LIQUID CRYSTAL POLARIZATION CAMERA

Lawrence B. Wolff
Todd A. Mancini
Computer Science Department
The Johns Hopkins University
Baltimore, Maryland 21218

Abstract

We present a fully automated system which unites CCD camera technology with liquid crystal technology to create a *polarization camera* capable of sensing the polarization of reflected light from objects at pixel resolution. As polarization affords a more general physical description of light than does intensity, it can therefore provide a richer set of descriptive physical constraints for the understanding of images. Recently it has been shown that polarization cues can be used to perform dielectric/metal material identification, specular and diffuse reflection component analysis, as well as complex image segmentations that would be immensely more complicated or even infeasible using intensity and color alone. Such analysis has so far been done with a linear polarizer mechanically rotated in front of a CCD camera. The full automation of resolving polarization components using liquid crystals not only affords an elegant application, but reduces the amount of optical distortion present in the wobbling of a mechanically rotating polarizer. In our system 2 twisted nematic liquid crystals are placed in front of a fixed polarizer placed in front of a CCD camera. The application of a series of electrical pulses to the liquid crystals in synchronization with the CCD camera video frame rate produces a controlled sequence of polarization component images that are stored and processed on Datacube boards. We present a scheme for mapping polarization states into hue, saturation and intensity which is a very convenient representation for a polarization image. Our polarization camera outputs such a color image which can then be used in polarization-based vision methods. The unique vision understanding capabilities of our polarization camera system are demonstrated with experimental results showing polarization-based dielectric/metal material classification, specular reflection and occluding contour segmentations in a fairly complex scene, and surface orientation constraints for object recognition.

1 Introduction

As human beings we naturally think of vision in terms of perception of intensity and color. Polarization of light might appear to be of little relevance or benefit to automated vision systems simply because the human visual system is almost completely oblivious to this property of light. In context of physics-based vision there is in fact a compelling motivation to study polarization vision - polarization affords a more general description of light than does intensity, and can therefore provide a richer set of descriptive physical constraints for the interpretation of an imaged scene. As intensity is the linear sum of polarization components, intensity images physically represent reduced polarization information. Because the study of polarization vision is more general than intensity vision there are polarization cues that can immensely simplify some important visual tasks (e.g., region and edge segmentation, material classification, etc...) which are more complicated or possibly infeasible when limited to using intensity and color information. A detailed description of a variety of polarization-based vision methods are contained in [11], [12], [14], [13], [3].

A criticism that has sometimes been leveled at polarization-based vision methods is the inconvenience of obtaining polarization component images by having to place a linear polarizing filter in front of an intensity CCD camera and mechanically rotating this filter by hand or by motor into different orientations. This inconvenience is simply a result of commercially available camera sensors being geared towards taking intensity images instead of polarization images. We feel that there are considerable advantages to building a camera sensor geared towards doing polarization vision, capable of taking polarization images without external mechanical manipulation of a filter. We call such a camera sensor a "polarization camera". There already exist polarization-based vision methods that can significantly benefit a number of application areas such as aerial reconnaissance, autonomous navigation, inspection, and, manufacturing and quality control. A polarization camera would make polarization-based vision methods more accessible to these application areas and others. It should be fully realized that as intensity is a compression of polarization component information, that a polarization camera can function as a conventional intensity camera, so that intensity vision methods can be implemented by such a camera either alone, or, together with polarization-based vision methods. As intensity-based methods are physical instances of polarization-based methods, a camera sensor geared towards polarization vision does not in any way exclude intensity vision, it only generalizes it providing more physical input to an automated vision system! Adding color sensing capability to a polarization camera makes it possible to sense the complete set of electromagnetic parameters of light incident on the camera. However in this paper we discuss the implementation of a monochrome polariza-

0-8186-2840-5/92 $03.00 © 1992 IEEE

tion camera.

We discuss and demonstrate in this paper particular implementations of 2 different aspects for the design of a polarization camera. The first is obviating the need for mechanically rotating a polarizing filter in front of an intensity camera sensor by using twisted nematic (TN) liquid crystals that electro-optically "steer" the plane of polarization of light. These kinds of liquid crystals are more typically used for producing graphical displays such as in laptop personal computers. In this paper we are proposing to apply the same physical principles of liquid crystal technology, but in a different way, to polarization-based computer vision applications. By electrically switching the states of TN liquid crystals in front of a fixed linear polarizer, measuring the components of polarization of reflected light from objects can be placed under full computer control. Whereas before the polarizer had to mechanically rotate relative to the plane of polarization of light to resolve components, the polarizer now remains fixed while instead it is the plane of polarization that is rotated by the liquid crystals. This also reduces optical distortion caused by a rotating polarizing filter particularly near high contrast intensity edges where small subpixel shifts can produce a false detection of high partial polarization. A special driver was built to make sure the switching of the liquid crystals is properly synchronized with video frames produced by the intensity camera.

