A Hand-held Optical Surface Scanner for Environmental Modeling and Virtual Reality

R. B. Fisher, A. Fitzgibbon, A. Gionis, M. Wright, D. Eggert
Department of Artificial Intelligence
University of Edinburgh

1 Introduction

One of the first impressions one has when observing or using current "virtual reality" systems is that of impoverished environments. Based largely on CAD design systems, the surface shapes are simple, and the coloring and shading is limited to the current sophistication of computer graphics. Unfortunately, the real world is much more complex and random than can be expressed by such systems. What is missing includes stains, dents, random outgrowths, natural and individual variations, subtle colors, natural shape and reflectance textures, and anything else that makes the world "real" rather than "ideal". While one could construct a VR model of e.g. a real office, it would be very time-consuming to create all of the real detail.

Computer-based environmental scanning has addressed some of these issues, but the sensors used are typically vehicle-based (i.e. for vehicle navigation) and thus not suited for easy use, nor can they readily access all viewpoints around a part or in a portion of a scene in order to obtain a complete description. Range sensors are generally used to obtain surface shape measurements, and some of these sensors also acquire a reflectance description of the scene at a single frequency by analyzing the strength of the observed range signal (e.g. as in the Perceptron sensor[21]). Using an additional color camera to acquire "true" surface color descriptions (e.g. as in the Cyberware sensor[5]) introduces two problems: 1) registration of the color image data with the surface range information and 2) how to extract the "true" color or reflectance from the observed light, which is a function of the light source characteristics, light position and surface shape.

This paper describes a "logical sensor" [7] whose purpose is to recover the shape and color characteristics of a scene in a form usable for processes such as virtual reality and environmental modeling. The simplest non-technical description of the sensor is that it is an inverse paintbrush, acquiring into a computer database the shape and color of surfaces as one "paints" over them with the "brush" of the sensor.

The sensor is hand-held rather than robot-held (unlike e.g. the 3D Scanners [1] sensor), which makes it easy for a human to use, and thus avoids robotics issues that have helped confound previous research (e.g. robot manipulator or vehicle control, navigation, collision avoidance, position estimation, etc.). Surface shape is measured by extracting range information from structured light triangulation. The sensor uses the same color camera for both the range and color measurements, thus eliminating the problem of having to register two images. The pose of the hand-held sensor is determined using a global position detection device, which then allows transformation of the sensor's measurements into the global coordinate system, thus simplifying the problem of having to correspond and integrate multiple measurements from a moving sensor. The sensor projects broad-spectrum illumination from a close distance in conjunction with camera automatic gain control and white balancing in order to obtain reasonable R/G/B color estimates.

The main innovations of the scanner are:

1. It is hand-held, whereas previous scanners have been bulky items.
2. A single color camera is used instead of two or more.
3. Surface color is obtained as well as surface shape.

2 Theory

2.1 Block Diagram and Overview

The sensor is a hand-held device shaped much like a telephone handset. Figure 1 shows the block diagram of the sensor in cross-section (not to scale). At the bottom of the sensor is a compact two-dimensional RGB color CCD TV camera. Next to this is a directional bright broad-spectrum light source. The middle of the handset is the grasping point. At the top of the handset is the structured-light projection device. The device used for global position sensing is attached to the handset.

The working area of the sensor depends on the overlap between the camera's field of view, the area illuminated by the broad-spectrum light source and the projected structured light pattern. For the components used here, this results in a working area of about 10-15 cm in diameter at a distance of about 20 cm from the handset. The stripe analysis region and color acquisition region are slightly offset from each other.

The handset is connected via cabling to other instruments to return video from the color camera and measurements from the position sensing devices. The cable also carries power for the sensors, light sources and position sensing devices.

The video signal is connected to a PAL to RGB converter, which is then connected to a digitizer. The digitized image is processed by a standard computer, which also acquires the measurements from the position sensor. A color graphics display shows the 3D surface points as they are acquired. The actual components used are given in Section 3.

