October 1992

A Proposal Concerning the Analysis of Shadows in Images by an Active Observer (Dissertation Proposal)

Gareth D. Funka-Lea
University of Pennsylvania

University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-92-78.

A Proposal Concerning the Analysis of Shadows in Images by an Active Observer (Dissertation Proposal)

Abstract

Shadows occur frequently in indoor scenes and outdoors on sunny days. Despite the information inherent in shadows about a scene's geometry and lighting conditions, relatively little work in image understanding has addressed the important problem of recognizing shadows. This is an even more serious failing when one considers the problems shadows pose for many visual techniques such as object recognition and shape from shading. Shadows are difficult to identify because they cannot be infallibly recognized until a scene's geometry and lighting are known. However, there are a number of cues which together strongly suggest the identification of a shadow. We present a list of these cues and methods which can be used by an *active* observer to detect shadows. By an active observer, we mean an observer that is not only mobile, but can extend a probe into its environment. The proposed approach should allow the extraction of shadows in real time. Furthermore, the identification of a shadow should improve with observing time. In order to be able to identify shadows without or prior to obtaining information about the arrangement of objects or information about the spectral properties of materials in the scene, we provide the observer with a probe with which to cast its own shadows. Any visible shadows cast by the probe can be easily identified because they will be new to the scene. These actively obtained shadows allow the observer to experimentally determine the number and location of light sources in the scene, to locate the cast shadows, and to gain information about the likely spectral changes due to shadows. We present a novel method for locating a light source and the surface on which a shadow is cast. It takes into account errors in imaging and image processing and, furthermore, it takes special advantage of the benefits of an active observer. The information gained from the probe is of particular importance in effectively using the various shadow cues. In the course of identifying shadows, we also present a new modification on an image segmentation algorithm. Our modification provides a general description of color images in terms of regions that is particularly amenable to the analysis of shadows.

Comments

University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-92-78.

A Proposal Concerning The Analysis Of Shadowns In Images By An Active Observer (Dissertation Proposal)

MS-CIS-92-78
GRASP LAB 335

Gareth D. Funka-Lea

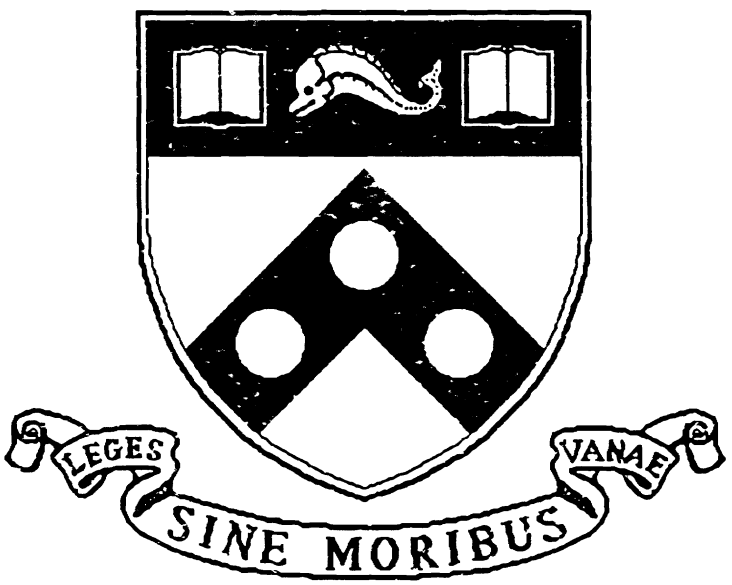

University of Pennsylvania
School of Engineering and Applied Science
Computer and Information Science Department
Philadelphia, PA 19104-6389

October 1992

A PROPOSAL CONCERNING THE ANALYSIS OF SHADOWS IN IMAGES BY AN ACTIVE OBSERVER

Gareth D. Funka-Lea

GRASP Laboratory
Dept. of Computer and Information Science
University of Pennsylvania
Philadelphia, Pennsylvania

Dissertation Proposal
Supervised by Dr. Ruzena Bajcsy

October, 1992

ABSTRACT

A Proposal Concerning the Analysis of Shadows in Images by an Active Observer

Gareth D. Funka-Lea

Shadows occur frequently in indoor scenes and outdoors on sunny days. Despite the information inherent in shadows about a scene's geometry and lighting conditions, relatively little work in image understanding has addressed the important problem of recognizing shadows. This is an even more serious failing when one considers the problems shadows pose for many visual techniques such as object recognition and shape from shading. Shadows are difficult to identify because they cannot be infallibly recognized until a scene's geometry and lighting are known. However, there are a number of cues which together strongly suggest the identification of a shadow. We present a list of these cues and methods which can be used by an *active* observer to detect shadows. By an active observer, we mean an observer that is not only mobile, but can extend a probe into its environment. The proposed approach should allow the extraction of shadows in real time. Furthermore, the identification of a shadow should improve with observing time. In order to be able to identify shadows without or prior to obtaining information about the arrangement of objects or information about the spectral properties of materials in the scene, we provide the observer with a probe with which to cast its own shadows. Any visible shadows cast by the probe can be easily identified because they will be new to the scene. These actively obtained shadows allow the observer to experimentally determine the number and location of light sources in the scene, to locate the cast shadows, and to gain information about the likely spectral changes due to shadows. We present a novel method for locating a light source and the surface on which a shadow is cast. It takes into account errors in imaging and image processing and, furthermore, it takes special advantage of the benefits of an active observer. The information gained from the probe is of particular importance in effectively using the various shadow cues. In the course of identifying shadows, we also present a new modification on an image segmentation algorithm. Our modification provides a general description of color images in terms of regions that is particularly amenable to the analysis of shadows.

Contents

CHAPTER 1

Shadows in Image Understanding

1.1 The Importance of Shadows

Artists interested in realism have for centuries relied on shadows to give a scene a sense of depth and atmosphere. In computer graphics a great deal of work has been dedicated to the accurate generation of shadows as an aid to verisimilitude (see [Woo et al. 1990] for a review or see [Takita et al. 1991] and [Thirion 1992] for more recent work). Despite the importance of shadows in generating realistic images, relatively little work has been done until recently on the role of shadows in image interpretation.

The recognition of a shadow within a scene reveals a considerable amount of information about that scene. First, that there is a directional, localized light source in the scene. For instance, shadows are not present outdoors on over-cast days. Second, knowing the correspondence between a shadow and the object causing the shadow constrains the scene geometry [Waltz 1975] [Shafer 1985a]. Third, the difference in appearance between the same surface material lit and in shadow can tell us something about the difference between the characteristics of the direct light and the light that illuminates the shadow. The information that can be gathered from shadows will be discussed in more detail throughout this work.

1.2 The Nature of Shadows

Shadows result from the obstruction of light from a source of illumination. As such, shadows have two components: one spectral and one geometric.

The spectral nature of a shadow derives from the characteristics of the light illuminating the shadow as compared to the additional light that would illuminate the same area if there

was no obstruction. Hence, shadows reveal themselves as a spectral change in radiance due to a change in the local irradiance. Shadows are often remarked to be illuminated by *ambient light.* Ambient light is generally used to refer to the light that fills an environment without having a particular localized source. The psychologist Gibson defined ambient light as the light passing through a point in space from many different directions [Gibson 1966]. However, in this work we will refer to ambient light with respect to a given location in space as being all the light striking the location except that light which emanates from a particular light source of interest. Hence the light illuminating a shadow is the ambient light. Note that by our definition, ambient light may include light from strong localized light source and that ambient light may be capable of casting shadows. However, for a scene with only one source of illumination, ambient light will be strictly reflected or scattered light.

The geometry of a shadow is determined by the nature of the illumination obstruction and the scene geometry. A light source may be only partially obstructed. In fact, for any non-point light source, the outer portion of a shadow results from the partial obstruction of the light source. This is the *penumbra* of the shadow, while the *umbra* is the part of the shadow where the light source is completely obstructed. See Figure 1.1 for an example of the shadow geometry for an extended light source.

In this work we will be dealing with shadows at particular *intermediate* scale. We will assume that part of the shape of a shadow is visible. And generally, the more of a shadow's shape that is visible, the better we will be able to recognize it as a shadow. Shadows which are individually not visible can still effect the appearance of an object or scene. At the small scale, unresolvable shadows within the microstructure of a surface will darken the appearance of the surface. At the large scale, for example, on any overcast day, an observer under the clouds is within the shadow of those clouds. However, we do not recognize this effect as shadowing unless the boundary of the shadow can be seen. Shadows at the small and large scale are not addressed in this work.

1.3 Shadow Cues

Unfortunately, recognizing shadows in a scene is a difficult problem. Shadows can only be confidently recognized once the scene geometry, materials, and spectral flux are known. By spectral flux we mean a characterization of the light at any point in the scene. This is more than just the characterization of sources of illumination because it includes the effects of inter-reflections between surfaces and the transmission properties of the environment.

