Interactive Stereo Electron Microscopy Enhanced with Virtual Reality

E. Wes Bethel[a], S. Jacob Bastacky[b], Kenneth S. Schwartz[a]
Lawrence Berkeley National Laboratory
University of California, Berkeley
Berkeley, CA

ABSTRACT

An analytical system is presented that is used to take measurements of objects perceived in stereo image pairs obtained from a scanning electron microscope (SEM). Our system operates by presenting a single stereo view that contains stereo image data obtained from the SEM, along with geometric representations of two types of virtual measurement instruments, a 'protractor' and a 'caliper.' The measurements obtained from this system are an integral part of a medical study evaluating surfactant, a liquid coating the inner surface of the lung which makes possible the process of breathing.

Measurements of the curvature and contact angle of submicron diameter droplets of a fluorocarbon deposited on the surface of airways are performed in order to determine surface tension of the air/liquid interface. This approach has been extended to a microscopic level from the techniques of traditional surface science by measuring submicrometer rather than millimeter diameter droplets, as well as the lengths and curvature of cilia responsible for movement of the surfactant, the airway's protective liquid blanket. An earlier implementation of this approach for taking angle measurements from objects perceived in stereo image pairs using a virtual protractor is extended in this paper to include distance measurements and to use a unified view model [1].

The system is built around a unified view model that is derived from microscope-specific parameters, such as focal length, visible area and magnification. The unified view model ensures that the underlying view models and resultant binocular parallax cues are consistent between synthetic and acquired imagery. When the view models are consistent, it is possible to take measurements of features that are not constrained to lie within the projection plane.

The system is first calibrated using non-clinical data of known size and resolution. Using the SEM, stereo image pairs of grids and spheres of known resolution are created to calibrate the measurement system. After calibration, the system is used to take distance and angle measurements of clinical specimens.

1. INTRODUCTION AND PROBLEM STATEMENT

One of the objectives of current biomedical research is a better understanding of the form and function of the biological components that constitute organisms. One approach used to gain such understanding is observation of the system or biological component in question. Observations from different viewpoints and at different magnifications provide data about the structural and spatial relationships of small microstructures, and this data is often used in research and clinical applications. In particular, taking measurements of microstructures is important in understanding the pathophysiology at a cellular level of resolution.

A Scanning Electron Microscope (SEM) is commonly used to obtain high-resolution images of the surfaces of specimens. The SEM generates as output a 2D image of the specimen. As such, it is a straightforward matter to use a ruler or "paint program" to count pixels, and obtain a measurement of two-dimensional feature width or height from images of specimens. When the target features and desired measurements needed for research and clinical study are 3D, as are those described in our case study, 2D measurement methods lack sufficient accuracy. In this paper, we use the term "instrument" to refer to the SEM.

1.1 Escaping Flatland

In order to take 3D measurements, we need to somehow "escape the flatland" of 2D imaging in a way that allows us to use virtual sensors to take accurate measurements of features perceived in stereo image pairs. Previous work with similar goals, namely acquisition of 3D measurements, describes the use of 3D "virtual tape measures" for obtaining measurements of distance in CAD models of aircraft [2]. In our application, we use a similar approach for measurement, namely a virtual measuring device that is manipulated with one or more input devices. However, our task is somewhat more difficult since we are

a. [ewbethel,ksschwartz]@lbl.gov, Mail Stop 50F, 1 Cyclotron Road, Berkeley CA, 94720.
b. sjbastacky@lbl.gov, Children's Hospital Oakland Research Institute, 5700 Martin Luther King Jr. Way, Oakland CA, 94609.

measuring features that are visible in 2D images, as opposed to measuring parts of a geometry-based CAD model. In this paper, we use the term “virtual sensor”, or just “sensor,” to refer to the virtual measuring device employed within a virtual world to obtain measurement of distance, as with a caliper or tape measure, or of angles, as with a protractor. One of our goals in using virtual sensors combined with carefully designed view frusta is to avoid many of the troublesome issues associated with computing points in space using methods described in earlier microscopy literature [3].

To escape the flatland of 2D imaging, we need a source of 3D data to measure. Using eucentric tilting of the specimen in the instrument, we can obtain images from varying and known viewpoints. When the tilt angle is small, the viewing geometry produced by eucentric tilting closely approximates a binocular viewing geometry. The resulting images can be viewed in stereo, and 3D structure perceived by a viewer. Stereo imaging helps a viewer to escape flatland, but doesn’t directly address the problem of taking 3D measurements. So, we face a fundamental dilemma: do we attempt to derive 3D information from the images and use that derived information as part of the measurement process, or do we rely instead on a viewer’s ability to perceive depth from stereo and visually “fuse” stereo image data obtained from the SEM with stereo views of the measurement device generated by rendering software?