2.2 Surface Position Measurements Relative to Sensor

The surface position portion of the sensor is based on the standard technique of triangulation of a projected light pattern [3]. In this subsection, all quantities are represented in the

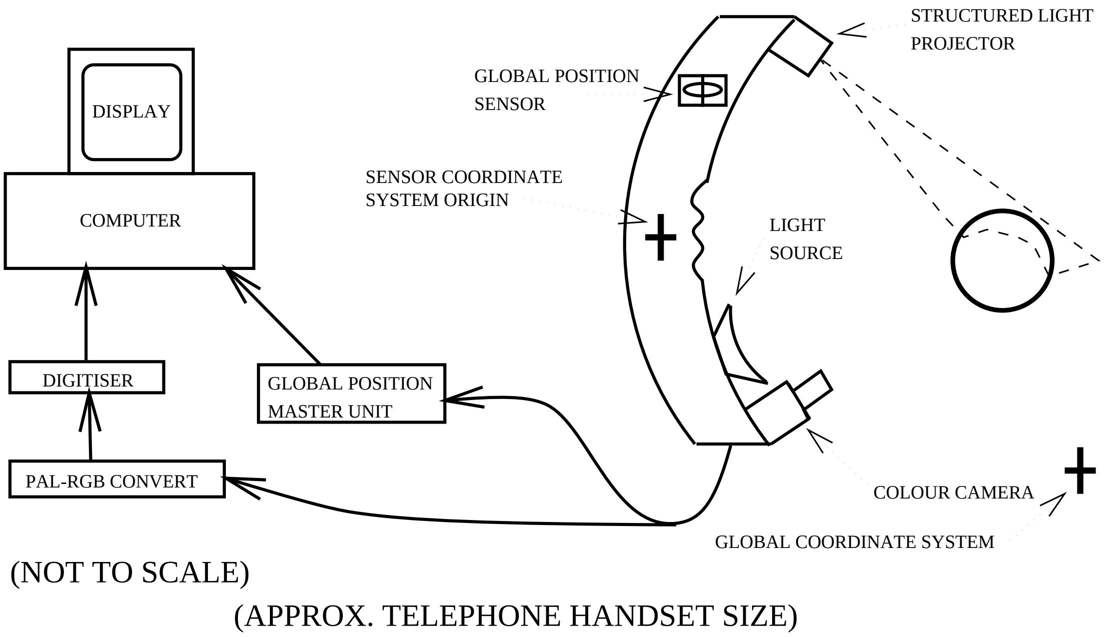

Figure 1: Block diagram of hand-held surface scanner

sensor coordinate system. Here, the sensor system is a coordinate system rigidly associated with the hand-held unit. It is not the same as the traditional "camera coordinate system", which is associated with the color camera. However, as there is a rigid transformation between the camera, the laser, the global position sensor, and the "sensor coordinate system", we shall just use the sensor coordinate system for describing all features. Finding the relation between these components, and between the sensor and world coordinate systems is a calibration problem not described here.

Suppose that the projected light pattern is a set of light rays

$$\vec{f}(\alpha, \delta) = \vec{l}_\alpha + \delta \vec{m}_\alpha$$

where $\vec{l}_\alpha$ is the center of the α^{th} light ray projection, $\vec{m}_\alpha$ is the direction of the α^{th} ray and δ is the distance along the ray.

Suppose that this light pattern is observed at a set of points (i, j). Here we use the term "image points" instead of "pixels" to allow non-integer image position measurements made through subpixel techniques [19]. For simplicity, the image points are described by their 3D position in the image plane, rather than 2D pixel image plane coordinates. Standard camera calibration methods can be used to relate 2D pixel positions to the 3D positions.

Each observed image point has a known viewing direction $\vec{v}(i, j)$ through a central pixel $\vec{c}$ (here we use a pinhole camera model), so that the true surface point projected onto each image point lies on the ray

$$\vec{g}(i, j, \mu) = \vec{c} + \mu \vec{v}(i, j)$$

where μ is the distance along the ray.

Intersecting $\vec{f}(\alpha,\delta)$ and $\vec{g}(i,j,\mu)$ gives the 3D coordinates $\vec{x}(i,j)$ of the surface point simultaneously illuminated by the structured light and observed by the sensor. The position $\vec{x}(i,j)$ is given in the scanner coordinate system, as are all other positions used above.

Using a light pattern instead of a single point of light allows simultaneous acquisition of many points from a single image. Most triangulation range sensors similar to this one use one or more planes of light generated using a spread laser beam.