The second aspect of a polarization camera that is addressed in this paper is the implementation of a natural visual representation for the polarization state of light at each pixel in a polarization image. How can humans be made to "see" reflected polarization, a phenomenon that nobody has ever directly observed except with the aid of special filters ? This representation for polarization images should of course be amenable to quick computational processing by existing polarization-based algorithms that extract visual constraints from polarization information. It turns out that practically all light that is reflected in most environments is *partially linearly polarized*, meaning that it can be represented by the sum of unpolarized and completely linearly polarized states. This is an important simplification as in the most general case the completely polarized component of polarization can be any state of elliptical polarization. We demonstrate a natural one-to-one mapping of a state of partially linear polarization into a hue, saturation (i.e., excitation purity), and, intensity, derived respectively from the orientation of the plane of the completely linear polarized component, the partial polarization (i.e., the percentage of complete linear polarization content), and the intensity of the light. A similar system for mapping polarization states into color was proposed by Bernard and Wehner [1] for studying biological polarization vision systems. Therefore, in a polarization image, unpolarized light appears achromatic and regions that are significantly partially polarized appear chromatically saturated. The intensity of light in a polarization image is simply the pixel intensity itself, regardless of color, and can be easily processed by intensity-based vision methods. This distinctly demonstrates how a polarization image is a generalization of a gray level intensity image. A number of polarization images taken with our liquid crystal polarization camera are shown.

Liquid crystals can provide a low cost convenient way of converting a standard intensity CCD camera into a fully automatic polarization camera. It is certainly possible to manufacture a compact optical head consisting of a linear polarizer and TN liquid crystals in series that can be mounted on the end of a lens. Electrical contacts would be provided on such an optical head for cables leading to the camera video sync and power supply.

While physics is our primary motivation for studying polarization vision and building polarization cameras, it is important to note that there are a large number of biological visual systems within a variety of insects [7], [8], and fish [4], [6], that primarily rely on polarization sensing. In fact not only is polarization vision in biological systems historically older than human vision, but in sheer number, there are more natural polarization vision systems existing today than there are humans.

2 Polarization Vision Methods: Background

The *polarization state* of light characterizes its complete description as an electromagnetic wave, apart from wavelength. The electric field oscillation for light in general can be represented by the superposition of 2 mutually orthogonal waveforms. When this superposition is deterministic, the time evolution of the tip of the electric field traces out an ellipse in the plane perpendicular to the light wave's direction of travel, and such light is said to be *completely polarized*. The shape of this ellipse is dependent upon the phase difference of the deterministic superposition of the 2 mutually orthogonal waveforms. Linear polarized light results from a degenerate ellipse when these waveforms are either at 0° or 180° relative phase. *Unpolarized light* results from the completely non-deterministic superposition of 2 equal amplitude mutually orthogonal waveforms. Most light emitted from incandescent light sources is of this type. The time evolution of the tip of the electric field for unpolarized light is a non-deterministic isotropic orientation distribution in the plane perpendicular to the light wave's direction of travel. For a formal mathematical description of polarization of light see Born and Wolf [2] or Clarke and Grainger [5].

Unpolarized light reflected off a material surface becomes *partially polarized* meaning that the polarization state can be represented as the sum of the unpolarized state and a completely polarized state. Unpolarized light reflected from objects is almost always *partially linearly polarized* meaning that the completely polarized component is linearly polarized. An exception to this occurs when there is multiple specular intereflection amongst metals producing somewhat of an elliptical completely polarized component. However the assumption of partially linearly polarized reflected light is very accurate almost all the time. The polarization state of partially linearly polarized light can be measured using only a linear polarizing filter. Figure 1 shows what happens to the isotropic electric field orientation distribution of an unpolarized light wave passing through such a filter. Only the component of the electric field parallel to the transmission axis is transmitted. It turns out

that the radiance of unpolarized light passing through a polarizing filter is constantly one-half the original light radiance regardless of the orientation of the polarizing filter. The transmitted radiance for a partially linearly polarized light wave in general is sinusoidal as a function of polarizer orientation oscillating between a minimum and a maximum every 90°. We will term this the *transmitted radiance sinusoid.* Figure 2 depicts a transmitted radiance sinusoid with minimum occurring exactly at the zero reference. Referring to Figure 2 it can be shown [5] that the difference $I_{max} - I_{min}$ represents the magnitude of the linearly polarized component of the original light, while the sum $I_{max} + I_{min}$ represents the total radiant intensity of the original light. The ratio $\frac{I_{max}-I_{min}}{I_{max}+I_{min}}$ varying between 0 and 1 inclusive therefore represents the fraction of light that is linearly polarized, and is called the *partial polarization.* At 0 partial polarization (i.e., unpolarized light) the transmitted radiance sinusoid is a flat horizontal line and the partial polarization is zero. For linearly polarized light, $I_{min} = 0$ (occurring when the transmission axis of the polarizer is perpendicular to the orientation of the linearly polarized light) and therefore the partial polarization is 1.