1.2 Depth Perception

There exists a large body of work on the subject of human perception of depth relationships. Registration of parallax cues is only one part of stereo depth perception. A recent survey work describes several interrelated factors [4]. These factors include binocular disparity or steropsis, motion parallax, and shape-from-shading. That these cues have cognitive synergy is well accepted, but the mechanics and relative weights of each is more of a mystery. In this paper, we have built our system around registration of binocular disparity, or parallax, in order to achieve a simple goal: point measurements of features perceived in stereo image pairs. In this paper, we focus on formulating a view model that ensures correct binocular parallax when viewing stereo-rendered geometry concurrent with stereo image pairs. We have not explored use of additional cues, such as shape-from-shading or motion parallax.

Motion parallax refers to the change in depth relationships or cues, such as occlusion, that occur when objects in a 3D scene are moved, or when the viewer changes location. Delving into this topic would require use of vision-based tools that compute a depth field or 3D models from two or more images obtained from different viewpoints. An alternative approach for implementing motion parallax would make use of Image Based Rendering (IBR) techniques [5]. IBR takes as input a number of images acquired from known viewpoints, then allows navigation in the scene. Depth cues could further be enhanced by coupling an estimation of depth disparity with IBR in order to implement per-pixel depth occlusion of observed features and virtual sensors.

Because the SEM imaging process is not based upon visible light, and is not changeable, there is little we can do to improve shape-from-shading cues. The shading apparent in SEM images is similar to the shading resulting from illuminance, but different is several respects. The images produced by SEM are the result of interactions between the electrons from the SEM beam and the atoms of the sample at and below the surface - some beam electrons are scattered, and some specimen electrons are dislocated by collisions with beam electrons. All of these scattered and dislocated electrons are collected to produce the most common type of SEM image, called the secondary electron image.

Regions of the specimen that are perpendicular to the electron beam do not produce many electrons, so these surfaces appear to be dark in shade. At edge regions, there is significant electron production, and this effect produces regions of brightness. Surfaces that are oblique to the electron beam produce intermediate numbers of electrons, and appear brighter than those perpendicular to the electron beam, but appear darker than edges. Another type of SEM image, created only from high-energy electrons, primarily those of the beam scattered by the specimen, gives more faithful shading than the secondary electron image. However, this backscattered electron image has the disadvantage that valuable information from the “shadows” is lost, and secondary electron images are almost universally preferred for interpretation. Secondary electron images contain features not present in backscattered images. The difference, using an illuminance analogy, is similar to the absence of ambient light in an image.

In this project, we have opted for an approach that capitalizes upon depth perception from binocular parallax cues. In the sections that follow, we first describe our “unified view model,” including its derivation and why it is central to generating correct binocular parallax, which is a prerequisite for accurate 3D measurement. We then describe how this system has been used in a research setting, first through a calibration phase, then more directly as part of a medical research study. We conclude with some discussion of the limitations of our approach and provide some suggestions for how the approach can be strengthened in future work.

2. UNIFIED VIEW MODEL

Our software creates stereo images from two types of information. The first type of information consists of stereo image pairs, which are obtained via eucentric tilting of a specimen inside the SEM. The second type of information consists of one or more geometric models used to create a visual representation of a virtual sensor. Our goal is to render both types of information within a single window, and in such as way as to have the depth cues resulting from binocular parallax appear to be consistent across both types of information. If this goal is achieved, the user of the system will be able to perform measurements of distance or angular relationship between features present in the stereo image pairs using the virtual sensors. In addition, a view formulation that is consistent across SEM tilting and sensor rendering will allow us to calculate distance and angular measurements in the world coordinate system of the sensors, without the need for coordinate system transformations. In order to achieve this goal, we must render the sensor geometry using a binocular viewing formulation that is exactly the same as that formed by the eucentric tilting of the specimen in the microscope. We refer to such a binocular viewing formulation as a "unified view model."