Offline calibration methods are used to estimate the values $\vec{l}_\alpha$, $\vec{m}_\alpha$, $\vec{c}$ and $\vec{v}(i,j)$.

2.3 Integration of Multiple Shape Point Measurements

Each set of point measurements is acquired using a single camera image, but the sensor is moving over time. In order to construct a complete description of an object or a scene, the measurements must be collected in a common reference frame.

The global position sensor attached to the hand-held scanner provides a global position and orientation estimate for the scanner's local coordinate system origin. Then, if, at time t the position of the scanner is $\vec{p}^{(t)}$ and the transformation that maps orientations from the local to the global coordinate system is $R^{(t)}$, then the global position of every point $\vec{x}^{(t)}(i,j)$ measured by the scanner at time t is

$$R^{(t)}\vec{x}^{(t)}(i,j)+\vec{p}^{(t)}$$

These transformed points and their corresponding color information can be stored to construct a global model of the scene or object.

2.4 Surface Fitting

In order to create a surface for the scanned object, the range measurements are combined to generate a triangular patch mesh surface. The surface is represented in object coordinates, and allows the arbitrary repositioning of the shape for rendering. The surface is also needed to provide a place to attach color information.

Triangulated meshes can be generated from multiple 3D datasets in several ways. Each set can be first triangulated using, e.g. Delaunay triangulation. Then, these surfaces can be integrated through iterative averaging techniques [22], merging camera canonical sets based on Venn diagrams [24], or by zippering together the meshes along overlapping positions [25]. The alternative is to register all of the point sets all together. Such a point set can then be fitted using an inflationary balloon optimization process [4], or by triangulating the level set of an implicit surface through the points [15]. The latter approach was used here.

One key issue in the surface fitting is the large amount of noise in the 3D data - because of the inaccuracies in the 3D global position sensor orientation readings (as compared to 1-2 millimeters in the translation estimates) and about 1 mm from the range estimation process. The global position sensor produces orientation estimates that vary by about 3 degrees, so this means that estimated surface positions can vary by about 1 cm at a 20 cm scanning distance. This entails using techniques that find a global surface representation that passes through a cloud of range measurements. While this means surface metrical accuracy is not good enough for reverse engineering applications, this is not a serious problem for most VR applications.

To get a good surface estimate from the Hoppe [15] algorithm, we needed to reduce the measurement error. To do this, the algorithm clusters local groups of points to form better surface point estimates, which also removes outliers:

```
// make initial ball center list
ball_center_list = []
for each point P,
    if within distance D of any ball center
        then ignore point
        else append P to the ball_center_list

// assign points to balls
for each point P,
    assign to ball in ball_center_list with closest center

// delete outlier balls
delete all balls with less than T range points

// form output points
average points in each ball
```

The averaged output points are used as the inputs for the Hoppe algorithm. The pre-processing algorithm removes much of the noise from the surface, as well as the outliers, and reduces the number of sample points to triangulate. The pre-processing and Hoppe algorithms have the side-effect of rounding fold edges on objects, especially if the triangulation chosen for the Hoppe algorithm is coarse. On the other hand, if the triangulation is fine, then the surface fit will be affected by the considerable amount of noise here, causing many wild patch orientations. We also found it necessary to spend about 5 minutes of hand editing (using xgobi) the output of the preprocessing to remove outlying points before input into the Hoppe algorithm.

2.5 Color Measurement

Acquiring the color of a surface is a complicated matter, in that the observed image measurements are a function of the illumination onto the surface (including both the illumination direction and spectral distribution). Surface orientation and type affects how much light is reflected, as well as the color that we see. The image measurements are also a function of the spectral sensitivities of the sensor elements and the gamma (intensity compression) function of the camera.