FIGURE 1

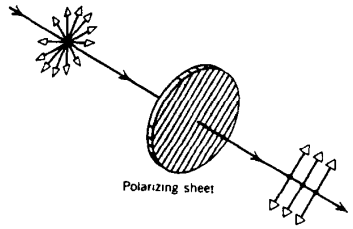

FIGURE 2

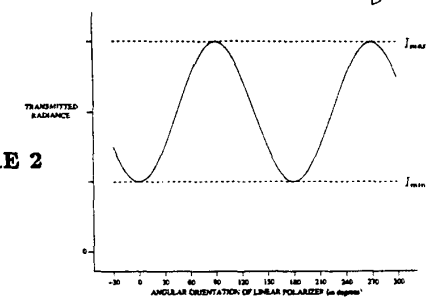

The polarization state of partially linearly polarized light can be characterized uniquely by the transmitted radiance sinusoid. The transmitted radiance sinusoid can be completely described by the parameters, I_{min}, I_{max}, and, the *phase*, θ, of the sinusoid with respect to some reference zero. Alternatively, and more directly related to visual features extracted from polarization-based methods, the transmitted radiance sinusoid can be completely described by the parameters:

$$\begin{aligned} &\text{(partial polarization) } \frac{I_{max} - I_{min}}{I_{max} + I_{min}}, \\ &\text{(total intensity) } I_{max} + I_{min}, \\ &\text{(phase) } \theta\,. \end{aligned} \tag{1}$$

A very simple way of implementing the computation of the polarization state of light incident upon each pixel in an image is to place a polarizing filter in front of an intensity camera as in Figure 3 and derive the transmitted radiance sinusoid for each pixel. As 3 points determine a sinusoid, at least 3 images are required to be taken respective to 3 unique orientations of the polarizing filter within a 180° range. Using more than 3 images overconstrains the determination of the transmitted radiance sinusoid and a nonlinear optimization technique such as Levenberg-Marquadt is quite easy to apply. We have found that using 3 images alone gives good accuracy in measuring reflected partially linear polarization states. We have had good success deriving the transmitted radiance sinusoid by obtaining polarization component images at polarizer orientations 0°, 45°, and, 90° on the polarizer ring vernier without even knowing what these angles are relative to the transmission axis of the polarizer (i.e., all we know is that these angles are 45° apart). The image irradiances obtained at each pixel are I_0, I_{45}, I_{90} respective to each of the relative polarizer orientations. If θ represents where I_{min} occurs relative to 0° on the polarizer ring vernier, then the derivation of the 3 parameter expressions in equation 1 are:

$$\theta = (1/2)\tan^{-1}\left(\frac{I_0 + I_{90} - 2I_{45}}{I_{90} - I_0}\right),$$

$$\text{if } (I_{90} < I_0) \text{ [if } (I_{45} < I_0)\ \theta = \theta + 90 \text{ else } \theta = \theta - 90]$$

$$I_{max} + I_{min} = I_0 + I_{90}, \tag{2}$$

$$\frac{I_{max} - I_{min}}{I_{max} + I_{min}} = \frac{I_{90} - I_0}{(I_{90} + I_0)\cos 2\theta}\,.$$

A good deal has been published about how important physical constraints relevant to image understanding can be extracted from these measured polarization parameters [11], [12], [14], [13], [3]. We will briefly describe the principles involved in polarization-based vision methods that are utilized by our polarization camera, and reference the more quantitative details in these publications.

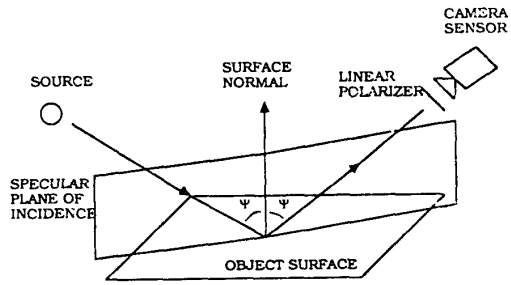

FIGURE 3

Figure 3 depicts the *specular plane of incidence*, also known simply as the *specular plane*, determined by the incident light direction and the viewing direction of the camera sensor. This is the plane in which specular reflection occurs into the camera sensor. Unpolarized light

can be represented as the non-deterministic superposition of a waveform parallel to the specular plane, with an equal amplitude waveform perpendicular to the specular plane. Upon specular reflection both waveforms are attenuated according to the Fresnel reflection coefficients, [9], [11], [14]. The parallel waveform is attenuated more than the perpendicular waveform so that specular reflection becomes partially linearly polarized with a larger waveform component perpendicular to the specular plane than parallel to it as shown in Figure 4. Hence, the polarization state of specularly reflected light is the sum of an unpolarized state with a completely linearly polarized state whose plane is perpendicular to the specular plane. The resulting measured transmitted radiance sinusoid should exhibit a minimum when the polarizer is oriented parallel to the specular plane, and exhibit a maximum when the polarizer is oriented perpendicular to the specular plane.