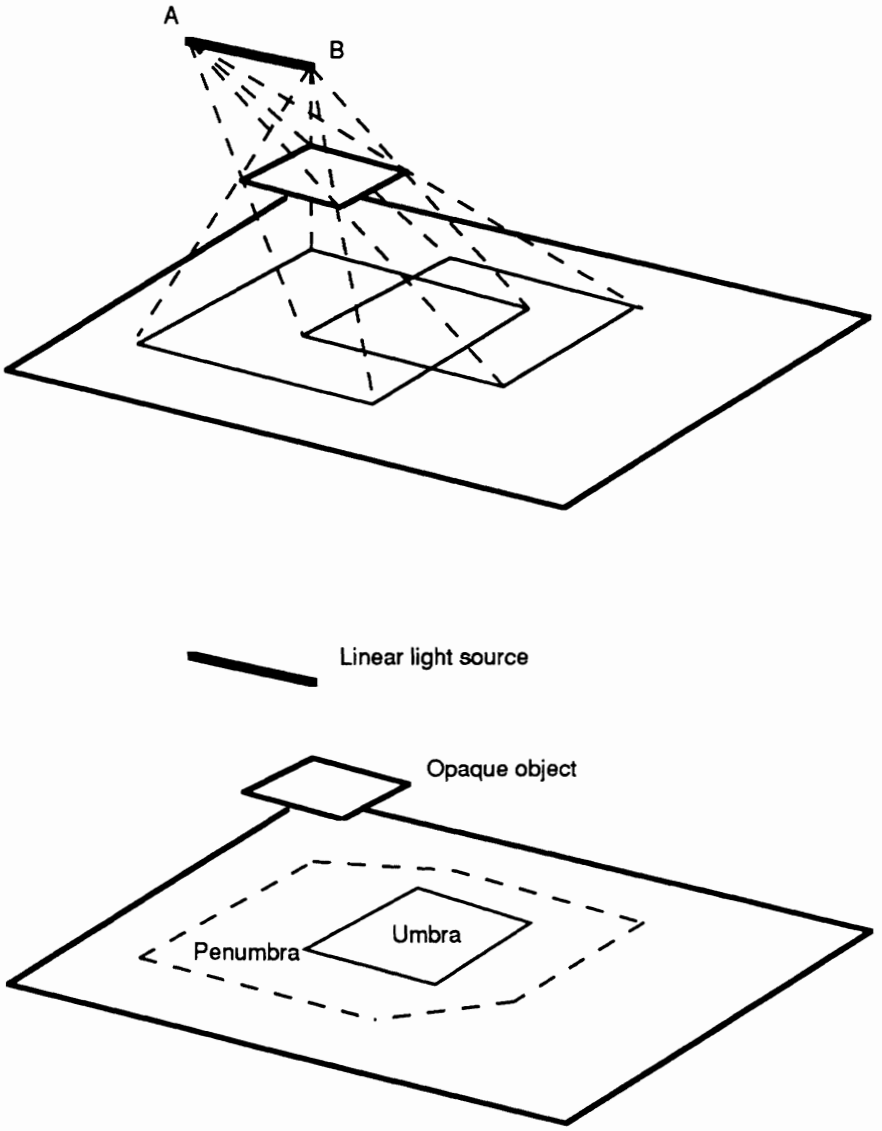

Figure 1.1: Shadow umbra and penumbra resulting from an extended light source. In this example, the scene is illuminated by a light source which has extent in only one dimension. At the top is shown the obstruction in illumination of the two end-points of the light source. At the bottom is shown the shadow umbra and penumbra. Note that the umbra is visible in the top part of the figure as the overlapping portion of the two squares cast onto the background plane. (This figure is based on figures in [Nishita et al. 1985] and [Woo et al. 1990].)

Knowing the scene's spectral flux and the material properties of a given surface we can then deduce that a change in the appearance of the surface is due to a change in irradiance. With this knowledge and the determination that light from a source of illumination has actually been obstructed, we can conclude that a shadow is present.

Detecting shadows falls into that large class of vision problems where, if most of the information about a scene is known then the remaining information can be deduced from an image of the scene. Although we cannot hope to distinguish shadows from material and geometric changes with certainty in an environment that conspires against us, there are a number of cues that suggest the presence of a shadow in a natural scene.

The most obvious spectral cue to the presence of a shadow is that a surface in shadow will appear darker than the same surface not in shadow because there is less light in a shadow. However, unless the source of illumination or the obstruction is moving we do not see the same individual points on a surface both lit and in shadow. Only, if the surface's geometry and material properties do not change rapidly at the boundary between shadow and not shadowed and if the ambient light is relatively constant across the shadow boundary, then one can be sure that a surface in shadow will be darker than than the adjacent area not in shadow.

Although the ambient light across a scene is a flux, far from highly specular materials, the ambient light is often slowly varying. Most researchers investigating color vision have made the assumption that ambient light is always uniform or the even stronger assumption characterized by [Rubin and Richards 1988]:

> **The gray world assumption:** The average of all the different albedos in the scene will be a spectrally flat *gray*, so that the ambient reflected light will have the same spectral character as the direct light.

Under the gray world assumption, the color of a surface in shadow lit only by ambient light will not differ in hue or saturation, only in intensity, relative to the same surface not in shadow. Consequently, hue and color saturation have often been used as a cue to detect one surface as whole despite any partial shadowing. However, no system currently tries to determine the local validity of assumptions about the ambient light. Such testing is an important part of the system we envision.

The changes in the irradiance of a surface that result in a shadow are unlikely to align with surface markings including surface texture. Consequently, the continuation of texture across an image region boundary is consistent with the presence of a shadow [Witkin

1982]. However, the texture detection must be unaffected by the types of spectral changes associated with shadows.

For an environment in which it has been determined that an extended light source exists, shadows can be expected to have a penumbra and umbra structure. This means that a shadow on a uniform surface material should show a decrease in intensity at the outer boundary of the shadow and a uniformly darker center region. This has often been noted as the tendency for shadows to have "soft" edges (see for example [Marr 1982]). Note that the size of the penumbra will depend on the shadow geometry. Let S be the distance between a shadow making object and the surface on which a shadow is cast. Let L be the distance between the shadow making object and the light source. The width of a penumbra will vary with $\frac{S}{L}$ (See Appendix A for a derivation).

The most obvious geometric cue to the presence of a shadow is when an object can be found between the surface on which the shadow is cast and the light source. However, this cue depends on the observer knowing where the light source and shadow are located. To use this cue to full advantage requires that the observer be able to determine three-dimensional locations of objects in the scene, which is often difficult. However, this cue can be used in a weaker sense to simple rule out the possibility of a shadow if no object can be found in the image plane between the shadow and the light source.

If the shape of the object casting a shadow is known, then the shadow must be a projection of a silhouette of the object. However, we rarely know the three-dimensional shape of objects in a scene. In addition, the shadow projection of an object's silhouette is unlikely to be a perspective projection for an extended light source. The nature of the projection can be complex. Consequently, finding the correspondence between a shadow and a known shadow making silhouette is still a difficult problem. We can see this in the shape of the penumbra cast by the square shadow making object in Figure 1.1.

Two cues follow from the fact that shadows are cast on objects in the scene. First shadow boundaries will change their direction across surface discontinuities under a general scene layout. Consider Figure 1.2 in which the shadow of a square is cast onto a rectangular solid. Also, because shadows are cast onto objects, they appear as markings on those objects. Consequently, as an observer moves about in a scene, shadows should remain stationary relative to the surfaces on which they are cast for a fixed scene geometry. The exception being when the observer casts its own shadow onto shadows in the scene. To our knowledge, this shadow cue has not been noted previously in the literature.

Below we summarize the cues that suggest the presence of a shadow.

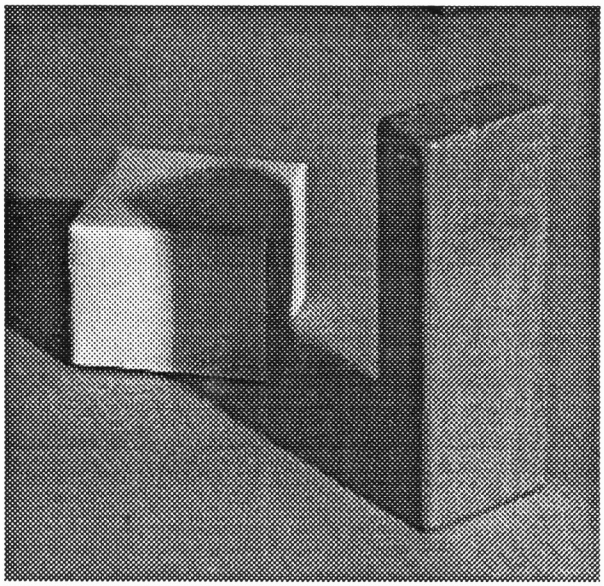

Figure 1.2: Shadow boundaries generally change with changes in geometry. Note how the boundary of the shadow of the tall block changes direction in the image across a change in the face of the small block.

- The intensity, hue, and saturation changes due to shadows tend to be predictable.
- Surface markings tend to continue across a shadow boundary and vice versa.
- For an extended light source, shadows can be expected to have a penumbra and umbra structure.
- Shadows are only possible if there is an object obstructing light from a light source.
- The shape of a shadow is the projection of a silhouette of the object obstructing light emitted from a light source.
- Shadow boundaries tend to change direction with changes in the geometry of the surfaces on which they are cast.
- Shadows remain stationary relative to the surfaces on which they are cast for a fixed scene geometry.

In this proposal we do not address methods to identify shadows by attempting to change the illumination conditions in a scene. For instance, introducing a new source of illumination

into the scene or attempting to cast a shadow where one is already expected to lie. These are powerful techniques for recognizing shadows but they are outside the scope of our current investigation.

Because there is no single image cue that indicates with certainty the presence of a shadow in a scene, shadow detection is difficult. As for certainty, the best we can hope for is that many image regions can be ruled out from the consideration of being shadows. Those image regions we choose to recognize as shadows must be those for which there are numerous pieces of supporting evidence without any contradictory evidence. However, none of the shadow cues are necessarily easy to detect and imaging uncertainties will always produce uncertainties in our scene hypotheses.

1.4 Spectral Model of Shadows

In this section, we examine in more detail the spectral characteristics of shadows as they appear in an image.

1.4.1 Model of Shadows Without Other Reflectance Effects

Let $D(\lambda)$ be the amount of energy put out at each wavelength by a source of illumination as measured at a given surface. $D(\lambda)$ is not the only illumination striking the surface. There is also the light from any other sources of illumination in the environment $L_1(\lambda), \ldots, L_n(\lambda)$ and the light $A'(\lambda)$ that has been reflected or scattered in the environment.

