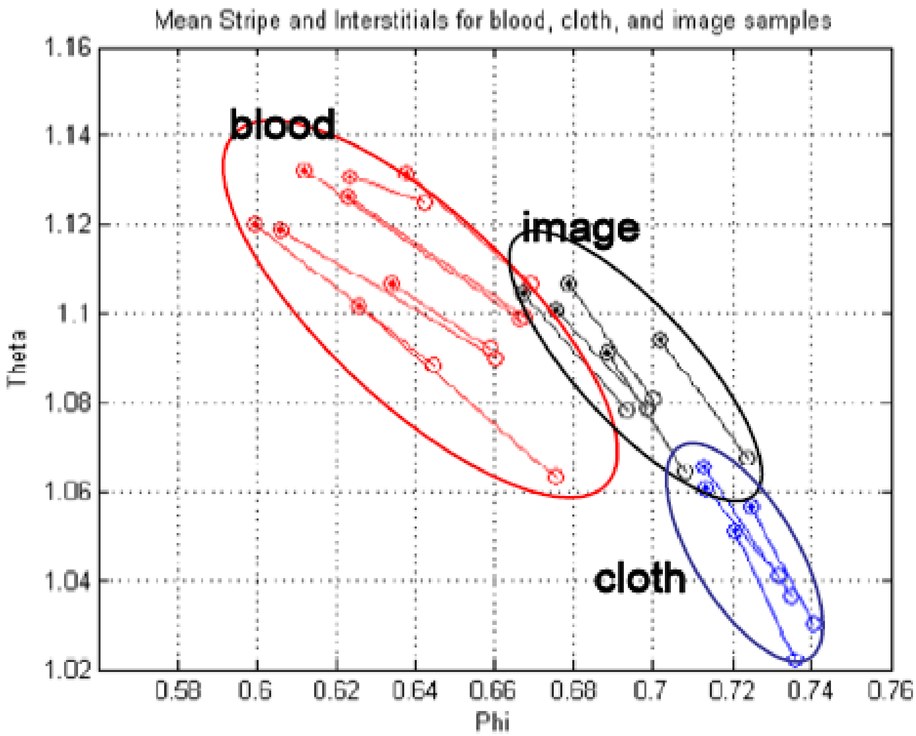

Figure 18 Mean of Stripe and Interstitials of blood, cloth and image samples plotted in Phi, Theta space. Circled dots are Stripe and circles without dots are Interstitials, lines connect Strip and Interstitial means of the same sample.

7.10 *Creating Stripe and Interstitial Metrics by Category*

Figure 18 plots the means in unit vector color space of the blood, cloth, and image stripe and interstitial samples. As noted earlier the color space is significantly compressed. The introduction of luminance data corrects this somewhat but it is clear that the color and luminance spaces overlap to some degree and to that extent using just that information will lead to multiple classifications for some pixels. This is an issue that will require further attention. Nevertheless there is sufficient separation to make it worthwhile to do an experiment in classifying pixels by category and injecting false colors that reflect the result.

7.11 *Using False Color and Closest Subcategory to False Color Images*

The data shown in Figure 18 was combined to produce a set of six criteria to specify the six categories: stripe and interstitial for blood, cloth and image. The measure used was the collective mean of the samples by category. The algorithm used was to assign each pixel to the category that it was nearest to. These were marked by false colors chosen to reflect the classification but also be reasonably representative of the category. The result is depicted in Figure 19.