The polarization state of diffuse reflection is most of the time unpolarized. However, near occluding contours where surface orientation normals are almost orthogonal to viewing, diffuse reflection from inhomogeneous dielectric surfaces (e.g., plastic, ceramic, rubber, etc...) can become significantly partially linearly polarized. Practically all of diffuse reflection arising from inhomogeneous dielectric surfaces results from penetration of light into the surface, multiple subsurface scattering, and then refraction back out into air. Figure 5 depicts the change in polarization state of diffuse reflected light as it is refracted from within the dielectric into air. Before refracting out into air the polarization state of light is assumed to be unpolarized. For most angles of refraction with respect to the surface normal, the refracted light making up diffuse reflection remains essentially unpolarized. However, above 60^o between the surface normal and the viewing direction, refraction produces a significant partial linear polarization. At these high angles the electric waveform perpendicular to the *emittance plane*, determined by the surface normal and the viewing direction, is much more attenuated than the electric waveform parallel to the emittance plane. The polarization state of diffuse reflection near occluding contours is the sum of an unpolarized state with a completely linearly polarized state whose plane is parallel to the emittance plane. The resulting measured transmitted radiance sinusoid should exhibit a minimum when the polarizer is oriented perpendicular to the emittance plane, and exhibit a maximum when the polarizer is oriented parallel to the emittance plane.

In summary, significant partial polarization (i.e., above 10%) in a scene can be due to specular reflection and/or diffuse reflection from inhomogeneous dielectric objects near occluding contours. For emittance and specular planes coinciding near an occluding contour, the transmitted radiance sinusoids for the specular and diffuse reflection components are respectively 90^o out of phase. This is an important physical principle that can be exploited to help distinguish between partial polarization due to specular reflection and diffuse reflection. While most of the time the specular and emittance planes are unknown, additional polarization principles can be exploited to identify specularity, and, occluding contour regions and edges in a scene, [13], [14]. Adding heuristics about physical size can help physical edge labeling [3]. This reveals an immense amount of physical information about a scene that would normally be difficult, or sometimes infeasible, to obtain from intensity and color information alone.

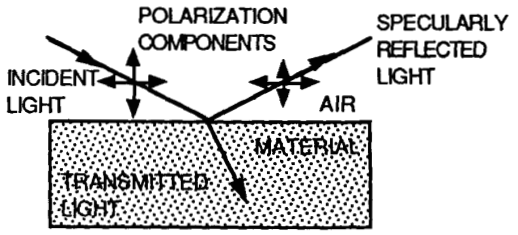

FIGURE 4

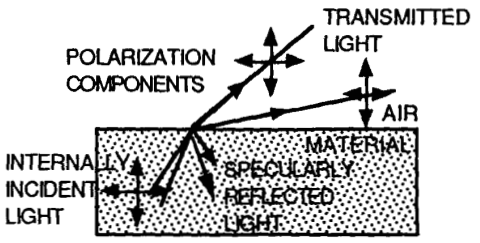

FIGURE 5

On smooth and mildly rough surfaces the phase of the transmitted radiance sinusoid gives surface normal constraint information [11], [13], [14]. Observing Figure 3 the surface normal is constrained to lie in the specular plane. The pattern of transmitted radiance sinusoid phases from specular reflection occurring at multiple surface orientations on an object gives physical shape cues that can be exploited in object recognition.

Another important mode of physical information for interpreting objects in a scene is identification of intrinsic material classification. The capability of determining whether parts of an object are metal (conductor) or dielectric (nonconductor) can be very useful to object recognition and material inspection in manufacturing (e.g., circuit board inspection, package inspection, etc...). Material classification is a difficult problem using intensity and color, but there exist polarization cues that immensely simplifies the problem of determining relative electrical conductivity of materials, with metal and dielectric at the extremes. The theory of this is explained in [12], [13]. It turns out that if the specular angle of incidence is between 30^o and 80^o, and the specular component of reflection is strong relative to the diffuse component, the quantity:

$$\frac{I_{max}}{I_{min}}, \tag{3}$$

derived from transmitted radiance sinusoid parameters,

is a very reliable discriminator for varying levels of electrical conductivity. This ratio for most metals varies between 1.0 and 2.0 while for dielectrics this ratio is above 3.0.

The section on Visualizing Polarization And Experimental Results shows how our fully automatic liquid crystal polarisation camera, using our new visualization scheme for polarization, can be used to implement these types of applications.

3 Theory of Operation of a Liquid Crystal Polarization Camera

Obtaining the transmitted radiance sinusoid by rotating a polarising filter in front of a CCD camera is a mechanically active process that produces optical distortion and is difficult to fully automate. Unless the axis perpendicular to the polarizing filter is exactly aligned with the optic axis of the camera, small shifts in projection onto the image plane occur between different orientations of the polarizing filter. At intensity discontiniuites in a scene, significant shifts in image intensity are observed giving the false interpretation of reflected partial polarization even if it does not exist. Fully automating the mechanical rotation of a polarizing filter would require a motor that would have to precisely rotate the filter in synchronization with video frame rates.