The total illumination striking the surface is

$$D(\lambda) + A'(\lambda) + \sum_{i=1}^{n} L_i(\lambda).$$

Assume now that an object is brought between the light source D and the surface. The reflected light in the scene changes due to reflections off the obstructing object, call it now $A(\lambda)$. So the illumination striking the surface is now

$$\alpha D(\lambda) + A(\lambda) + \sum_{i=1}^{n} L_i(\lambda)$$

where $\alpha \in [0 \ldots 1]$ indicates that the light source D will be only partially obstructed at some locations on the surface if D is not a point light source.

Assume for the moment that the surface is perfectly Lambertian and let $S(\lambda)$ be the surface reflectance (albedo). Also assume that there is no shading across the surface. For

instance, we might be concerned with just a small portion of a surface over which the amount of light striking the surface from the various sources is constant. Let $Q_j(\lambda)$ be the weighting function of the observer's camera system for the jth filter ($j = 1, \ldots, m$). Then, for a particular viewing angle, the light measured by the camera from the surface directly lit and in shadow for one filter is

$$I_j = \int_\Lambda \left(\alpha D(\lambda) + A(\lambda) + \sum_{i=1}^{n} L_i(\lambda) \right) S(\lambda)\, Q_j(\lambda)\, d\lambda. \tag{1.1}$$

Λ is the range in which $Q_j(\lambda)$ is non-zero.
We will use the following notation:

$$\mathbf{D} = \begin{bmatrix} \int_\Lambda D(\lambda)\, S(\lambda)\, Q_1(\lambda)\, d\lambda \\ \vdots \\ \int_\Lambda D(\lambda)\, S(\lambda)\, Q_m(\lambda)\, d\lambda \end{bmatrix}$$

$$\mathbf{E} = \begin{bmatrix} \int_\Lambda \left(A(\lambda) + \sum_{i=1}^{n} L_i(\lambda)\right) S(\lambda)\, Q_1(\lambda)\, d\lambda \\ \vdots \\ \int_\Lambda \left(A(\lambda) + \sum_{i=1}^{n} L_i(\lambda)\right) S(\lambda)\, Q_m(\lambda)\, d\lambda \end{bmatrix}$$

where $\mathbf{D}$ and $\mathbf{E}$ are m element vectors. From the above equation it follows that the image of the surface lit and in shadow is

$$\mathbf{I} = \alpha \mathbf{D} + \mathbf{E}. \tag{1.2}$$

The $Q_j(\lambda)$ span a sub-space of color space and Equation 1.2 is the parametric form of a line in this color sub-space with parameter α. Note that the line has end-points where $\alpha = 0$ or $\alpha = 1$. The end-point of the line at $\alpha = 0$ corresponds to the umbra of the shadow. The end-point where $\alpha = 1$ corresponds to the surface directly lit. The open interval of the line (where $0 < \alpha < 1$) corresponds to the penumbra of the shadow.

1.4.2 Shadows with Shading, Inter-Reflections, and Specularities

Shading, or variations in the amount of light striking a surface due to a change in geometry, complicate the model we have described above. For a scene with a single light source, no inter-reflections, and a uniformly colored Lambertian surface which receives varying amounts of illumination due to surface curvature or varying distance from the the light

source, the reflection from the surface will describe a linear cluster in color space [Shafer 1985b]. This linear cluster will be indistinguishable from a linear cluster in color space resulting from a shadow penumbra under the same illumination conditions.

The reflectance of a surface is further complicated when the amount of irradiance varies for each of multiple light sources. In the case where a uniformly colored Lambertian surface is illuminated by light in varying amounts from two distinctly colored lights, the reflection from the surface will describe a planar cluster in color space. If the surface is illuminated by light in varying amounts from multiple distinctly colored lights, the reflection from the surface will describe a volume in color space [Lee 1991]. For multiple, differently colored light sources, the reflection distortion in color space due to a shadow being cast on a surface will be super-imposed on the volumetric cluster due to shading. If the color volume due to shading includes a full range of received light from the obstructed light source (0% to 100%), then the distortion due to shadow will occur entirely within the volume due to shading. Otherwise, the distortion due to shadow will extend the color volume due to shading.

The light reflected from one surface onto a second surface serves as a source of irradiance for the second surface. As such, inter-reflections complicate our shadow model in the same way that multiple light source do.

As for specularities (highlights), the specularities due to the direct light source that were visible before the obstruction was introduced within the area now occupied by the umbra, will no longer exist. Specularities that fall within the penumbra will still be visible but their shape will be truncated at the boundary of the umbra. Diffuse specularities due to a rough surface [Torrance and Sparrow 1967] under an extended light source can however become dimmer even where they are still visible because the light that strikes only some of the microfacets at any given point of the surface may be obstructed.

1.5 Proposal Outline

We propose a general approach to scene interpretation for an active observer that takes into account shadows and utilizes shadows the observer casts into the scene.

We propose to have an active observer place a probe into the environment in order to cast new shadows, if possible. Because any shadow of the probe will be a new shadow in the environment, the difficulty of shadow detection will be greatly reduced. This will allow the observer to examine a known shadow in a particular environment. From a known

shadow, information about a scene's geometric and spectral properties can be recovered. In particular, the location and size of the light source and the location of the surface on which the shadow is cast can be determined. Also, an estimate can be made of the direct and ambient light in the environment.

With the information gained from the shadow probe and the shadow cues discussed above we plan to detect the naturally occurring shadows in a scene. Detection will never be a certainty because of the problems discussed, but we believe that the portions of the scene labeled as shadows will have a very high likelihood of being actual shadows.

In the course of this work, we will present a number of new methods for recognizing shadows and for interpreting actively cast shadows. We will present a new method for locating a light source and the surface on which it is cast. This method takes into account errors in imaging and image processing. This method takes special advantage of the benefits of an active observer. We will present a new technique for segmenting color images for shadow analysis. Also, our list of shadow cues presents the clearest exposition to date of what features can be used to detect shadows.

In the next chapter we briefly review work done in image understanding involving shadows. In Chapter 3 we introduce the use of a shadow probe. We show how to recognize the probe's shadow and how to use the spectral information gained from analyzing the probe's shadow to partially interpret the scene. In Chapter 4 we show how to use the shadow cast by the probe to locate the position in three dimensional space of a light source and the probe's shadow. In Chapter 5 we will present a strategy for hypothesizing the presence of shadows based on their spectral properties under conditions of limited shading and local inter-reflections. In Chapter 6 we discuss the use of cues other than color for recognizing shadows. Finally, in Chapter 7 we discuss the overall structure of our system for recognizing shadows by an active observer.

CHAPTER 2

Review of Work on Shadows

Work in image understanding involving shadows has fallen into two general categories: that which detects certain scene elements despite the presence of shadows (implicit shadow analysis) and work which tries to detect or interpret shadows in a scene (explicit shadow analysis).

2.1 Implicit Shadow Analysis

Some researchers have tried to take shadows into account by first trying to determine what their goal object should look like lit and in shadow and then using both sets of information to detect the goal object. For instance, in the problem of road detection for autonomous navigation, both [Turk et al. 1988] and [Crisman 1990] use multiple color clusters to define the appearance of a road. If a road is expected to be in shadow then at least one color cluster is used for the lit road and one for the shadowed road. However, what constitutes a road for a vehicle must initially be manually selected.

In other work, the goal object to be located is examined both lit and in shadow to determine if there is a particular spectral band or color model in which it can be easily located despite shadows. In [Ranson and Daughtry 1987] experiments were done to determine how shadows biased spectral samples taken from aerial images of fir trees. Images were taken from above of fir trees evenly placed on a large turntable. The turntable was rotated relative to the angle of the sun. Green band samples were found to be less sensitive to variations in the amount of shadows than red or infrared band samples.

Both the work on road and fauna detection is very domain specific and presents a highly impractical approach to dealing with shadows for a completely autonomous agent moving

in an unstructured environment.

As was mentioned above, hue has been suggested as an object cue that might circumvent shadow effects. For example, [Liu and Moore 1990] suggest using a three dimensional hue representation for satellite images in order to suppress (but not eliminate) shadow effects. However, hue is a successful way to ignore shadows only in the case where the light illuminating the shadows is proportional to the direct source of illumination:

$$Ambient(\lambda) = \beta Direct(\lambda)$$

and where the shadows are bright enough that hue information can still be reliably recovered.

Rubin and Richards looked for cues for material changes irrespective of shadows, highlights, surface orientation changes, or pigment density changes. They assume that a color image has been normalized and then segmented into regions which represent both material changes and all the changes listed above. Spectral samples from neighboring regions are then compared to determine if the edge between them represents a material change. From each region two spectral samples are used to define a line. If the slope of the lines from the two regions differ, then the regions are taken to be from different materials [Rubin and Richards 1988]. This works for disregarding shadows only if the gray world assumption holds.

Finally, there is a large body of work on *color constancy*. Color constancy is a term from the study of human vision, where it was noticed that we tend to recognize the color of a material despite changes in illumination. If this can be accomplished, then the change in illumination due to a shadow should not effect the recognition of a surface partly in shadow. However, under large changes of illumination the phenomenon breaks down in humans. For a review of attempts to artificially reproduce color constancy see [Bajcsy et al. 1989].