The problem of estimating the true reflectance of a surface from its appearance is well known to computer vision researchers. Early approaches follow that described by Horn ([14] Chapter 9), wherein the illumination gradients in the three color channels are removed independently, and then are normalized by assuming that the brightest patch is white. This technique assumes that the observed scene is planar, but can be extended to polyhedral worlds [23]. Recent methods have depended on making special assumptions such as observing surface points in or near specularities [12], observing metals or dia-electrics [11], or having spectral measurements across the full visible spectrum [13]. More general techniques depend on estimating the illumination [18, 10]. If polychromatic light sources are collimated or surface geometry and spectral properties are separable, then these approaches are believed to be valid [17]. [6] shows that the main general approaches can all be reduced to an independent rescaling of three color channels (though the three channels may be a linear function of the standard RGB measurements). The above research is mainly concerned with surfaces radiating as a result of direct lighting. If observed color is a function of mutual reflection between nearby surfaces, then work by [9] shows how the reflectance can be extracted.

Alternatively, it is possible to obtain surface reflectance from a laser-based range finder, by analyzing the return beam signal strength [20]. This provides reflectance at only the laser frequency. Multiple spectral measurements can be made at the same surface point using a polychromatic laser [2].

One difficulty with all of these methods is their dependence on the assumption of Lambertian surface models, whereas real surfaces do not conform very well to the model. The deviation from the model becomes more apparent as surfaces become more slanted away from the observer; that is, as the deviation from the Lambertian assumption becomes more significant. Further, shading, lightness and color analysis usually depend on point light source assumptions, which are awkward, but not impossible, to arrange without simultaneously causing unwanted shadows. Finally, camera operating parameters, such as gain, gamma correction and white balance all conspire to prevent obtaining consistent measurements of input intensity. (We have experimented with calibrating the camera parameters with the AGC disabled and using controlled illumination, but have found that reflectance estimates deteriorate significantly as surfaces slant away from the observer – defeated by the Lambertian assumption).

In this project, we acquire and texture map color rather than reflectance. This approach is not as problematic as might be expected because:

- Ambient light enhanced by the light source on the hand-held unit is used, with the scanner always about the same distance from the scanned object, so that source illumination is nearly constant.

- The light source is near the camera optical axis, so that shadows are not easily seen by the camera,.

- The camera AGC works well enough to maintain a reasonably constant, unsaturated color signal.
- Only surfaces whose normals point to within 45 degrees of the camera optical axis are texture mapped. At these angles, the amount of reflected light is very nearly consistent for a wide range of non-specular surfaces.

In this section, we assume that the sensor is at a given position and all measurements relative to sensor coordinates are converted into measurements in the global coordinate system using the transformation provided by the global position sensor.

The method of ensuring registered range and color images in our scanner is simple - the same camera is used to acquire both the range and color information. A standard color TV camera is used and the light pattern used for calculating the 3D position is first observed via the red color channel. Color is acquired and "painted onto the surface" in a second pass over the object, after the surface shape is found. The key concept is that in the second scan each suitable pixel in the color image has a known line of sight into the scene. These lines of sight can be intersected with the now known object surface, because the global position sensor provides a common coordinate system in which to integrate the information (as the object being scanned will not have moved between scans). Where intersection occurs, the color signal can be attached to the surface shape description. Note that not all lines of sight from the intensity image will intersect the object surface; that is, some pixels are observed from the background. These pixels are detectable as their lines-of-sight do not intersect the object, and so can be ignored.

A single pass could be used to acquire all of the data, but for our prototype we use two passes.

To attach the color, the following quantities are needed:

Variable	Name	Notes
image plane pixel	(i, j)	the index of the current image pixel
camera center	$\vec{c}$	the current camera optical center
viewing directions	$\vec{v}(i, j)$	the line of sight from the camera center through pixel (i, j)
measured 3D surface position	$\vec{x}(i, j)$	3D position of the surface point that images onto pixel (i, j)
measured 3D surface normal	$\vec{n}_{\vec{x}}(i, j)$	3D surface normal of the surface point that images onto pixel (i, j)
light center	$\vec{l}$	the current light source center
light direction	$\vec{b}$	the illumination direction relative to $\vec{l}$

We know $\vec{c}$ and $\vec{l}$ in global coordinates through geometric calibration or measurement, $\vec{v}(i, j)$ in global coordinates through camera calibration.

The current viewing ray $\vec{c}+\mu\vec{v}(i, j)$ through pixel (i, j) is intersected with the previously found surface. This provides the $\vec{x}(i, j)$ measurement and knowledge of the surface function at $\vec{x}(i, j)$ allows calculation of $\vec{n}_{\vec{x}}(i, j)$.