2.1 SEM View Formulation

To begin, we must first generalize the view model of the SEM. The generalization is based upon a general relationship between focal length, magnification and observable area. These three parameters are readily available from the instrument, and are sufficient to formulate a view frustum suitable for use in geometric rendering. Given a focal distance *f* and a magnification *x*, we can compute the field of view angle *a* as follows:

FIGURE 1. SEM View Frustum:

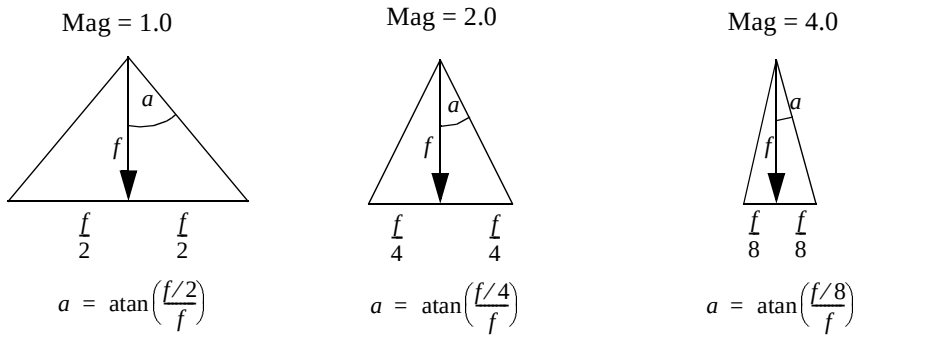

$$a = \operatorname{atan}\left(\frac{1}{2x}\right) \qquad (1)$$

(1) has the effect of decreasing the field-of-view half-angle *a* as magnification increases. However, (1) must be "tuned" to each individual instrument in order to produce the actual visible area of the instrument, rather than the visible area derived from (1), which is simply *2f/x* for a given magnification *x*.

In order to derive the per-instrument tuning constant, we require additional per-instrument information that consists of focal length *f*, and visible area *V* at some magnification *x*. Table 1, below, lists values for focal length, magnification, visible area computed by (1), along with the actual observable area from our instrument. In general, SEMs operate poorly at lower magnifications. As such, the magnifications of 1x and 10x are not used in practice, but are useful for the purposes of computing the field of view. The last column of Table 1 is the value of *a* required to produce the actual visible area.

TABLE 1. Comparison of Equation 1 Predictions with Actual Instrument Characteristics

Focal Length *f*	**Magnification *x***	**Visible Area (Eq. 1)**	**Actual Visible Area**	**Computed FOV, *2*a***
13mm	1	26.0 mm	88.9 mm	26.56*2 = 53.13 degrees
13mm	10	2.6 mm	8.89 mm	2.86*2=5.72 degrees

We can tune Equation 1 to match the characteristics of a given instrument with a constant scaling factor *w*:

$$a = \operatorname{atan}\left(\frac{w}{2x}\right) \qquad (2)$$

The constant *w* is used to define the proportion of an ideal view frustum to the actual frustum as implemented by each particular instrument. The constant term *w* may be computed on a per-instrument basis as

$$w = \frac{xV}{f} \tag{3}$$

where *f* is the focal length of the instrument, *x* is the magnification, and *V* is the length of one side of the square sample area visible by the instrument at the specified level of magnification. For the instrument used in this study, *w* evaluates to approximately *6.83*, where the focal length *f* is 13mm, and the area visible at single power magnification is a square with sides that are *88.9mm* in length. When *x=10*, representing a 10x magnification, the size of the area that is visible is a square that is 8.89mm on a side.

In this formulation, we make the following assumptions:

1. There exists a linear relationship between a change in magnification and a change in the amount of visible area seen at each magnification, i.e., *1/x*.
2. The shape of the area seen by the instrument at each magnification is a square.

The view formulation for a given SEM instrument is a tuned variant of the ideal view frustum. Users of our software must provide values for the focal length *f* and visible area *V* seen at some magnification *x*. These parameters are known constants for each individual instrument. In some cases, when the parameters are not known, they may be estimated using calibration images that contain features of known size that are obtained at known magnifications.

3. Eucentric Tilting and Binocular Stereo View Formulation

The previous section describes formulation of a monocular view frustum that matches the view frustum of the instrument. This frustum may also be used for "space-registered" geometric rendering, which means that occlusion and perspective foreshortening cues will be the same for artificially rendered geometry as they are for an image obtained from the instrument using some set of parameters. The computed frustum has the desired property of having the identical amount of viewing area as the instrument for given values of magnification and focal length. In this section, we describe the binocular stereo viewing geometry used in acquiring images from the instrument, and the corresponding view frustum used for geometric rendering.

Eucentric tilting, or rotation, occurs about an axis that intersects the area being observed within the specimen, and is intended to avoid the introduction of a translational component into the transformation [6]. In contrast, a traditional binocular view formulation consists of a translation of viewpoint. The following figure illustrates the conceptual difference between eucentric tilting, and a binocular view produced by translation.