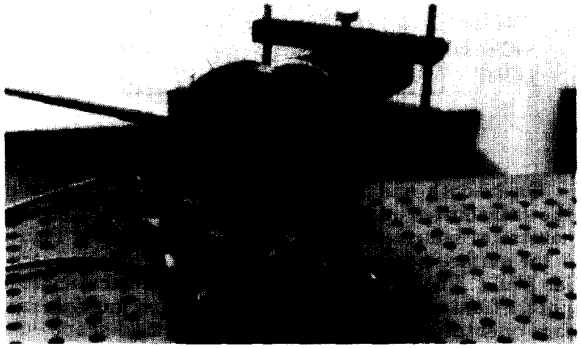

FIGURE 6

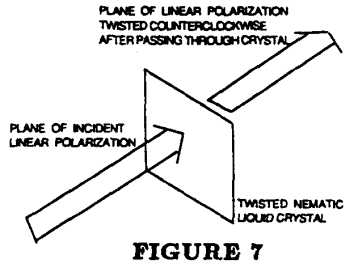

FIGURE 7

Figure 6 shows our liquid crystal polarization camera using a CCD camera with a fixed polarizer and two twisted nematic liquid crystals mounted in front. The idea behind a liquid crystal polarization camera is very simple which is why we feel it works so well. Nothing mechanically rotates; the polarizer remains fixed while the twisted nematic (TN) liquid crystals electro-optically rotate the plane of the linear polarized component of reflected partially linear polarized light. The unpolarized component is not effected. In general the transmitted radiance sinusoid can be recovered by the *relative* rotation of the plane of linear polarization with respect to the polarizer. Each TN liquid crystal is binary in the sense that it either rotates the plane of linear polarization by fixed n degrees, $0^o < n \leq 90^o$, which is determined upon fabrication, and, 0 degrees (i.e., no twist). We use 2 TN liquid crystals, one at $n = 45^o$, and the other at $n = 90^o$, to insure at least 3 samplings of the transmitted radiance sinusoid.

Liquid crystals come in different varieties and some of the theory behind them can be quite involved [10]. The molecular structure of the material in *twisted nematic* liquid crystals is helical, twisting slowly from one face of the crystal to the other face by a predesigned fixed amount n degrees. With no voltage applied across the liquid crystal faces, the plane of linear polarized light rotates along the helix by n degrees. See Figure 7. When an AC voltage is applied across the liquid crystal faces, the helices straighten out so that the plane of linear polarized light is not rotated in this state. The switching or "relaxation" time of twisted nematic crystals is slow compared to other types, on the order of 1/10 of a second, even though faster switching liquid crystals are being developed as time progresses. The switching of liquid crystal states leaves the geometry of the optical projection of the world scene onto the image plane virtually unchanged.

The driver for our polarization camera modulates a high frequency AC voltage so as to produce 4 states between the 2 TN liquid crystals, 0^o, 45^o, 90^o, and, 135^o. Each state lasts 4 frame times (1/7.5 seconds) in which the liquid crystals are allowed to switch and an image is grabbed on the last of these frames. We utilize only the first 3 liquid crystal states, while doing processing on our Datacube boards the last 4 frame times. Grabbing 3 polarization component images together with implementing lookup tables on the Datacube boards that perform the derivations of expressions 2, to produce a polarization image takes 16/30 of a second. At present our polarization camera therefore operates at slightly under 2Hz.

4 Visualizing Polarization and Experimental Results

A very important component of a polarization camera is its output. Instead of just outputting a set of intensity images corresponding to different components of polarization, a polarization camera should produce an image derived from these component images which naturally represents the physical characteristics of polarization. As human vision for the most part is oblivious to the phenomenon of polarization we are limited to image representation in terms of intensity and color. We present a way in which to map polarization information into a visual color space that not only makes it obvious what physical characteristics of polarization are being represented, but also makes it clear how these characteristics relate to important physical properties of the scene being imaged.

Consider the 3 polarization parameters of expressions 2 characterizing the transmitted radiance sinusoid.

The quantity $I_{max} + I_{min}$ is directly observed as the total intensity so it is natural to represent this as an image intensity. The other 2 parameters, phase of the transmitted radiance sinusoid, and partial polarization, are not directly observed by human vision, so these we propose to be mapped into the color domain. It is quite interesting that the phase of the transmitted radiance sinusoid is an angular quantity with range $0 - 180$ degrees, while the partial polarization is a fractional quantity with range from 0 to 1. This is exactly analogous with hue and saturation, respectively, except that the phase angle of the transmitted radiance sinusoid should be multiplied by 2 to cover the full 360^o range of hue. See Figure 8 for the case when $\theta = 0^o$ is represented by green. We therefore propose to map the parameters of expressions 2 respectively into hue, intensity, and saturation, of HSV color space, as just described, to represent a polarization image. Regions of a polarization image with 0 partial polarization will have no chromatic saturation and will appear as ordinary gray level intensity. Polarization with the same phase will have the same hue but possibly different saturation content according to how much it is partially polarized. The largest phase difference between two transmitted radiance sinusoids is 90^o and such polarization states will appear in a polarization image as complementary colors.