2.6 Attaching Color to the Surfaces

Each video image provides a set of image pixels that are stored and post-processed to add the color to the triangulated range surface patches. Color is texture mapped onto the triangular surface patches by:

- The working region near the optical axis of the intensity image is extracted.
- Working in a raster fashion, the intensity image lines of sight $\vec{c} + \mu\vec{v}(i,j)$ are intersected with the triangulated 3D surface to obtain the 3D point position $\vec{x}(i,j)$ and patch normal $\vec{n}_{\vec{x}}(i,j)$. Pixels that do not intersect the object surface are ignored.
- Pixels on patches whose normals are more than 45 degrees away from the line of sight $\vec{v}(i,j)$ are ignored, because a) this prevents blurring the color measurement along tangent lines at the edges of the objects and b) shading effects are less significant when the surface normals point towards the observer.
- Pixels must lie within 15 degrees of the optical axis to avoid significant optical attenuation ([14], page 208).
- Unsaturated RGB color values are texture mapped onto the surface patches.
- Multiple color estimates at the same point are averaged.
- When RGB texture is added, a 2D Gaussian blurring convolution (1-2-1/2-4-2/1-2-1) was applied to reduce noise.
- Color values were propagated into uncolored pixels to fill in cracks between patches.

The resolution of the texture mapping cells was $1mm^2$.

3 Sensor Hardware

The theory described above did not consider actual hardware of such a sensor. Our implementation uses these devices:

- **Global Position Sensor**: The "A Flock of Birds" sensor model 6DFOB [8] provides global position XYZ and orientation measurements. This sensor reports the position of small movable remote unit relative to a fixed base unit. This model has a working volume of about $1m^3$ and has accuracy of about 1-2 mm and 3 degrees.
- **structured light pattern**: a laser stripe projector with a non-Gaussian stripe distribution [16] is used. It produces a 635 nanometer wavelength red laser, 5 mW, 15 cm long, 0.1 mm thick single stripe 12 cm long at 20 cm distance. The intensity profile is nearly linear (i.e. non-Gaussian) along the stripe.

- **color camera**: an off-the-shelf CCD color camera (model CD1 manufactured by Panasonic) is used. A 7.5 mm focal length lens is used.
- **broadband light source**: a 20 watt halogen bulb attached to a variable tap transformer to adjust intensity is used.
- **PAL-RGB converter**: model TYPE-PD-84 manufactured by Electrocraft is used to convert the PAL video signal output from the color camera into the three R/G/B signals input into the digitizer.
- **digitizer/frame grabber**: model DASM-VIP8 manufactured by the Analogic Corp. was used.
- **display rendering**: the Renderware real-time 3D graphics library made by Criterion software was used, for the surface display and color texture mapping. Color dithering was used to overcome limited 8-bit color space limitations.

4 Example Results

Figure 2 shows one of the test objects used here. Figure 3 shows the collection of range stripes obtained from the surface shape measurement phase and the output of the surface fitting algorithm is shown in Figure 4. Figure 5 shows two views of the cow with the RBG texture mapped onto the surface. While this paper can only display the greyscale values, the reconstructed color of the model is quite close to the toy's indoor appearance.

Figure 6 shows an intensity image of a cluttered table top. Figure 7 shows the RGB data texture mapped onto the recovered nearly-planar surface, with protrusions for the objects lying above the surface.

The stripe data for each model took about 3 minutes to acquire on a Sun Sparc 10/51 workstation, and this allowed the acquisition of about 50,000 3D surface measurements. The color data was taken immediately after this, and this took about 3 minutes as well. Though fewer images were taken (on the order of 10-20), these images were taken one at a time (rather than out of a video stream), and it took a few seconds to position the scanner before each picture.

The offline processing took: a) 10-60 minutes for the preprocessing, depending on the number of data points and parameters, to reduce the number of data points used in surface fitting to 1000-3000. Then, 1-10 minutes of hand editting was needed to remove outlying points. The Hoppe algorithm takes about 1 minute on this point set size, but 5-10 runs were needed to obtain a satisfactory result (issues are: surface coverage without holes, smooth surfaces, no breakthrough from internal noise points and a reasonable triangulation scale). Finally, color was texture mapped at about 1 intensity image per minute.