FIGURE 2. Stereo Frustum from Eucentric Tilting and Translation-based Binocular Views

Eucentric Tilt Translation-Based Binocular View

In Figure 2, we have used the unlikely value of forty degrees of tilt in order to exaggerate the difference of tilting as compared to translation. In practice, a tilt value of five degrees, or interocular separation in the case of the binocular model, produces a strong stereo effect that is not overpowering[c]. The geometry of these two stereo views is similar in the geometric sense, i.e., similar triangles. They differ in actual focal length, or distance from the view point to the viewer. In the eucentric tilt model, the focal length remains constant at *f*. In the binocular view model, the focal length increases from *f* to approximately *f* +

c. Overpowering can be defined as such a high degree of separation that a viewer is unable to successfully fuse the stereo images.

$f*sin^2(a/2)$. For typically small values of *a*, this difference is negligible. We point out this difference because rendering software typically uses a binocular view formulation formed by translating the eyepoint for each stereo view, whereas stereo views are obtained from instruments by eucentric tilting.

It is also possible to use translation of the specimen perpendicular to the electron beam within the instrument in order to obtain a stereo pair that contains binocular parallax cues. Doing so has several practical drawbacks. One drawback is a significant sacrifice of resolution in the stereo pairs produced by the SEM. When using translation, the feature of interest will appear close to the left edge of one image, and close to the right edge of the other. The amount of resolution sacrificed is a function of the imaging resolution of the instrument and the magnification level. Another drawback is the additional calibration steps required in order to produce a targeted amount of interocular separation. This accuracy is required in order for our unified view model to be valid. We mention specimen translation as an option, as not all SEMs are capable of eucentric tilting.

4. UNIFIED VIEW MODEL IMPLEMENTATION

So far, we have described the formulation of a monocular view frustum that reproduces the frustum formed in an SEM by a set of instrument parameters. We have also described the difference between the stereo frustums that result from eucentric tilting as compared with translating the each eyepoint while maintaining a constant focal point in order to produce a traditional binocular view formulation. In this section, we will merge these two view formulations in the context of a software implementation. We shall see that two view formulations are needed in order to correctly render acquired stereo image pairs along with stereo views of sensor geometry. In addition, we shall show a simple transformation that is used to generate measurements from sensor coordinates that reflects distance in the coordinate system of the SEM.

Recall that the stereo image pairs from the SEM have been obtained via eucentric tilting of the specimen in the instrument. Such images can be viewed in an anaglyph stereo format using off-the-shelf paint programs and low-cost red-blue glasses. No special view formulation is needed to view these images. However, since we will be merging stereo views of sensor geometry with stereo image pairs, we need to define an appropriate 3D view frustum so that parallax cues are consistent between the stereo sensor geometry and features present in the stereo image pairs. The consistent cues will ensure accurate visual alignment of sensor geometry and perceived 3D features.

4.1 Stereo Rendering of Sensor Geometry

Our objective in deriving a binocular view formulation for rendering sensor geometry is to match the view frustum parameters of the instrument, including the stereo separation defined by the tilt or translation used to obtain stereo image pairs. By faithfully reproducing the eucentric tilt view formulation, within the limits of accuracy possible using a binocular stereo view formulation, we are able to create a high degree of accuracy, or registration, between the SEM/tilting and sensor geometry view models.

In deriving the binocular view frustum used for rendering stereo views of sensor geometry, we use Equation 2, along with the constants defined by instrument parameters and Equation 3, to define the shape of our view frustum. We reproduce the effect of increased magnification, and the corresponding reduction in perspective foreshortening, by multiplying the focal length of the frustum by the magnification scale used to obtain the images. The field of view narrows with increasing magnification. The combination of increasing the focal length and narrowing the field of view with increasing magnification has the effect of keeping the area of the projection plane nearly constant as measured in world coordinates. At focal length *d2*, the shape of the view frustum more closely approaches the shape of a parallelepiped than at focal length *d1*. At distance *d1*, the frustum shape is closer to being right pyramid (these frustums have not been truncated to reflect placement of near and far clipping planes). This type of monoscopic view frustum is used for rendering the stereo image pairs obtained from the SEM. Each source image is texture mapped onto a quadrilateral that is placed at the focal plane *P*.The view frustum shape approaches a cube as magnification increases to infinity, as illustrated in Figure 3.