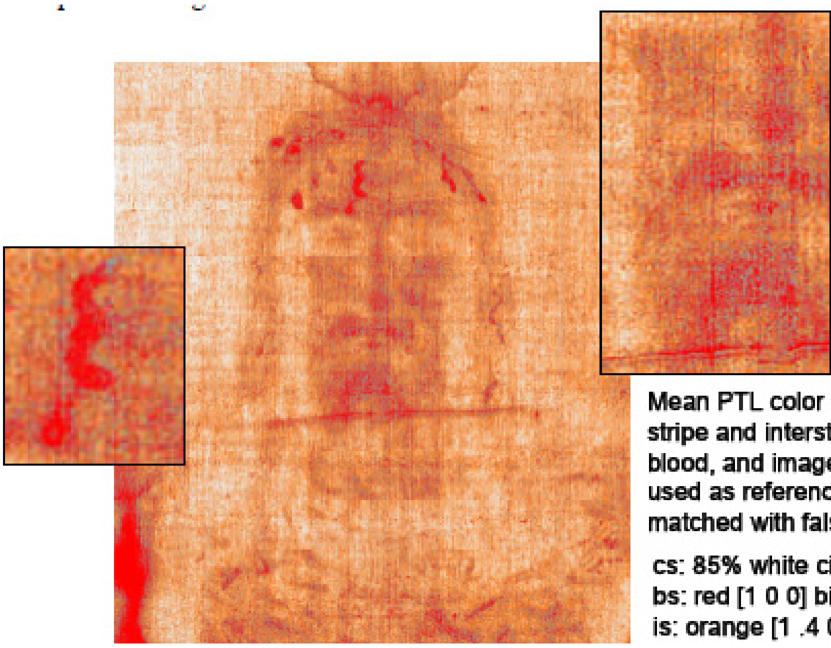

Figure 19 Pixel Classification by Nearest Category Mean and False Color Marking using false color assignment as noted.

The result of this first algorithm is rather satisfactory. There is no reason to believe that it is in any way optimum or that it can not be markedly improved. Notice that the scorch marks, for which there was no classifier, ended up being mostly classified as blood stripe pixels or blood interstitial pixels.

Figure 20 Pixel Classification by Nearest Category Mean and False Color Marking using false color assignment as noted in Figure 19.

8. Conclusions

The conclusions drawn from this exploratory effort are limited and raise a number of questions which cannot be answered at this point. The principal conclusions are:

8.1 *Shroud is characterized by a very narrow color/luminance space which makes classification by color alone difficult*

Figures 5.1, 5.2 and 18 all illustrate the fact that the color and luminance space is relatively small. This make categorization at the data/pixel level somewhat dubious and challenging. Nevertheless a reasonably good categorization was achieved by separating the pixels into stripe and interstitial categories.

8.2 *Contrast Stretching May Ameliorate this Problem (requires further work)* Reducing the image to the unit vector (Phi, Theta) color space and then contrast stretching it produces images that seem to have reasonably good color separation by category. See Figure 11 which illustrates this rather well. Blood, Cloth, Image and Scorch all appear to occupy different color ranges in this image. This suggests further research may provide a luminance independent classification algorithm.

8.3 *Region Analysis of Stripe and Interstitials Separately May Improve Segmentation* The process of separating category samples into stripe and interstitial subcategories substantially improved the pixel-based segmentation. This suggests that further analysis moving from pixel based to region based analysis could improve the classifiers by looking at the local region around the pixel being classified. This will require further research into transitions between regions and statistics within regions.

8.4 *The Image Area Shows a Strong Affinity with the Interstitial Blood Modes as well as having pixels that are likely evidence of blood on nose, mustache, and beard*
This observation that it is hard to distinguish image pixels from interstitial blood pixels may have some image formation implications. The interstitial blood mean and the image interstitials intimately overlap in the color space. It is unclear whether the pixels that classify as blood on the nose, mustache, and beard are actually blood possibly flowed onto these areas and then transferred to the shroud. Blood pixel classification in the crease by the chin may be fragments of blood trapped in the crease. It might also simply be a misclassification. Further investigation of this issue is suggested.

9. Recommendations for Further Work
Some suggestions for further work have been commented upon in passing during the account of this exploratory research. A few recommendations are singled out for special attention, but most are just recommendations to continue and deepen the current work which has been fruitful.

9.1 *Extend work by exploring more selective substitution schemes*
Expanding the selectivity of the category classifiers would include finding ways to further normalize the data, extending the analysis to include regional statistics and transitions between stripe and interstitial and category changes. This might include using the contrast stretched unit vector color space in conjunction with the original color space.

9.2 *Explore Fine Tuning using color and luminance gradients*
Related to 9.1 color and luminance gradients signal transitions between stripe and interstitial and between categories. Distinguishing between these would allow building segmentation contouring algorithms.

9.3 *Explore Stripe/Interstitial Relationship Further by Category*
The current work only provided some initial exploration of stripe and interstitial discontinuities for a limited number of categories, primarily blood, cloth, and image with a couple of examples of scorch in the mix. The markings on the shroud have numerous other potential categories which could be examined including hypothetical markings of speculative entities such as flowers, berries and seeds.