Figure 2: Cow toy used for experiment.

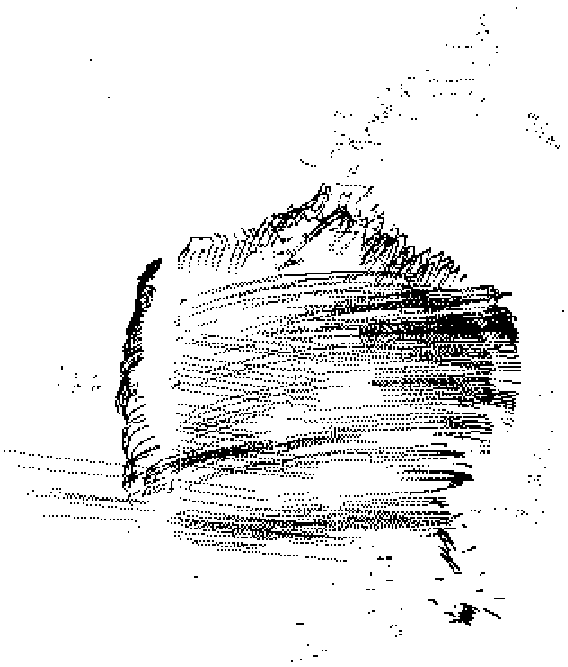

Figure 3: Stripe data covering cow toy.

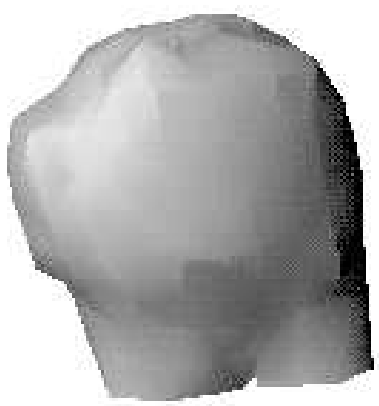

Figure 4: Untextured view of recovered cow toy surface shape.

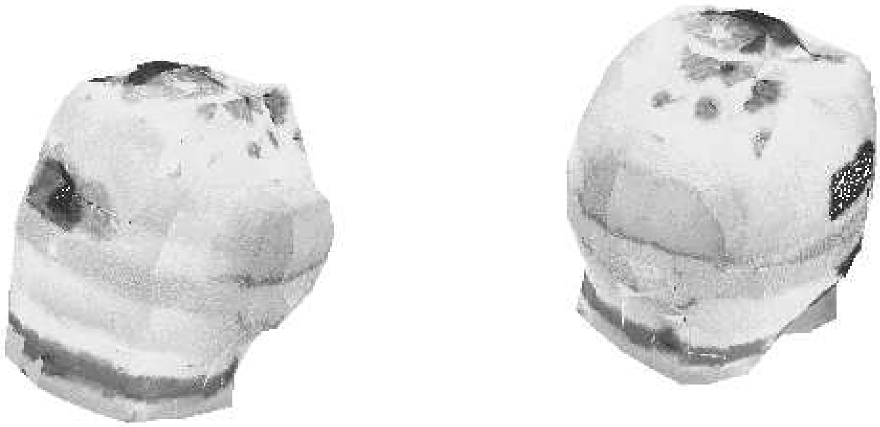

Figure 5: Two views of the texture-mapped cow model.

Figure 6: Intensity image of cluttered table top.