FIGURE 3. View Frustum Shape as a Function of Focal Length

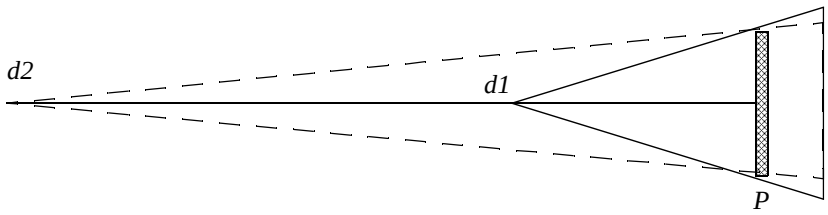

Features of the OpenRM Scene Graph software[d] used for scene management and rendering allow us to have the left eye image texture and geometry appear in the left channel, and the right eye image texture and geometry in the right channel, even when using a monoscopic view formulation. If no inter-image registration is required, then the textured quads are rendered at exactly the same location in space. There are no depth occlusion issues to resolve, as each quad is rendered into a separate stereo channel. If manual image registration is needed, one or both of the quads can be independently translated within the view plane in order to achieve the desired registration.

Because we are using a traditional binocular viewing formulation when rendering stereo views of sensor geometry, the eye position for each view is offset from the center of projection by an amount that is proportional to the focal length. In other words, in order to maintain a fixed interocular separation angle with respect to the focal point, each eyepoint must be translated by an amount that is proportional to focal length. This is illustrated in the following figure. Assuming an interocular separation angle of *a* and focal length of *d1*, each eye point is offset from the line-of-sight vector by a distance of *d1 * sin(a)*. When the focal length is *d2*, each eye is offset from the line-of-sight vector by a distance of *d2 * sin(a)*. Special adjustments to the near and far clip planes are required that take into account the amount of interocular separation and the focal length. Without these adjustments, sensor geometry that is placed at the same location as the texture-mapped quads containing SEM stereo images would be clipped away by the renderer.

FIGURE 4. Maintaining a Constant Interocular Separation Angle with Increasing Focal Length

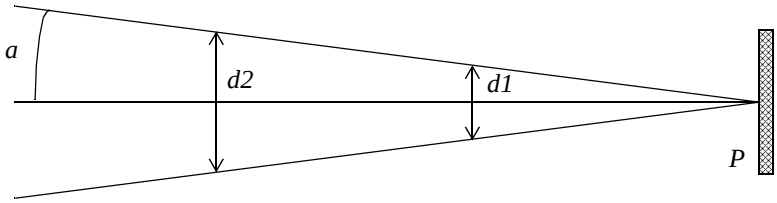

5. MEASUREMENTS WITH VIRTUAL SENSORS

Thus far, our discussion has focused on maintaining accurate proportions between two view models. One view model results from eucentric specimen tilting, needed to obtain stereo image pairs from the SEM. The other is the binocular view model needed to create stereo views of sensor geometry. The formulations we have presented represent an accurate registration of these two view models. For the purposes of measuring only angles, no additional work is required. In order to accurately measure distances, we need to define a transformation between the pixel resolution of the SEM images and the virtual sensor world coordinate system.

5.1 SEM to Sensor Coordinate System Transformation

The explicit transformation between sensor coordinates and SEM coordinates can be avoided by defining appropriate vertex values for the texture-mapped quadrilaterals. The size proportions for each quadrilateral are adjusted to reflect the proportions of the corresponding input image. Thus, input images need not be exactly the same size. We do require an additional parameter that indicates the height in pixels of the image created by the instrument. The height in pixels is used based upon the assumption that the eucentric tilt axis is vertical, in order to accommodate left and right eye views. We also assume that the eucentric

d. OpenRM Scene Graph is an LGPL Open Source scene graph system for use on all OpenGL systems, including Win32, Linux and Unix that supports multistage, pipelined-parallel rendering. Source code and documentation can be obtained at the OpenRM website, http://openrm.sourceforge.net/.

tilt axis is located in the center of the image. Accordingly, the center of the texture-mapped quads are placed directly at the center of projection, and in the projection plane. Alternate formulations are possible, although we have not explored them.

TABLE 2. Instrument Parameters Required for Scene and Frustum Construction

Parameter	Description
Focal Length *f*	The focal length of the instrument in millimeters.
Visible Area *V*	For a magnification value of 1.0, the value of *V* is the square root of the visible area. If the amount of space visible at magnification 1.0 is V^2 millimeters, we require a value of *V* to be used for this parameter.
Pixel resolution *P*	Instruments typically produce images of a constant size, regardless of magnification. We assume the eucentric tilt axis is positioned in the center of the output image, and that it is tilt axis is vertical in orientation. The parameter *P* height in pixels of the image produced by the instrument.
Magnification *x*	The magnification value used when obtaining a set of SEM stereo images must also be provided.