2.2 Explicit Shadow Analysis

2.2.1 Shadow Recognition

Recently within the field of image understanding, a number of researchers have begun to address the problem of recognizing shadows and utilizing the information inherent in shadows. Unfortunately, much of this work has been simplistic in nature. Often, all the dark regions of an image, as determined by a threshold, which lie next to an "object" in the direction of the the light source have been labeled as shadows [Nagao et al. 1979], [Huertas and Nevatia 1988], [Irvin and McKeown 1988], [Liow and Pavlidis 1990]. However, when the

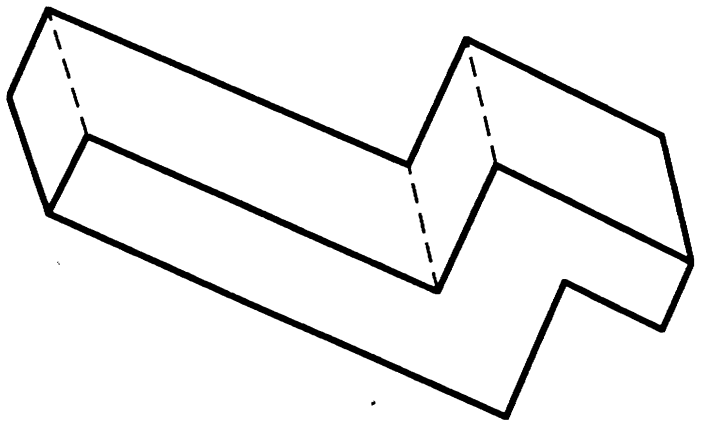

Figure 2.1: Lowe and Binford: Shadows create parallel virtual lines. When the geometric boundaries of an occluding object cast a shadow onto a surface, corners in the occluding object will lead to corners in the cast shadow. The correspondence between these corners will be found as virtual lines that are parallel or converge to a common point.

shadows in an image conform to these guidelines, these systems are reasonably successful in achieving their goal of finding buildings in aerial images.

Tables 2.1 and 2.2 review the shadow detection methods of a variety of systems.

[Gershon et al. 1986] consider two cases for shadows in their recognition scheme. For *ideal* shadows, the light illuminating the shadow is taken as proportional to the direct illumination

$$Ambient(\lambda) = \beta Direct(\lambda)).$$

For this case the measured reflection values for the same surface material under the same viewing conditions lit (R, G, B) and in shadow $(R_{shad}, G_{shad}, B_{shad})$ will be proportional:

$$\begin{aligned} R_{shad} &= \beta_R R, \\ G_{shad} &= \beta_G G, \\ B_{shad} &= \beta_B B, \end{aligned} \tag{2.1}$$

$$\beta_R = \beta_G = \beta_B.$$

The other case is the more interesting. For *non-ideal* shadows, a reflected illumination is taken to be irradiating the scene in addition to the proportional ambient irradiation in the ideal case. The reflection from the directly lit surface is now

Reference		
Input Data	**Spectral Methods**	**Geometric Methods**
[Adjouadi 1986]		
BW	Threshold based on histogram. 1D or 2D correlation across edge. Power spectral comparison across edge.	
[Gershon et al. 1986], [Lee 1991]		
Color	Change in intensity with limited or no change in hue or saturation.	
[Huertas and Nevatia 1988]		
BW	Threshold based on histogram.	Match object and shadow corners. Shadow on opposite side from light source.
[Irvin and McKeown 1988]		
BW	Threshold based on dark regions near to initial building hypotheses.	Building adjacent to shadow relative to light direction.
[Jiang and Ward 1992]		
BW	Threshold defined by offset from line fit to pixels at scan-line endpoints. (Assume endpoints on scene background and not in shadow.)	Penumbra present, Cast shadow / self-shadow structure, or Shadow on opposite side from light source.

Table 2.1: References on Detecting Shadows: Part I. BW stands for black and white (gray-scale) images.

Reference		
Input Data	**Spectral Methods**	**Geometric Methods**
[Liow and Pavlidis 1990]		
BW	Threshold on average gradient across edge.	Shadow on opposite side from light source.
[Lowe and Binford 1985]		
Curves		Corners in correspondence relative to light direction. Constraint propagation. See Figure 2.1.
[Nagao et al. 1979]		
Color	Threshold based on histogram of image intensity.	Object adjacent to shadow relative to light direction.
[Scanlan et al. 1990]		
BW	Threshold = median of local image means.	
[Thompson et al. 1987]		
Stereo BW		Shadows move when illumination direction changes.
[Witkin 1982]		
BW	Correlation across an edge with a shift in regression parameters for curves fit on either side of the edge.	

Table 2.2: References on Detecting Shadows: Part II.

$$(r + \beta R + R, g + \beta G + G, b + \beta B + B)$$

and from the shadowed surface is

$$(r + \beta R, g + \beta G, b + \beta B).$$

The authors define the *pull factor* as a measure of the deviation from the ideal shadow case. The pull factor is the normalized magnitude of the ambient illumination not proportional to the direct light in the direction perpendicular to (R, G, B). In the two-dimensional $(Red, Green)$ space, the pull factor is

$$pull\_factor = \frac{(r, g) \cdot (-G, R)}{||(R, G)||^2} = \frac{|gR - rG|}{R^2 + G^2}. \tag{2.2}$$

The authors assume that the pull factor can be determined by a higher-level process. The pull factor is used as a bounds on the difference in proportionality between two regions if they are to be considered as shadows. The authors use double-opponent filters to measure the relative change in $(Red, Green, Blue)$ across a color edge but the shadow criteria is approximately:

$$|\beta_R - \beta_G|, |\beta_R - \beta_B|, |\beta_G - \beta_B| < pull\_factor. \tag{2.3}$$

Although, we make use of a more general shadow model then the non-ideal shadow model of [Gershon et al. 1986], a measure of the difference between ambient and direct light like the pull factor plays a role in our recognition of shadows. This will be shown in Chapter 5.

2.2.2 Shadow Interpretation

Although, his work was limited to the analysis of "block worlds," [Waltz 1975] was able early on to demonstrate the advantage of introducing shadow interpretation into computer vision systems. By adding shadow labels to his curve classification scheme, Waltz was able to improve the performance of his constraint satisfaction system for interpreting line drawings. This improvement results from the added scene constraints shadows provide. In examining the particular constraints shadows places on a scene, he also identified many of the principles later shadow analysis systems would use.

The most important work on the interpretation of the geometric information inherent in shadows is [Shafer 1985a]. Shafer assumes that shadows have already been detected in an image and that the correspondence between the shadow and the shadow making object is known. With this information, Shafer studies what three-dimensional information can be derived about the object casting the shadow and surfaces on which it is cast. The study is done in terms of a case analysis. Initially, only simple scenes are considered. Using the results gained from these cases, more complex scenes are examined.

The simplest case Shafer examines he calls the *basic shadow problem.* In this case, there are two flat surfaces, one of which casts a shadow on the other due to a single light source. The light source is assumed to be infinitely far away so that light rays emanating from the source are parallel. The flat surface casting the shadow is assumed to be a polygon. The goal of the analysis is to derive a description of the two surface planes in terms of their surface gradients and to determine the direction of illumination. The problem has six unknowns: two for each surface gradient and two for the direction of illumination. However, Shafer shows that there are only three constraints provided by the correspondence of the shadow and its shadow making object. Hence, additional information is needed to solve the problem.

The basic shadow problem is shown to be linear under orthographic projection. Under perspective projection, however, the problem involves quotients of quadratic equations.

Shafer considers a number of extensions to the basic shadow problem. Under none of the extensions does the problem become fully constrained. For instance, multiple light sources provide no additional constraint on the problem. Others cases considered include: shadows cast on polyhedra, shadows cast by polyhedra, and shadows cast by curved surfaces.

It is important to note that knowing the relative gradients of surfaces only partially describes the three dimensional relationship between objects. For most tasks, one also needs to know information about the relative position of objects, such as whether or not two objects touch. Fortunately, this type of structural information about objects can also be gained by examining shadows [Waltz 1975]. See Figure 2.2 for an example.

Waltz and Shafer are both important works in enunciating what geometric information about a scene can and cannot be gained from shadows under ideal circumstances. Neither, however, addresses the issue of recognizing shadows or recovering an image segmentation that will support their analysis. Both work with perfect line drawings. Consequently, it is not clear that their analysis is practical. For a more recent work on the shadow interpretation of edges see [Hambrick et al. 1987].

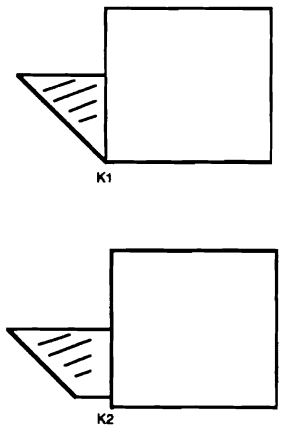

Figure 2.2: Shadows and T-Junctions. Two figures are shown above of a square and the shadow it casts on a background plane. In the top figure, the shadow and the square touch at corner $K1$ and we can conclude that the square is resting on the background plane. In the bottom figure, the shadow and the square do not touch at corner $K2$. Instead, the shadows touches the square at a T-junction just above $K2$. From this we can conclude that the square is not in contact with the background plane [Waltz 1975].

The geometric scene information from shadows that has been recovered in practice is much simpler or involves more tightly constrained environments than in Waltz and Shafer's work. In aerial images where the location of the ground and of the light source is known, shadows have been used to determine the approximate height of buildings [Lowe and Binford 1985], [Huertas and Nevatia 1988], [Irvin and McKeown 1988]. In [Kender and Smith 1987], images were taken during the strictly controlled motion of a single light source in order to recover three dimensional structure. The key idea of the method is that a surface will first be lit when the angle of the illumination becomes tangential to the surface. The method requires a very strictly controlled environment and a large number of images.

2.3 Conclusion

The usefulness of recognizing shadows has been amply shown. This is revealed by work in computer vision in two ways. First, many computer vision modules that have been developed assume that the effects of shadows have already been taken into account before they begin processing. Second, shadows have been shown to provide useful information about a scene in their own right. What is needed is better methods for identifying shadows

and more successful techniques for utilizing the information inherent in shadows. The latter depends in part on a deeper understanding of what shadows mean for a scene. This proposal addresses these key issues.