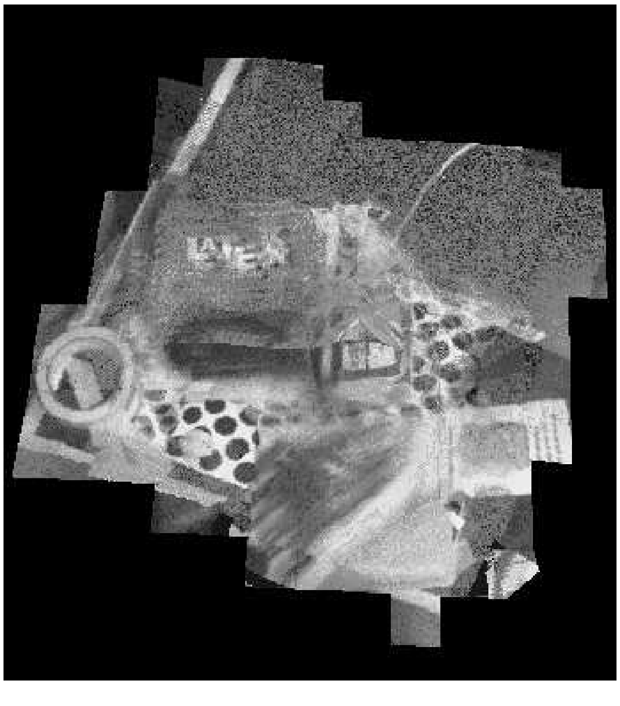

Figure 7: Texture-mapped surface model of cluttered table top.

5 Discussion and Conclusions

Adding a global position sensor to a range sensor clearly allows the range sensor to be portable and to be able to integrate surface shape measurements into a common global reference frame. Through the use of AGC and sensor controlled illumination, although true reflectance was not obtained, color data sufficient to reproduce the appearance of the test objects was obtained.

While the acquired range data are rather noisey and have a moderate amount of outliers (arising mainly from the noisey orientation measurements), the shape resulting from the Hoppe algorithm is quite good. The color information obtained is also reasonable in color, but there are problems with accurate registration of the color onto surface patches that slant away from the camera position (again, we believe, a consequence of errors in sensor orientation coupled with the surface slant).

The real-time portion of the sensor is reasonably fast, mobile within 0.5 meter of the base unit and is easily hand-held. The equipment that we have assembled allows us to produce a complete 3D colored surface model of an irregular natural object in a reasonable time, which offers a considerable improvement over existing methods. While texture mapping of intensity data to surfaces derived range images is not new, here we have shown how this can be made into a practical device. All these make this approach attractive for Virtual Reality applications, because of the complexity of real world scenes makes it impossible to reproduce them using hand-built models.

6 Future Work

There are a number of areas where improvements can be made:

- Better surfacing: because the data is particularly noisy, an algorithm that fits a constrained smooth surface through a cloud of points would have better performance than the algorithm used, which tries to stay close to the 3D data. For now, noisy surface data means that object edges are likely to be rounded.

- Because of variation in the global position sensor's orientation estimates, estimated surface positions vary about 5% in the direction perpendicular to the optical axis of the camera. Improved orientation estimates would reduce the 3D surface measurement errors, which would improve the surface fitting. The orientations could be improved (at some expense) by mounting three global positioning sensors on the scanner and computing orientation from the more accurate translation estimates and/or averaging the orientation estimates.

- The blue channel of the camera used is considerably less sensitive, and produces noisier image data.

- A brighter laser would allow the simultaneous acquisition of the color and shape information, although the color would have to be saved until a surface was estimated.

Alternatively, one might be able to estimate color several pixels to each side of the laser stripe, and interpolate these onto the stripe position. At the moment, the laser stripe is rather dim and the sensor illumination swamps the laser. On the other hand, the current laser is eyesafe and increasing the power would mean that eye protection would be needed when scanning.

- An improved surface display should be created as the sensor sweeps over a surface, to give the user better information about what portions of the scene still need to be covered.
- The current sensor is designed to work only at close range with small objects. For larger applications, such as scanning building interiors and exteriors, some straightforward re-engineering is needed: use of a global position sensor with a larger working volume (currently $1m^3$), a longer laser stripe better matched to the lens field of view and a higher resolution camera.
- Images that contain only a limited range of reflectances might cause a mis-match of colors across patches (slightly noticeable in the examples shown above). Improved color estimation is useful.
- The texture mis-registration due to position sensor orientation errors is noticeable, so some averaging or blending would be useful.
- Currently, best results are obtained with one texture view per patch; better averaging of multiple views (perhaps hundreds of views) might produce more reliable models.
- Variable surface patch size would be useful for providing large patches where little shape change is occurring, and small patches where there is much image detail.

Acknowledgements

The work presented here was supported by a grant from the UK Engineering and Physical Science Research Council (GR/H86905).