The parameters *f, x, P* and *V* are combined to produce a scale factor that can be used to transform units in sensor coordinates to micrometers, the unit of measure in the SEM.

5.2 Calibration and Verification

Prior to using the caliper sensor on actual medical data, we first performed a calibration step to verify that the view frustum models were indeed correct. Figure 5 shows measurement of a 1-micrometer calibration marker inserted by the SEM imaging software. At this resolution, we see our measured caliper distance is 1.0058 micrometers. Visually, the length of the caliper and the calibration marker are identical, but the measurement differs by less than one percent.

FIGURE 5. Measurement of a Calibration Marker

Figure 6 shows a zoomed-out view of the same SEM stereo pair used in Figure 5. This set of calibration images was obtained by imaging an etched copper mesh at 5000x magnification. The copper mesh contains 2000 grid squares per inch, for a grid spacing of about 12.7 micrometers. Our caliper measures the vertical spacing at about 12.766 micrometers. This measurement differs from the expected results, 12.7 micrometers, but by less than one percent. Unlike measurement of the calibration marker, the copper screen appears in stereo, hence there is some amount of ambiguity introduced into the measurement process as human perception begins to have some influence on the accuracy of the measurement. In addition, at this level of resolution, the boundaries of the copper mesh appear to be quite rough and bumpy.

FIGURE 6. Measurement of Etched Copper Screen Spacing

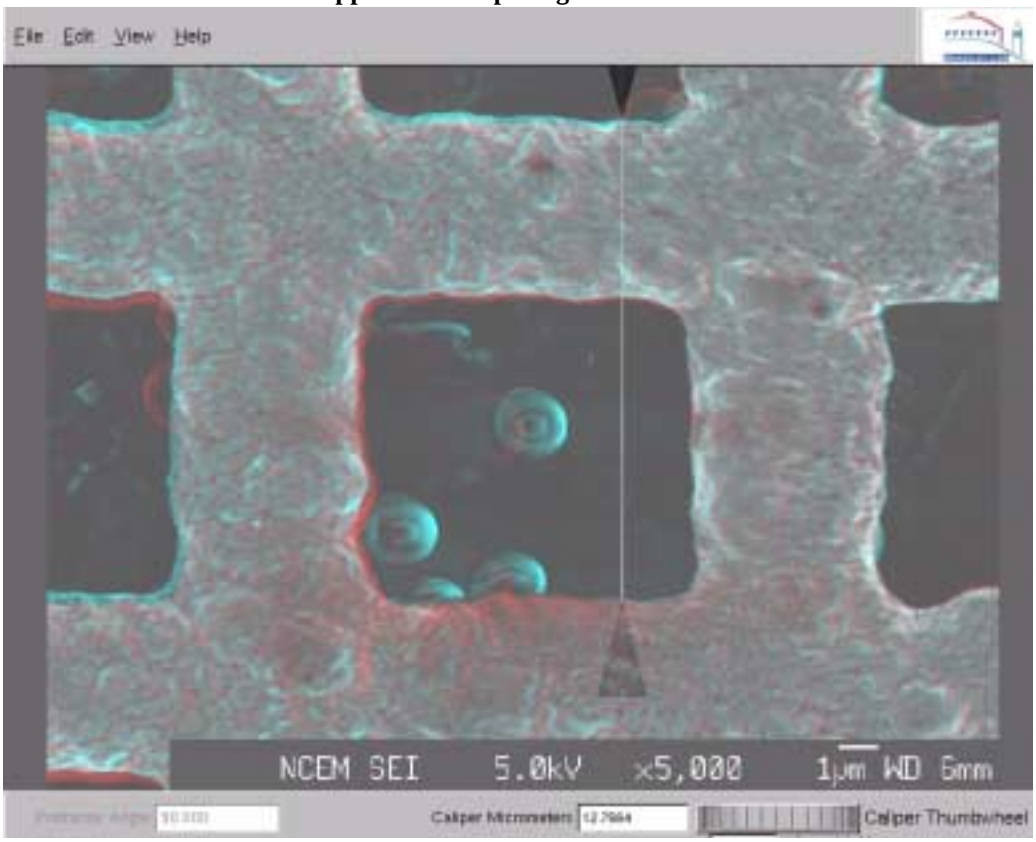

To increase the accuracy of magnification calibration in both X and Y directions, we acquired images of a calibrated 2000 mesh per inch grid at magnifications of 500 and 100, which provide 10 and 50 grid squared per row. This reduces the percentage error due to the irregularity of the etched grid edges. Polystyrene spheres of calibrated diamter are used to validate Z measurement. In these instances, the caliper measurements were within one percent of the expected results. With this level of accuracy, we proceeded to perform measurements of samples used as part of a medical study. The medical study focuses upon surface tension within bronchial passages.