CHAPTER 3

Shadow Probe

We propose that an autonomous agent should place a probe into the environment in order to try to make its own shadow. This probe could be separate from the agent's other actuators or the functions of the shadow probe could be combined with other functions in a multi-purpose actuator. For instance, a gripper can be used to make a shadow. However, grippers tend to have complex silhouettes and hence produce shadows with complex shapes. A complex shape can make shadow identification more difficult. Therefore, for this current work we propose to use a square for the shadow probe. This square may be attached along a robotic arm with a gripper at its end or to an independent appendage – this is not a concern of our work. We require only that the shadow probe can be extended into the environment from a recessed place on the agent where it does not cast a shadow. That the agent can move is essential; that the agent can move the shadow extended probe independently of itself is helpful but not essential. Only some issues of how to make the best use of a shadow probe that can be moved independently will be addressed here. Many of the practical issues of an independently movable shadow probe will depend on the architecture of the agent.

The shadow probe should also have at least one side that is or can be made visible to the agent. This side can then be used to judge whether or not the light within the environment has changed by monitoring changes in the appearance of the shadow probe.

3.1 Detecting the Shadow of the Probe

We assume that the environment does not change during the time it takes to extend the shadow probe from its recess in the agent. Consequently, if the probe casts a visible shadow then the shadow can be found simply by examining the difference between images taken

before the probe was extended and after it was extended. See Figure 3.1 for an example.

Interesting problems arise in detecting the probe's shadow when the shadow has a penumbra and umbra structure or when there are multiple shadows cast by the probe. If a penumbra and umbra is present in the shadow, then the agent should have a way of locating the umbra versus the penumbra.

If the shadows cast by multiple light sources are distinct, then they can be found by simple region growing until all the detected shadows are labeled. The problem of disambiguating over-lapping shadows cast by a single object due to multiple light sources is not addressed here. It is a difficult problem. Below we assume that each shadow is due the obstruction of a single light source.

3.1.1 Detecting the Umbra and Penumbra

If a shadow is cast onto a uniformly colored surface and there are no other direct sources of light illuminating the shadow, then the only image structure within the shadow is due to the penumbra and umbra dichotomy. However, it is not always possible to tell if a region of an image corresponds to a uniformly colored surface because of shading. Note that if the shadow is illuminated only by ambient light without a strong direction, then there will be no shading within the shadow. Shading due to a light source that illuminates the probe's shadow can be accounted for since its effect will be constant before and after the probe is extended into the scene. Rather than trying to determine if the surface on which a shadow is cast is uniformly colored, we look for a technique to find the umbra and penumbra of a shadow for any type of surface.

Color changes on the surface on which a shadow is cast will show themselves as variations in color within the shadow. The color changes could be confused with the penumbra of the shadow. One possible technique to address this issue involves looking at the ratio of the the images before and after the probe casts its shadow. Let $S(\lambda, x, y)$ be the surface reflectance as measured within each viewing cone defined by pixel (x, y). Let $D(\lambda, x, y)$ be the spectrum of the source of illumination measured at the surface on which the shadow is to be cast for the viewing cone (x, y). Let $A(\lambda, x, y)$ be the other light illuminating this surface. Assume that $D(\lambda, x, y)$ and $A(\lambda, x, y)$ are locally constant over (x, y). The ratio of the light reflected from the surface after and before the probe is introduced is

$$\frac{(\alpha(x,y)\,D(\lambda,x,y) + A(\lambda,x,y))\,S(\lambda,x,y)}{(D(\lambda,x,y) + A(\lambda,x,y))\,S(\lambda,x,y)} = \frac{\alpha(x,y)\,D(\lambda,x,y) + A(\lambda,x,y)}{D(\lambda,x,y) + A(\lambda,x,y)} \tag{3.1}$$

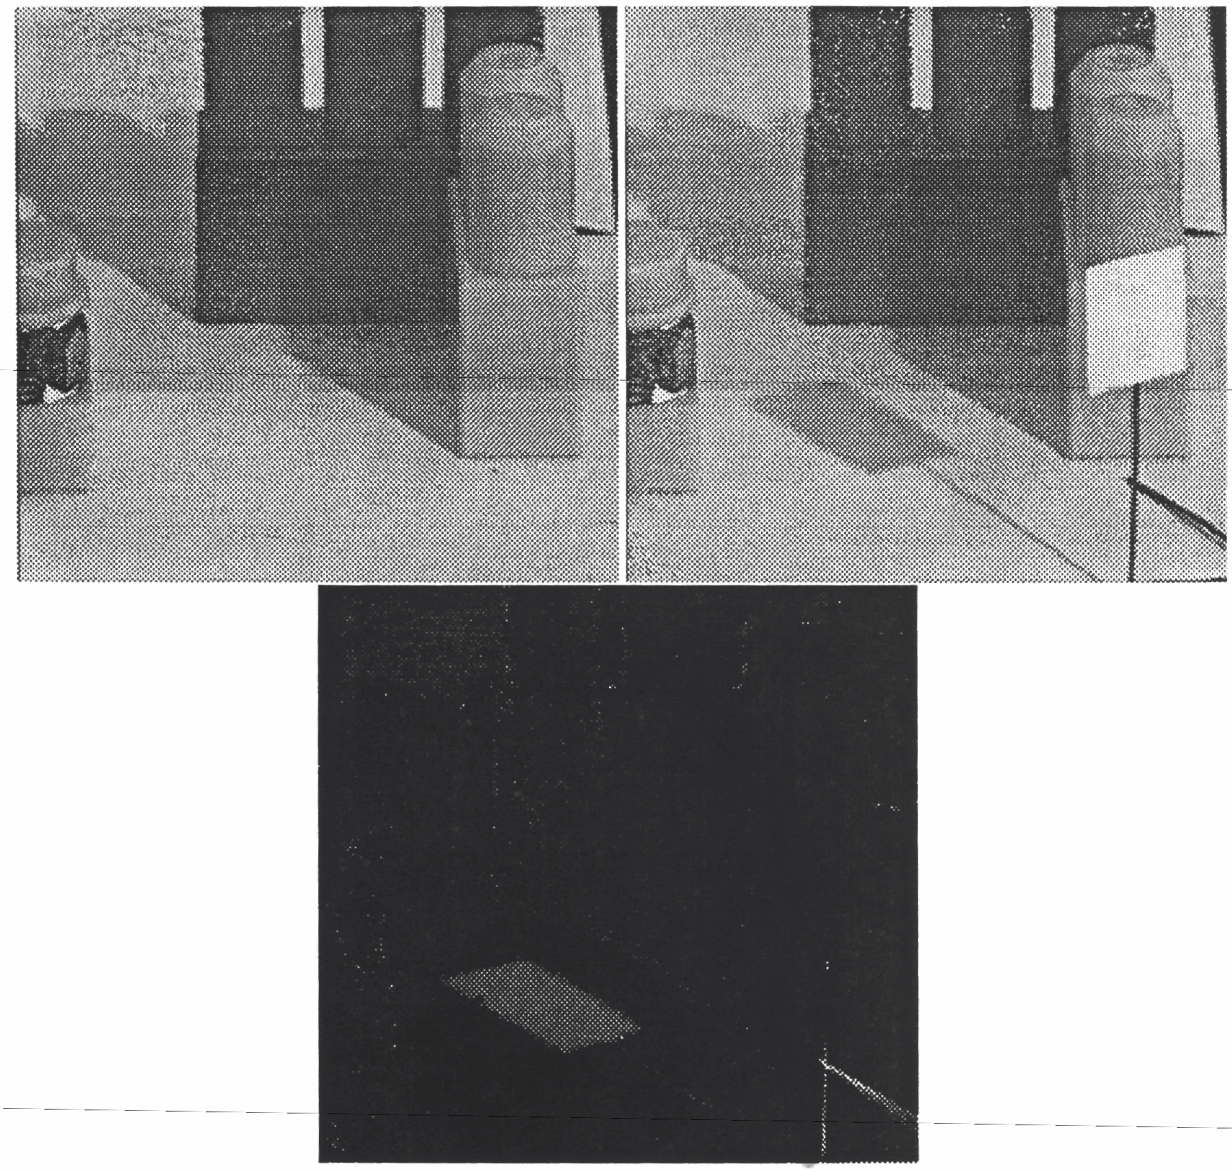

Figure 3.1: Detecting the shadow of the probe. At the top left is the original image of the scene. At the top right is the scene after the introduction of the shadow probe. At the bottom is the difference of the two images where the probe image is darker than the original. This has had the effect of removing the probe itself, although in general the probe should be removed based on the knowledge of its position. The probe arm is still visible but can be discounted because of its width.

where $\alpha(x,y) \in [0 \ldots 1]$ indicates the degree of partial occlusion of the light source on the surface within the viewing cone defined by pixel (x,y). Note that the ratio in Equation 3.1 is independent of the surface reflectance and in fact varies only with $\alpha(x,y)$.

Unfortunately, we cannot directly measure this ratio, we have only the measurements that the camera takes. CCD cameras are integrators and the measurements taken are integrals over wavelength λ for a given filter $Q_j(\lambda)$. Consequently, the ratio of the images before and after the probe is introduced is

$$\frac{\int_\Lambda (\alpha(x,y)\, D(\lambda,x,y) + A(\lambda,x,y))\, S(\lambda,x,y)\, Q_j(\lambda)\, d\lambda}{\int_\Lambda (D(\lambda,x,y) + A(\lambda,x,y))\, S(\lambda,x,y)\, Q_j(\lambda)\, d\lambda} \tag{3.2}$$

This ratio is only independent of $S(\lambda,x,y)$ if the light illuminating the surface and the surface reflectance are separable in the integrals, in other words that:

$$\int_\Lambda (D(\lambda,x,y) + A(\lambda,x,y))\, S(\lambda,x,y)\, d\lambda = \int_\Lambda (D(\lambda,x,y) + A(\lambda,x,y))\, d\lambda \int_\Lambda S(\lambda,x,y)\, d\lambda \tag{3.3}$$

This is the case when either of $(D(\lambda,x,y) + A(\lambda,x,y))$ or $S(\lambda,x,y)$ are uniform over λ. However, the ambient light illuminating a surface is always partly correlated with that surface and so we cannot expect Equation 3.3 to hold. Figure 3.2 demonstrates how far from uniform the ratio in Equation 3.2 can be. This demonstration is even more convincing in the original color images than in the black-and-white reproductions presented here.