References

[1] 3D Scanners Ltd., South Bank Technopark, 90 London Road, London SE1 6LN, UK.

[2] R. Baribeau, M. Rioux, and G. Godin, Color reflectance modeling using a polychromatic laser range sensor, IEEE. Trans. Pat. Anal. and Mach. Intel., 14, 1992, 263-269.

[3] P.J. Besl, Active, Optical Range Imaging Sensors, Machine Vision and Applications, Vol 1:2, 127-152, 1988.

[4] Y. Chen, G. Medioni, Surface description of complex objects from multiple range images, Proc. Int. Conf. Comp. Vis. and Pat. Recog., 153-158, 1994.

[5] Cyberware Inc., Rapid 3D Color Digitizer Model 3030 Product Information, 2110 Del Monte Avenue, Monterey, CA 93940, 1995.

[6] G.D. Finlayson, M.S. Drew, and B.V. Funt, Diagonal transforms suffice for color constancy, Proc 4th Int. Conf. on Comp. Vis., 164-171, 1993.

[7] R. B. Fisher, A. Fitzgibbon, A. Gionis, M. Wright, D. Eggert, A Hand-held Optical Surface Scanner for Environmental Modeling and Virtual Reality, UK patent application 9506484.6.

[8] Ascension Technology Corp. PO Box 527, Burlington VT 05402.

[9] B.V. Funt, M.S. Drew, and J. Ho, Color constancy from mutual reflections, Int. J. Comp. Vis. 6, 1991, 5-24.

[10] R. Gershon, A.D. Jepson, and J.K. Tsotsos, From [R, G, B] to surface reflectance: computing color constant descriptors in images, Proc. Int. Joint Conf on Artificial Intelligence, 755-758, 1987.

[11] G. Healey, A color reflectance model and its use for segmentation, Int. Conf. on Computer Vision, 460-466, 1988.

[12] G. Healey, T. Binford, Color Algorithms for a General Vision System, Proc. Int. Joint Conf on Artificial Intelligence, 759-762, 1987.

[13] J. Ho, B.V. Funt, and M.S. Drew, Separating a color signal into illumination and surface reflectance components: Theory and applications, IEEE. Trans. Pat. Anal. and Mach. Intel., 12, 1990, 966-977.

[14] B.K.P. Horn, Robot Vision, MIT Press, 1986.

[15] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, W. Stuetzle, Surface reconstruction from Unorganized points, SIGGRAPH 92, 71-78.

[16] Lasiris Inc. 3549 Ashby, Saint-Laurent, Québec, Canada H4R 2K3. Laser part SNF-501L-635S-5-30.

[17] H.C. Lee, E.J. Breneman, and C.P. Schulte, Modeling light reflection for computer color vision, IEEE. Trans. Pat. Anal. and Mach. Intel., 12, 1990, 402-409.

[18] L.T. Maloney and B.A. Wandell, Color constancy: A method for recovering surface spectral reflectance, Image Understanding, 215-224, 1989.

[19] D.K. Naidu, R.B. Fisher, A comparative analysis of algorithms for determining the position of a stripe to sub-pixel accuracy, British Machine Vision Conference, 217-225, 1991.

[20] Nitzen, D., Brain, A. E., Duda, R. O., The measurement and use of registered reflectance and range data in scene analysis, Proceedings IEEE, Vol 65, pp206-220, 1977.

[21] Perceptron Inc., LASAR Hardware Manual, 23855 Research Drive, Farmington Hills, MI 48335, 1993.

[22] M. Rutishauser, M. Stricker, M. Trobina, Merging range images of arbitrarily shaped objects, Proc. Int. Conf. Comp. Vis. and Pat. Recog., 573-580, 1994.

[23] P. Sinha and E. Adelson, Recovering reflectance and illumination in a world of painted polyhedra, Proc 4th Int. Conf. on Comp. Vis., 156-163, 1993.

[24] M. Soucy, D. Laurendeau, A General surface approach to the integration of a set of range views, IEEE. Trans. Pat. Anal. and Mach. Intel., vol. 17, 4, 344-358, 1995.

[25] G. Turk, M. Levoy, Zippered polygon meshes from range images, SIGGRAPH 94, 311-318.