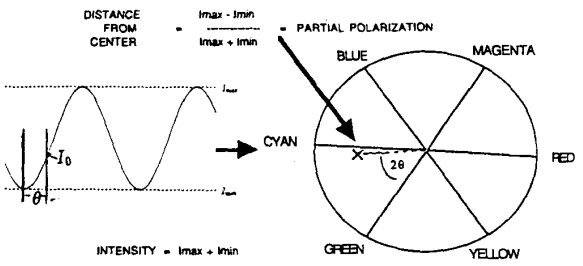

FIGURE 8

Unfortunately the color polarization images described below could only be reproduced in this printing in black-and-white. Full color reproductions are available upon written request.

Figure 9 shows how a polarization image provides important information about a scene that would be very difficult and perhaps impossible to deduce from an intensity image. The left intensity image of Figure 9 shows what apparently are 2 mugs in a scene. Looking closely at the intensity image reveals that there is some difference between the 2 mugs; the left mug has its letters reversed. The only visual cues telling that the left mug is simply a reflection are very high level features such as the reversal of recognizable high level features (e.g., alphabet letters) or the edge of the glass mirror. Otherwise the reflected intensity (and color) of the 2 mugs look essentially the same. This type of problem occurs in vision fairly frequently such as when stray specular glare from objects give the false interpretation that real edges actually exist there. Consider the problem of an autonomous land vehicle viewing a scene part of which is reflected by a lake or river. How does the vehicle know which are the "real" elements of the scene ? How does a mobile robot know when it is running into a glass door, or if navigating according to edge cues, which are geometric edge cues opposed to specular edge cues ? The right polarization image in Figure 9 was obtained with our liquid crystal polarization camera showing that the left mug has Cyan chromaticity implying significant partial polarization. Cyan chromaticity is also observed at specular highlights on the right mug as well. (The very bright center of specularities saturate the camera so that pixels record gray level 255 regardless of the state of the TN liquid crystals. This gives a flat transmitted radiance sinusoid, and hence, the appearance of unpolarized light, when in fact the reflected light from these areas are significantly partially polarized. This is a limitation of the dynamic range of the SONY XC-77 CCD camera being used, and NOT our polarization vision algorithm.) Significant partial polarization is also observed at the occluding contour of the right mug as Red color. Note that the hue colors Cyan and Red are complementary colors indicative of transmitted radiance sinusoids 90^o out of phase.

When taking the polarization image in Figure 9, the fixed polarizer analyzer on our liquid crystal polarization camera is oriented so that the transmission axis is horizontal. The hue Cyan represents when the plane of the linear polarized component of reflected light is vertical with respect to the image (i.e., when the transmitted radiance sinusoid is observed minimum when both TN liquid crystals are in the zero twist "ON" state). Using the physical principles discussed in the Background section, Cyan hue corresponds to specular reflection when the specular plane is horizontal relative to the image. As it turns out the specular planes in this image are all approximately horizontal. The Red color hue representing a 90^o transmitted radiance sinusoid phase difference relative to the Cyan color hue is indicative of partial polarization from diffuse reflection near an occluding contour. Note the thin edge of specular reflection occurring at the very edge of the occluding contour on the right mug. Even when the specular planes are not known it is possible from polarization images to deduce which are the specular reflections and which are the occluding contours, from a number of physical principles [13], [14]. Polarization vision makes segmentation of these types of regions and edges immensely easier than when using intensity.

Figure 10 shows the intensity and polarization images of a cylindrical cup illuminated with an extended light source so as to produce specular reflection from a number of different surface orientations. The different color hues shown in the polarization image correspond to specular plane surface orientation constraints. See Figure 3. In this example, Cyan color hue corresponds to specular planes oriented vertically in the image while the complementary color hue, Red, would correspond to specular planes oriented horizontal in the image. Almost the entire spectrum of color hues is displayed here. Figure 11 shows intensity and polarization images of one hemisphere of a plastic sphere illuminated with an extended light source. While the polarization image does not give completely unique surface orientation information, the pattern of specular plane constraints gives enough rudi-

mentary shape information to distinguish different shape classes for object recognition. For instance, on a cylindrical shape the lines of constant color hue are parallel to one another (Figure 10) while on a spherical shape lines of constant color hue mutually intersect at a point (Figure 11). Besides being useful in sorting by shape systems in manufacturing, outdoor objects illuminated by skylight serving as an extended illuminator may be able to be distinguished by shape class as well.