The work should be extended to other features such as scorch, water stain margins, detritus (dirt, droppings) generally overlapping or intruding into other regions or embedded as patterns of dirt in otherwise pristine regions. Further work in this area could extend both the categories and the criteria used to classify them. Applying objective criteria to discriminate and delineate sub-regions of the shroud by category would raise the overall objectivity of the analysis to a more analytical level.

9.4 *The Larger Effort*
The larger effort suggested in the abstract of comparing multiple images has hardly been touched

in the present paper. Two images were used, Schwortz 1978 and Durante 2000 with the latter getting the most attention. Comparing results using multiple images would strengthen the results supported only by single images. It would also contribute to a greater understanding of the properties of the shroud to explore features as they appear in different lighting on different occasions. Of particular interest would be a deeper understanding of banding as exhibited in both reflected and transmitted light images. Scanning classifiers that look for well defined interfaces between categories would be interesting. The scope of the studies done by the present investigator were limited by the equipment available. Recently an extraordinary set of images were made of the shroud
by HAL 9000 of Novara. [11] Images of this resolution will certainly give new insight into the nature of the shroud images and markings.

10. Acknowledgements
Four individuals should be acknowledged in helping the author advance this work.
Mario Latendresse [3] of the University of Montreal wrote a small measurement JavaScript application and called for more quantitative analysis of shroud images, especially applying the algorithm to multiple images. This got the current author thinking and working towards that objective.
Barrie Schwortz, the documenting photographer of the 1978 STURP expedition has been a never failing source of images and encouragement as well as friendship. His selfless and never tiring work to explore the shroud has been an inspiration.
Giulio Fanti provided the Durante 2000 high resolution scan which has been the source of most of the imagery used in this paper. He too has been a source of inspiration especially his work with the Shroud Science Group.
Dan Scavone is a historian, a friend, someone who has generated all sorts of inspirational ideas about the shroud which has helped flesh out a broader understanding

11. End Notes and References
1. Secundo Pia's shroud photograph is discussed at http://www.shroud.com/vanhels4.htm.
2. http://dsc.discovery.com/news/2008/02/28/shroud-of-turin.html
3. Mario Latendresse's JavaScript application can be viewed at his website http://www.iro.umontreal.ca/~latendre/shroud/shroudCal.html or at Ray Schneider's website http://www.bridgewater.edu/~rschneid/FocusProjects/Shroud/ShroudMeasure/shroudCal. html
4. This hypothesis is related to discussions conducted by the Shroud Science Group on the topic of the late Ray Rogers' hypothetical imaging mechanism that called for a chemical reaction in a thin layer of surface impurities resulting from the yarn processing.
5. The work here was done on a Sony VAIO laptop with 2 GB of memory. Several software tools were used including: Python and the PIL (Python Imaging Library), CVIPtools, PhotoShop Elements, ImageJ a processing tool written in Java, and MatLab (Matrix Laboratory Image Processing Toolkit). All of these except for PhotoShop Elements and MatLab are open source tools.
6. http://en.wikipedia.org/wiki/Segmentation_(image_processing) gives a short overview from the viewpoint of computer vision.
7. Durante 2000 is a 300 dpi scan 7632x24197 pixels, 541,031 KB RGB TIFF image of the shroud. The scan was provided to the author for research purposes by Giulio Fanti.
8. Umbaugh, Scott E., *Computer Imaging—Digital Image Analysis and Processing*, pg. 49-51 briefly discusses the SCT (Spherical Coordinate Transform). The SCT/Center color segmentation algorithm is described on pgs. 159-161.
9. The Python Imaging Library (PIL) can be downloaded from http://www.pythonware.com/products/pil/ The library is somewhat limited but it is easy to use.
10. Matlab image processing function definitions are derived from the *Image Processing*

Toolbox™ *User Guide* which can be found on-line at
http://www.mathworks.com/access/helpdesk/help/toolbox/images/ and in pdf form at
http://www.mathworks.com/access/helpdesk/help/pdf_doc/images/images_tb.pdf
11. See http://www.shroud.it/NEWS.HTM On 22 January 2008 an extraordinary mosaic scan of the shroud was made which when composited will be the most detailed image of the shroud ever taken. It is composed of 1649 individual images according to some accounts.
12. Mario Latendresse, http://www.iro.umontreal.ca/~latendre/ Barrie Schwortz, http://www.schwortz.com/bio.pdf
Giulio Fanti, http://www.dim.unipd.it/fanti/fanti-ingl.html
Dr. Daniel Scavone, Professor Emeritus of History at University of Southern Indiana.