Instead of looking at the values of a ratio, we choose to re-formulate the relationship so that we can look for a signal with a certain form. Let

$$\mathbf{D}(x,y) = \begin{bmatrix} \int_\Lambda D(\lambda,x,y)\, S(\lambda,x,y)\, Q_1(\lambda)\, d\lambda \\ \vdots \\ \int_\Lambda D(\lambda,x,y)\, S(\lambda,x,y)\, Q_m(\lambda)\, d\lambda \end{bmatrix},$$

$$\mathbf{A}(x,y) = \begin{bmatrix} \int_\Lambda A(\lambda,x,y)\, S(\lambda,x,y)\, Q_1(\lambda)\, d\lambda \\ \vdots \\ \int_\Lambda A(\lambda,x,y)\, S(\lambda,x,y)\, Q_m(\lambda)\, d\lambda \end{bmatrix}.$$

Then the image of the surface when not shadowed is $I(x,y) = \mathbf{D}(x,y) + \mathbf{A}(x,y)$ and the image of the surface in shadow is $I_{sh}(x,y) = \alpha(x,y)\,\mathbf{D}(x,y) + \mathbf{A}(x,y)$. Consider

$$R(x,y) = \frac{I_{sh}(x,y)}{I(x,y) - I_{sh}(x,y)} \tag{3.4}$$

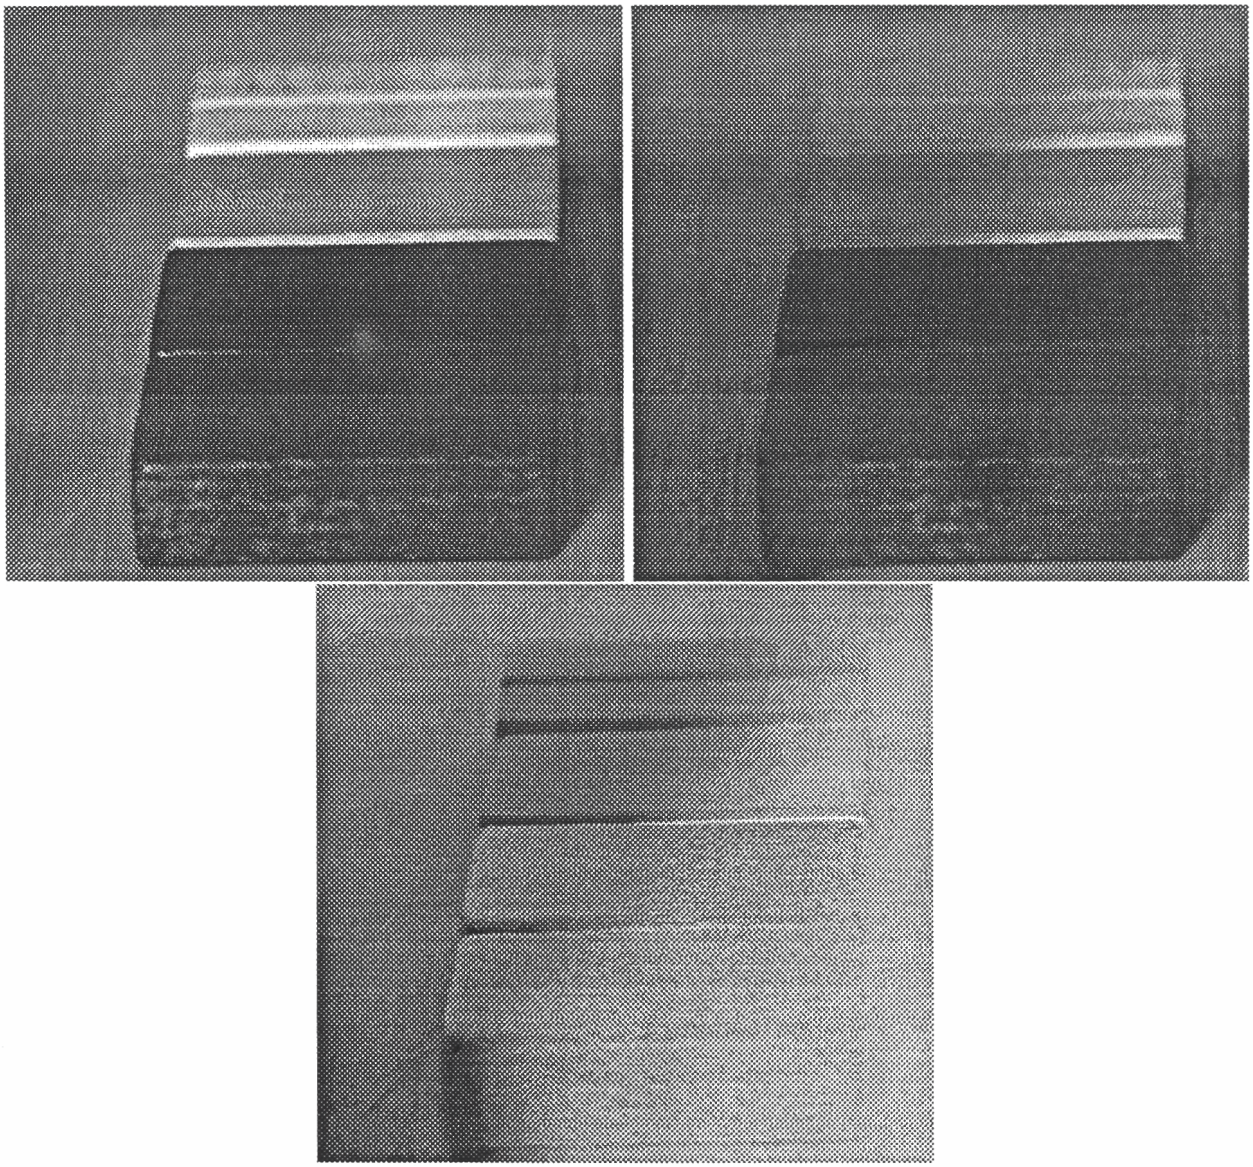

Figure 3.2: Using the ratio of shadowed to unshadowed images as a cue to surfaces in shadow. At the top left is the original image of a scene containing 5 wood blocks with 4 different colors. At the top right is the scene after the introduction of a new shadow to the left part of the scene. At the bottom is the ratio of the two images. Note that the ratio varies with the color of the blocks. In fact, the ratio generally has the complementary hue of the blocks in the original image.

$$R(x,y) = \frac{\alpha(x,y)\,\mathbf{D}(x,y) + \mathbf{A}(x,y)}{\mathbf{D}(x,y) - \alpha(x,y)\,\mathbf{D}(x,y)} \tag{3.5}$$

$$R(x,y) = \frac{\alpha(x,y)}{1-\alpha(x,y)} + \frac{\mathbf{A}(x,y)}{(1-\alpha(x,y))\,\mathbf{D}(x,y)}. \tag{3.6}$$

$R(x,y)$ is monotonically increasing with $\alpha(x,y) \in [0\ldots1]$ and

$$\lim_{\alpha(x,y)\to 1} R(x,y) = \infty, \tag{3.7}$$

$$R(x,y)\,|_{\alpha(x,y)=0} = \frac{\mathbf{A}(x,y)}{\mathbf{D}(x,y)}. \tag{3.8}$$

Consequently, the outer edge of the penumbra can be found when $R(x,y)$ approaches infinity. To find the inner edge of the penumbra (outer edge of the umbra) we need to be able to determine when

$$R(x,y) = \frac{\mathbf{A}(x,y)}{\mathbf{D}(x,y)}. \tag{3.9}$$

$\mathbf{D}(x,y)$ and $\mathbf{A}(x,y)$ vary with surface reflectance and with the amount of shading. We assume that surface reflectance and shading generally vary slowly and so the ratio in Equation 3.9 will be locally uniform. Consequently, we propose to test for the shadow umbra by determining where $R(x,y)$ ceases to be locally uniform. See Figure 3.3 for a demonstration.

In practice $R(x,y)$ is very sensitive to noise in the images. $R(x,y)$ is often infinite for small differences between the two images that arise from noise in the camera system. The effect of the noise can be greatly reduced through a few simple steps. First, when taking the difference of two images we take the minimum difference found within a 3x3 window centered at each pixel. Here we assume that the images can move by as much as a pixel. Second, we can ignore much of the noise in the difference of the images by suppressing all increases in pixel values between the images because we know the shadow will be darker. In the case where noise is as likely to increase as to decrease the value of a pixel, we can expect to remove half the noise with this technique. Finally, we expect the shadow of the probe to occupy more than a few isolated pixels in the image if it is present at all, so we can suppress isolated non-zero pixels.