Figure 12 shows material classification on a circuit board according to various levels of electrical conductivity. On this circuit board there are 3 basic classes of materials, bare metal which appears both bright and dark, plastic dielectric which serves as the substrate of the board, and metal on top of which there is a translucent plastic dielectric coating producing a combined reflected polarization signature between that of metal and dielectric. The bottom "polarization image" produced by our polarization camera is not the standard one used previously. Instead each pixel represents the ratio of expression 3 from the Background section. The bright ratios represent dielectric, intermediate ratios represent metal coated with plastic dielectric, dark ratios represent bare metal. In this way we are directly visualizing material type instead of polarization state and this exhibits the diversity that is possible with a polarization camera. Any mathematical combination of the transmitted radiance sinusoid parameters can be represented in intensity and color. As these combinations of polarization parameters relate to physical properties of parts of a scene, so can the actual physical properties of the scene be visualized in an output image of a polarization camera.

5 Conclusion and Future Work

We presented the design and a particular implementation of a camera sensor that is geared to do polarization vision. This type of camera sensor, which we term a "polarization camera", subsumes the capabilities of existing intensity cameras in its ability to fully automatically resolve polarization components and output a polarization image which generalizes the information content of a standard intensity image. We exploited the principles of liquid crystal technology to fully automate the process of resolving polarization components, and, presented a visualization scheme for mapping states of partial linear polarization into HSV color space. Datacube boards digitise polarization component images and process these images to produce a polarization image. Experimentation with our liquid crystal polarization camera demonstrated some of its capabilities with respect to segmentation of specular reflection and occluding contours, obtaining shape constraints, and, material classification.

Liquid crystals are low cost and convenient to use, so they provide an elegant way of quickly converting a standard intensity CCD camera into a fully automatic polarization camera. The design of our liquid crystal polarization camera lays the groundwork for faster and more self-contained polarization camera designs. The more general capabilities of polarization vision motivates the building of polarization cameras. There already exist polarization based algorithms for extracting important visual information that transcend some key limitations encountered by intensity and color based vision algorithms. It is felt that building polarization cameras will aid in the development of even more effective polarization vision algorithms.

Acknowledgements

This research was supported in part by NSF grant IRI-91-11973 and DARPA contract F30602-92-C-0191. We would like to thank very much Prof. Andreas Andreou and Philippe Pouliquen of Electrical and Computer Engineering at Johns Hopkins for the time consuming task of building the driver for the twisted nematic liquid crystals. We also thank Pat Dunn and Phil Bos for their immense help in designing the liquid crystals needed for our application.

References

[1] G. D. Bernard and R. Wehner. Functional similarities between polarization vision and color vision. *Vision Res.*, 17:1019–1028, 1977.

[2] M. Born and E. Wolf. *Principles of Optics.* Pergamon Press, 1959.

[3] T.E. Boult and L.B. Wolff. Physically-based edge labeling. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Maui, June 1991.

[4] D.A. Cameron and E.N. Pugh. Double cones as a basis for a new type of polarization vision in vertebrates. *Nature*, 353:161–164, September 1991.

[5] D. Clarke and J.F. Grainger. *Polarized Light and Optical Measurement.* Pergamon Press, 1971.

[6] C.W. Hawryshyn. Polarization vision in fish. *American Scientist*, 80:164–175, March-April 1992.

[7] G.A. Mazokhin-Porshnyakov. *Insect Vision.* Plenum Press, New York, 1969.

[8] S. Rossel and R. Wehner. Polarization vision in bees. *Nature*, 323:128–131, September 1969.

[9] R. Siegal and J.R. Howell. *Thermal Radiation Heat Transfer.* McGraw-Hill, 1981.

[10] E.B. Priestly P.J. Wojtowicz and P. Sheng. *Introduction to Liquid Crystals.* Plenum Press, New York, 1975.

[11] L.B. Wolff. Surface orientation from polarization images. In *Proceedings of Optics, Illumination and Image Sensing for Machine Vision II, Volume 850*, pages 110–121, Cambridge, Massachusetts, November 1987. SPIE.

[12] L.B. Wolff. Polarization-based material classification from specular reflection. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 12(11):1059–1071, November 1990.

[13] L.B. Wolff. *Polarization Methods in Computer Vision.* PhD thesis, Columbia University, January 1991.

[14] L.B. Wolff and T.E. Boult. Constraining object features using a polarization reflectance model. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 13(7):635–657, July 1991.