6. Application to Biomedical Research

Surface tension is the force in the plane of a surface that tends to minimize the area of the surface. Stretching the surface, as in breathing, requires work. The work of breaking can be considerable: we would be unable to keep up with the energy demand if it were not for a substance called pulmonary surfactant, which reduces surface tension by two orders of magnitude. This substance is produced by the epithelial cells lining the alveoli, which are airsacs at the terminal portion of the lung's bronchial tree. The active component of surfactant is a medium-small molecule of phospholipid that has a water attracting and water repelling end. The phospholipid rapidly inserts itself into the surface of the liquid film that lines the alveoli and airways in the lung. The molecules line up and create a film that resists reduction in surface area of the film. This is the way that soap bubbles, for example, are stabilized and prevented from collapsing.

In certain diseases, such as Respiratory Distress Syndrome of Prematurity, the lung does not make sufficient surfactant. The airsacs and airways of babies suffering from this disease tend to collapse and fill with water. As a result, a significant amount of effort is required in order to breathe. Mechanical respirators help take over some of the work of breathing, but patients benefit significantly from therapeutic surfactant administered to their lungs.

The development of the next generation of pulmonary surfactants is facilitated by being able to measure surface tension in the lung. We can evaluate the effect of various synthetic surfactants on surface tension along the airways and alveoli. We measure surface tension in the lung by depositing indicator droplets of liquid fluorocarbon at the surface of the airway using a mist of aerosolized liquid administered to the trachea and lung. We then freeze the tissue and examine it in the low temperature SEM

at -180^{o}C. At this temperature, the fluorocarbon drops are solid, and their shape can be imaged with the electron beam. Magnifications in excess of 10,000 times can be readily obtained of uncoated, frozen drops on the surface of the cold airway.

Figure 7 shows a fractured droplet of fluorocarbon deposited on the surface of the trachea. The height of the relatively flat drop is measured with the virtual caliper to be approximately 15.5µm high. The width of the droplet is considered the average of the major and minor axes of the roughly elliptical drop, whose border can be barely distinguished in this micrograph. The width of this droplet is approximately 256 µm. The ratio of height to diameter is inversely proportional to surface tension of the surface upon which the droplet rests. In this case, surface tension of the surface of the trachea is higher than the surface tension of the fluorocarbon (16 dynes/cm) so the drop is pulled flat. We expect surface tension to be different at other places in the lung, and are beginning to use the software described in this paper to measure the variation in surface tension as a function of position along the airway in the mammalian lung.