3.2 Spectral Samples for One Location

With the umbra of the probe's shadow located, the observer has a spectral sample of one location in space illuminated without a direct source of light. From the image prior to

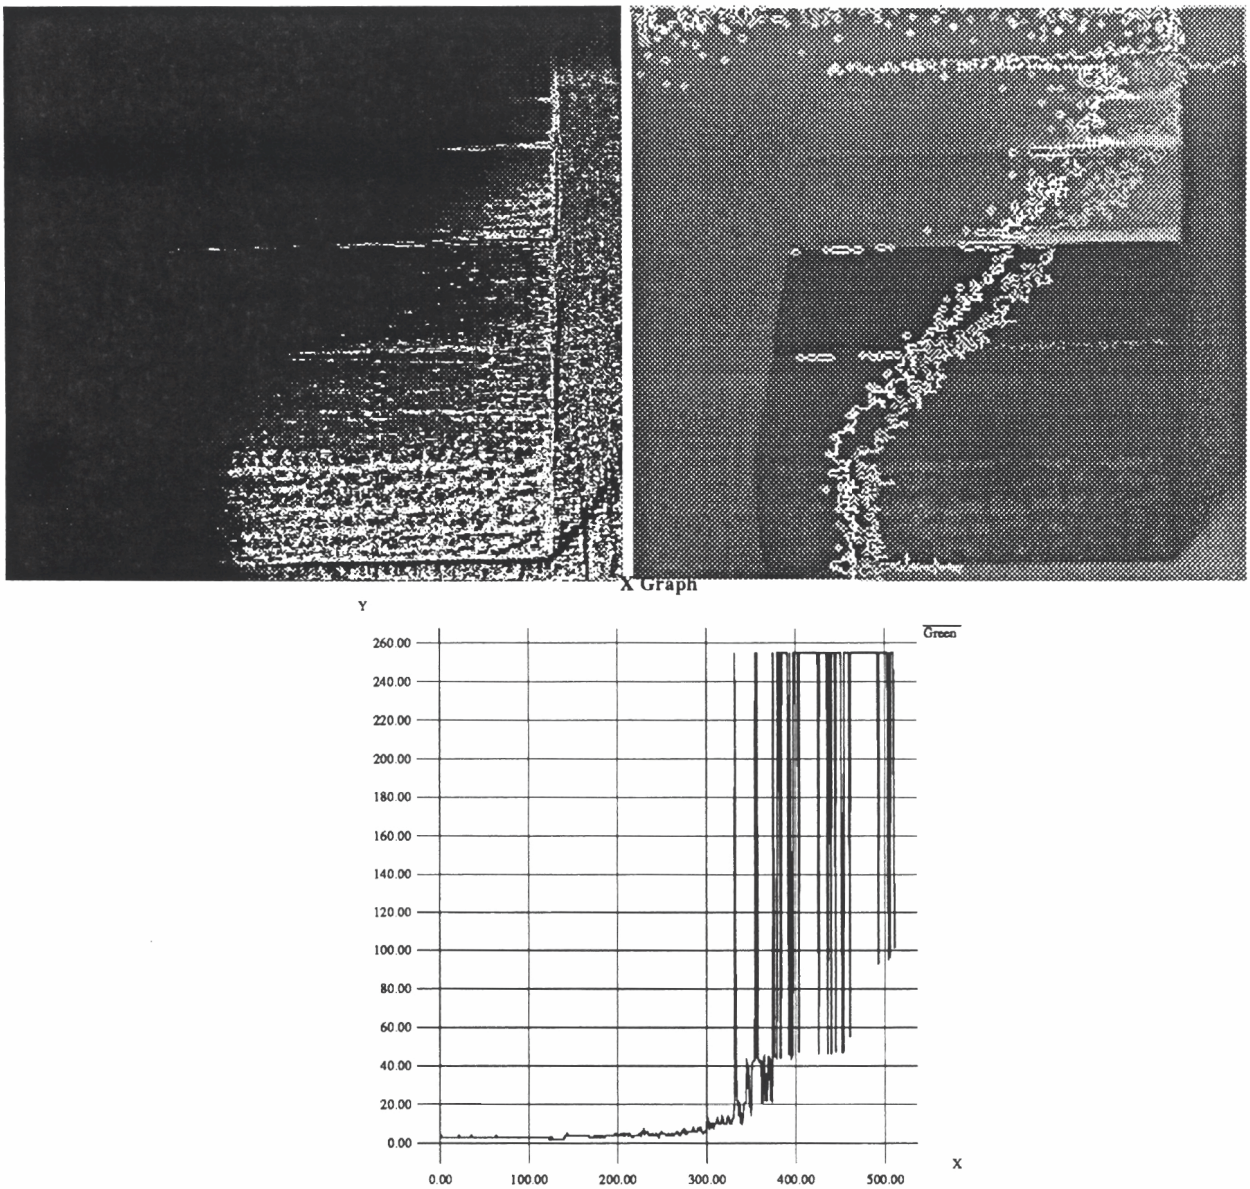

Figure 3.3: Locating the umbra and penumbra in the probe's shadow. At the top left is the image $R(x, y)$ as defined in Equation 3.4. At the top right, the boundary for the umbra is shown in white and the boundary for the penumbra is shown in gray on the shadow image. At the bottom is shown a horizontal slice along the green plane from the RGB image $R(x, y)$. Infinity is a value of 255.

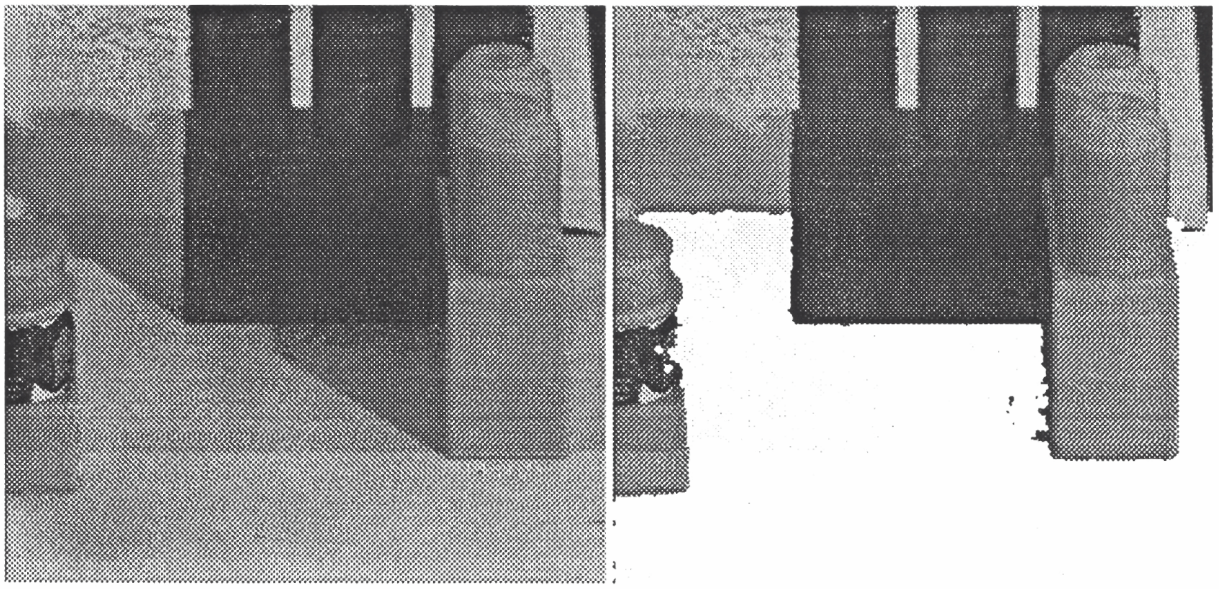

Figure 3.4: Initial image segmentation. At left is the original image and at right is the image region found based on the probe spectral sample.

introducing the probe, the observer also has a spectral sample of the same location with the direct source of light. We run a segmentation algorithm on the area in the image with the direct light source where the probe shadow subsequently appears. This is done in order to determine if multiple regions are present at this location. Then for each region, we extend the segmentation to the surrounding image. In extending the segmentation, we take into account the appearance of the each region lit and in shadow so that the segmentation will not separate other shadows on the same region. The segmentation is based on [Leonardis et al. 1990] where the models for each region are determined by the two spectral samples. This segmentation technique is briefly described in Chapter 5.

Figure 3.4 presents the results of this segmentation for the sample image used in Figure 3.1. Note that the pink piece of paper on which the probe's shadow was cast has been successfully found despite the shadow cast by the light colored block. However, a few pixels have not been recognized as belonging to the paper in the shadow from the block because of the strong reflection from the specular surface of the plastic block.

3.3 Conclusion

Here we have described how to locate the shadow of a probe and how to find the penumbra and umbra of the shadow. We have also described how to use the spectral/color information gained from the shadow the agent casts to provide a partial segmentation of an image of the agent's environment. In the next chapter, we describe how to use the probe to determine the location of the light source and the location of the shadow. We will also discuss some issues in the placement of the probe.

In Chapter 5 we will propose to use the spectral samples found from the shadow of the probe as data from which to estimate any trends in the appearance of shadows in a scene. Observing one specific location directly lit and in shadow, we would like to estimate the differences in spectra of the ambient and direct light. And hence, how shadows should change the appearance of a surface. It will help to have a large collection of data samples of different colored surfaces lit and in shadow. This can be achieved by moving the observer through the environment and doing repeated experiments with the probe. If this is not possible then it may be necessary to augment the shadow probe with a plate onto which the agent casts the probe's shadow. On this plate, a collection of color samples could be provided so that the agent would be guaranteed a good data set. The plate adds to the complexity of the shadow probe and to the complexity of its placement so that a shadow can be observed. In this proposal we do not intend to address the issues of having a two part probe.

CHAPTER 4

Shadow Probe Geometry

4.1 Locating a Light Source

In order to decide if an image region corresponds to a shadow in a scene, one must determine if a light source is being obstructed in a manner consistent with a shadow at that location in space and that there is some object onto which the shadow can be cast. Consequently, determining the location of any light sources within a scene is an important precursor to shadow identification. It is also important to determine the extent of a light source relative to the obstruction and relative to the location of the shadow cast by this light source. Point light sources produce shadows with strong edges while extended light sources may produce shadows with broad edges or may produce no shadows at all.

The next four sections provide the motivation and the high level details of our method for using shadows to reliably located a light source. Following these sections, in the Strategy Review section, the details of our method are mapped out.