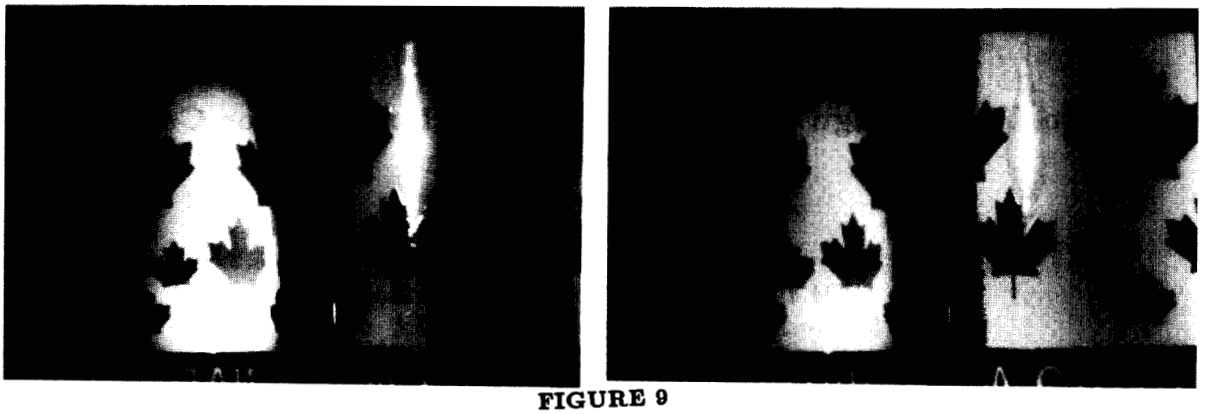

FIGURE 9

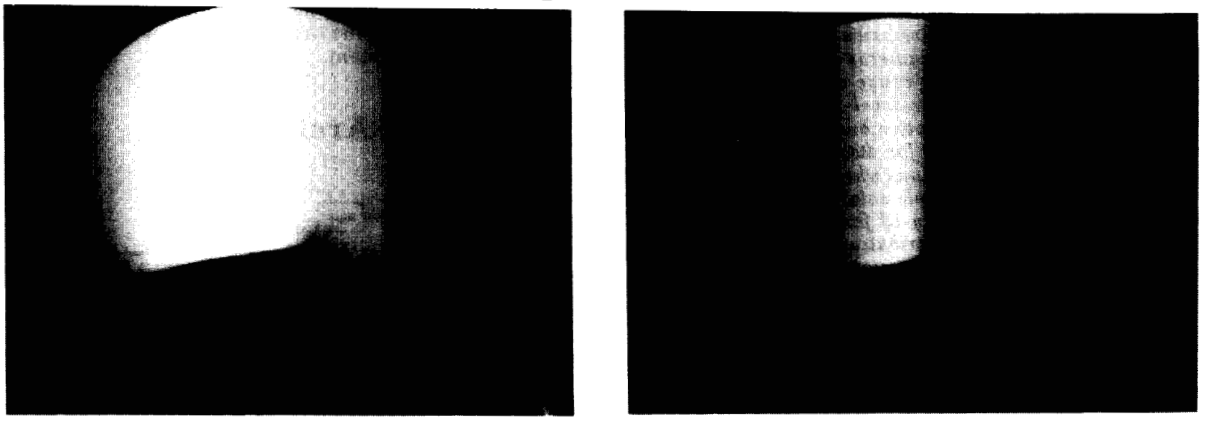

FIGURE 10

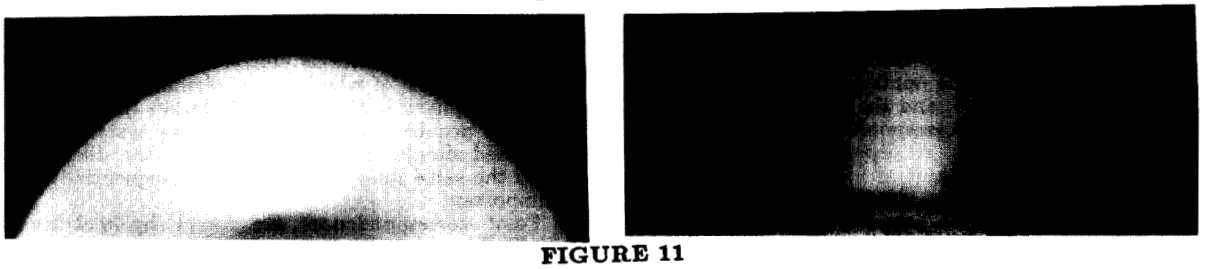

FIGURE 11

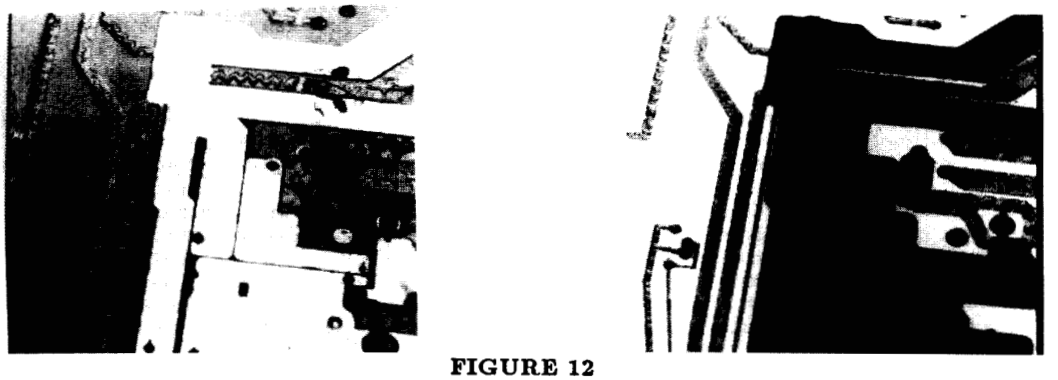

FIGURE 12