FIGURE 7. Height Measurement of Fractured, Frozen Fluorocarbon Droplet

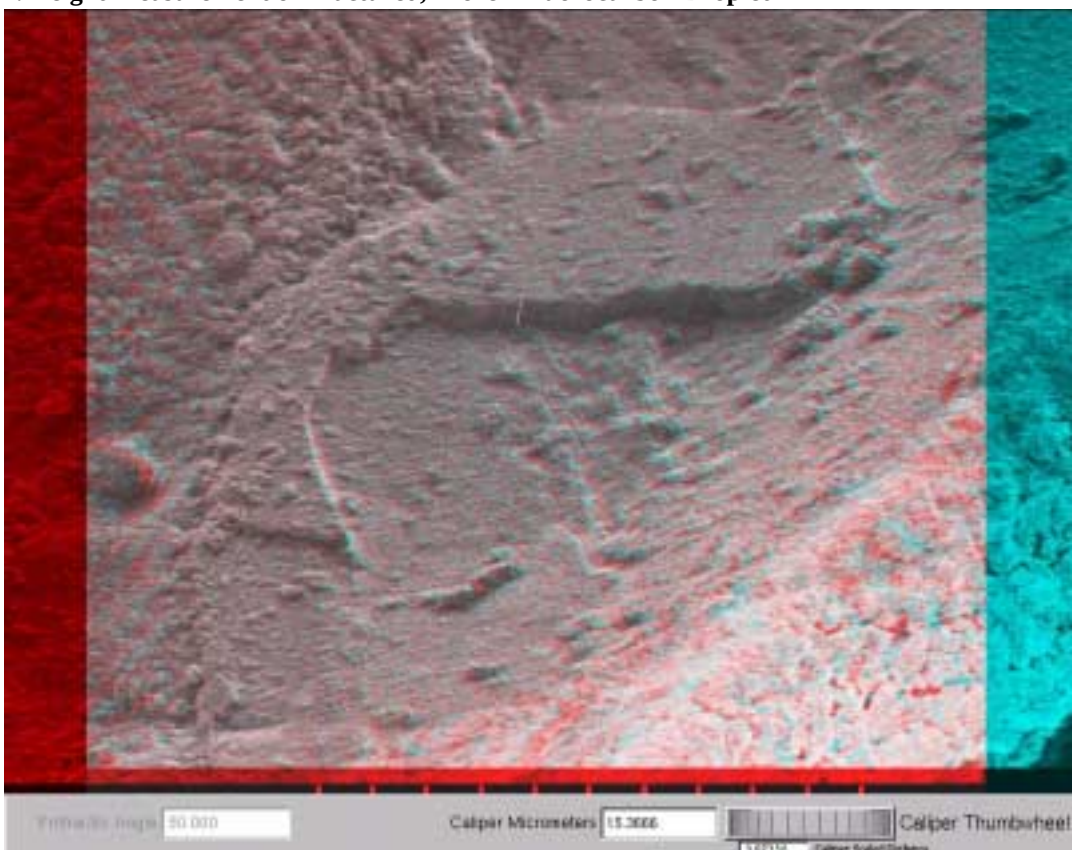

7. CONCLUSIONS AND FUTURE WORK

In this paper, we have focused upon formulation of a framework for ensuring a correct binocular parallax when viewing two different types of information. One type of information consists of stereo image pairs obtained from an SEM. The other type of information is a stereo rendering of a geometric model of virtual sensors. The sensors are used to interactively take point measurements of distances and angles of features observed in SEM stereo image pairs.

The notion of a unified view model is central to our system. The unified view model defines a binocular stereo view frustum that ensures the binocular parallax cues are consistent across SEM images and stereo views of geometry. When the binocular parallax is the same for both types of information, it is possible to use a single scale factor to transform between units in the world coordinate system of the sensor geometry and the SEM imaging coordinate system. Such an approach avoids a significant amount of coordinate system conversion.

We have narrowly focused upon binocular parallax to produce depth cues for a human observer. There are other depth cues that contribute to depth perception. These include shape-from-shading and motion parallax. Improvements to our system

would first include use of motion parallax. To create motion parallax, additional work would be required. Image-based rendering techniques, including use of dense depth maps, would aid greatly in providing motion parallax cues to improve depth perception. Shape-from-shading improvements are not possible at the SEM source, but may be possible if applied to 3D models extracted or derived from SEM source images.

The system we have described is useful for obtaining point measurements of 3D structures perceived in stereo image pairs produced by an SEM. The unified view model ensures that the binocular parallax is consistent across image acquisition and stereo geometry rendering. With such consistency, distance measurements are straightforward to perform. During use of the system, we often found ourselves wanting additional stereo pairs of specimens in order to cross reference and further validate results. Initial results of our system are promising, and we intend to perform further testing and validation with a larger set of calibration images.

8. ACKNOWLEDGEMENTS

This work was supported by the Director, Office of Science, of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098, and by the California Tobacco-Related Disease Research Program, Grant Number 7RT-0133A.

9. BIBLIOGRAPHY

[1] W. Bethel and S.J. Bastacky, “Measurement of Perceived Objects,” in Proceedings of IEEE Visualization 99, Late Breaking Hot Topics, IEEE Press, 1999. LBNL-43446.

[2] H. Sowrizal, “Interacting with Virtual Environments Using Augmented Virtual Tools,” in Stereoscopic Displays and Virtual Reality Systems, SPIE 2409, San Jose CA, 1994.

[3] S. Ghosh, “Electron Microscopy: Systems and Applications,” in H.M.Karara (ed.); Non-Topographic Photogrammetry, 2nd ed., Falls Church, VA: American Society for Photogrammetry and Remote Sensing, 1989. Pp. 187-201.

[4] D. Boyd, “Depth Cues in Virtual Reality and Real World: Understanding Individual Differences in Depth Perception by Studying Shape-from-shading and Motion Parallax,” BA Thesis, Department of Computer Science, Brown University, Providence RI, 2000.

[5] L. McMillan and G. Bishop, “Plenoptic Modeling: An Image-Based Rendering System,” in Proceedings of *Siggraph 95*, Computer Graphics Annual Conference Series, Addison-Wesley, 1995.

[6] E. Chatfield, “Introduction to Stereo Scanning Electron Microscopy,” in: M.A. Hayat (ed.): Principles and Techniques of Scanning Electron Microscopy:Biological Applications. New York: Van Nostrand Reinhold, 1974. Pp. 47-88.