4.1.1 Shafer's Contribution

It has long been recognized that a point to point correspondence in an image between a shadow and the terminator along a shadow making object indicates the direction of the obstructed light source projected onto the image plane. [Shafer 1985a] provides a careful consideration of the case where the corners of a polygon or polyhedra cast a known shadow. The corners are used for the point to point correspondence. Four cases are considered by Shafer for locating a point light source:

1. Orthographic projection with the light source infinitely far away,

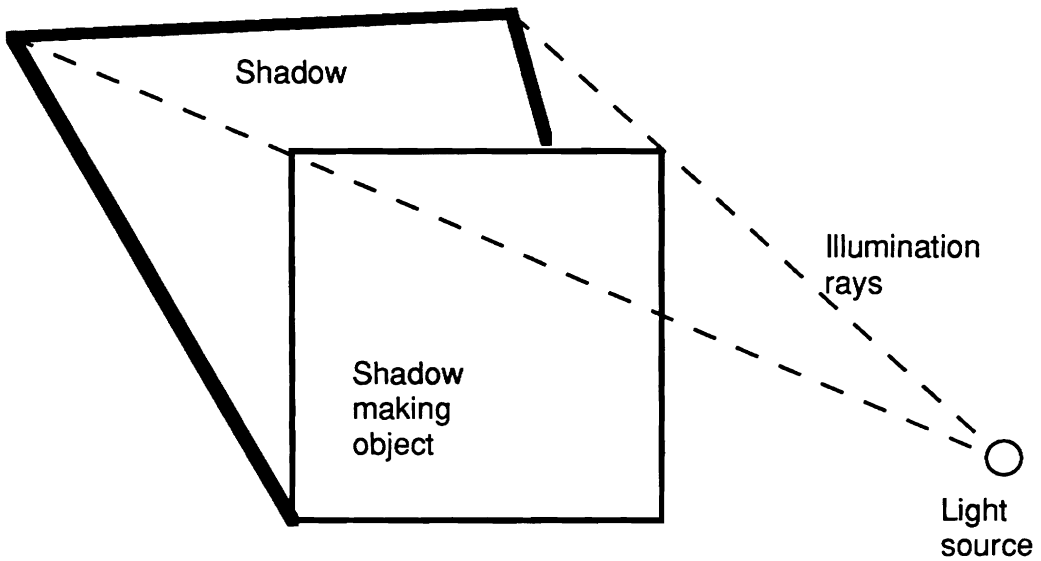

Figure 4.1: Illumination rays for a light source at a finite distance.

2. Orthographic projection with the light source a finite distance away,

3. Perspective projection with the light source infinitely far away,

4. Perspective projection with the light source a finite distance away.

The direction of a light source infinitely far away from a viewer can be completely specified by two angles (its slant and tilt) while a light source located at a finite distance must be specified by three values (for example, its X, Y, Z coordinates). Call the line in an image from the corner of a polygonal face to the shadow cast by that corner an *illumination ray*. Under orthography, an illumination ray provides the tilt of a light source infinitely far from the camera. For a light source at a finite distance viewed under orthography, the illumination rays from two corners will intersect at the coordinates of the light source in the image plane (the X and Y coordinates). See Figure 4.1. Consequently, to determine if a light source is infinitely far from the viewer under orthography, it is sufficient to check that the illumination rays from two corners are parallel. Using a perspective camera model the case for the case of a light source infinitely far from the camera, illumination rays from two corners will converge at a *vanishing point*. A line through the focal point of the camera and the vanishing point on the image plane completely specifies the location of the light source (slant and tilt). For a light source at a finite distance under perspective viewing,

	The Light Source		
	Infinitely far	Finite distance	Distinguish finite from infinite
Orthography	2 parameters 1 known	3 parameters 2 known	Can
Perspective	2 parameters 2 known	3 parameters 2 known	Can not

Table 4.1: Shafer's four cases for locating a light source based on an image of a shadow.

the intersection of two illumination rays only constrains two of the three coordinates of the of the light source. However, under perspective viewing there is no way to determine if a light source is infinitely far away or not. Illumination rays always converge at a point light source under perspectivity. These results are summarized in table 4.1.

There are limitations with using Shafer's analysis of using shadows to locate a light source. In part these limitations follow from the fact that Shafer's analysis is done nearly entirely in gradient space. He is interested in recovering surface normals and vector directions. However, the absolute location in 3-space of the light source can be important in analyzing shadows. He also deals strictly with point light sources, which are rare in most environments. Finally, the triangulation he uses to locate a light source is very sensitive to errors in the determination of the location of a shadow and the shadow-making object.

4.1.2 Locating a Light Source in 3-Space

It is easy to see why in general the three-dimensional location of a light source cannot be determined from a single image. Consider the shadow of a square cast on a plane by a single, point light source. A plane can be defined by the shadow cast by one corner of the square, the corresponding actual corner of the square, and the focal point of the camera. The light source must lie in this plane. From a second corner of the square, the shadow cast by this corner, and the focal point we can define a different plane in which the light source must also lie. The intersection of these two planes is a line and the line includes both the light source and the focal point of the camera. Examining any third point on the square and the shadow it casts, we can define another plane, but the intersection of this third plane with the previous two planes results in the same line defined by the light source and the focal point of the camera. No additional constraint on the location of the light source is gained by examining more than two shadow points [Shafer 1985a]. In conclusion,

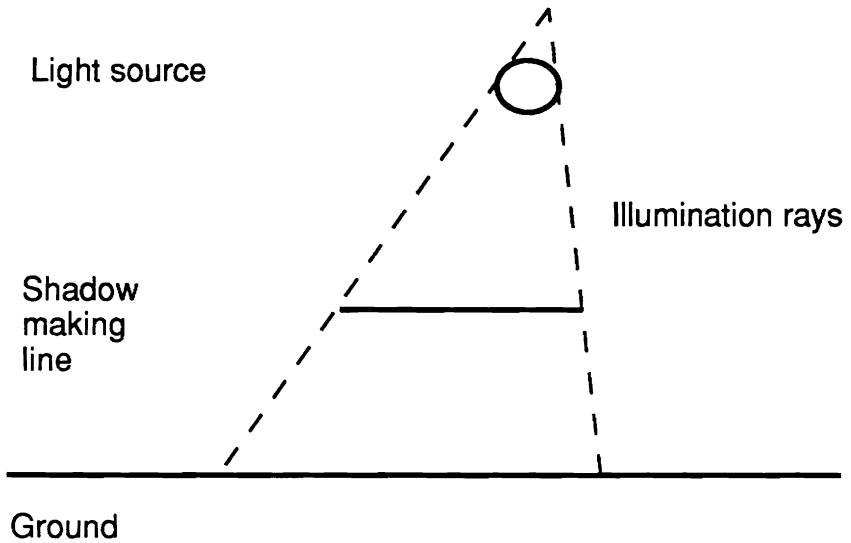

Figure 4.2: 2D Shadow Geometry. A line casts a shadow onto a line below it from a circular light source. Illumination rays are shown grazing the end points of the shadow making line.

the location of the light source can only be determined up to a line from a single image.

However, since our paradigm is active vision, we need not be satisfied with the information that can be gained from a single image. We can move the observer to acquire further constraints on the location of a light source. Moving the observer, moves the focal point of the camera and hence examining the intersection of a new plane as defined above with the two previous planes will now provide a distinct constraint on the location of the light source. Therefore, the location of a point light source can be uniquely determined by examining three illumination rays as long as at least one is from an image taken from a different viewing point. Hence our need for a mobile observer.

4.1.3 Extended Light Sources

The problem of locating a point light source is primarily one of triangulation. Two known illumination rays are extended and their intersection is found. Their intersection defines the position of the light source. However, for an extended light source this need not be the case. Consider the two dimensional shadow geometry shown in Figure 4.2. A line is casting a shadow onto another line below it. The light source is circular in extent. Although the illumination rays intersect, they intersect beyond the location of the light source because

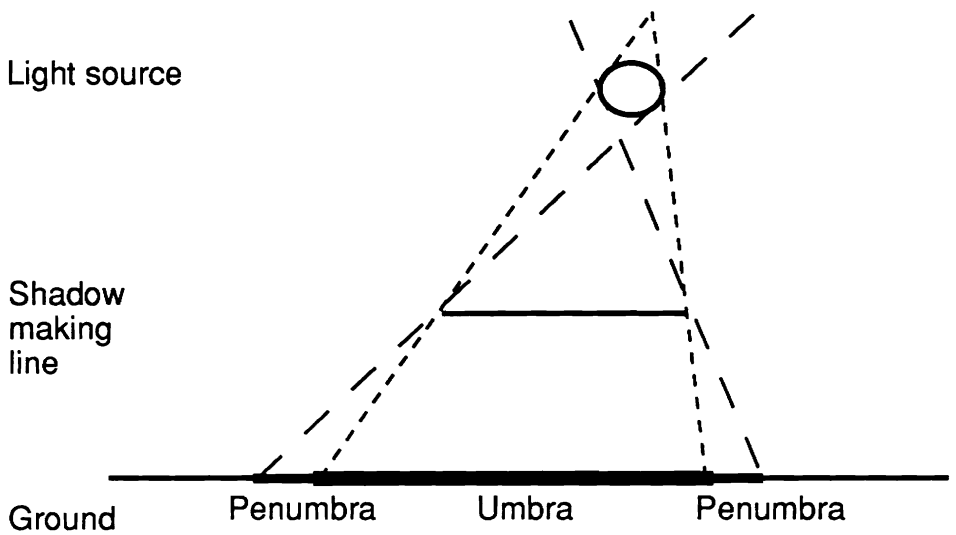

Figure 4.3: 2D Penumbra.

the two illumination rays graze the light source at different positions. For a light source with extent in 3-space, illumination rays need not intersect at all.

We also have to be careful in defining what we mean by the outer boundary of a shadow cast by an extended light source. Such a shadow will have an outer boundary where the light source is only partially obstructed by the shadow making object (at the penumbra) and an inner boundary where the the light source is completely obstructed (at the umbra). The penumbra and umbra provide different information about the location of the light source. This can be seen in Figure 4.3. Each of the illumination rays pictured grazes the light source at a different tangent point on the surface of the light source.

However, the important insight to be gained from Figure 4.3 is that the area between the various illumination rays confines the location of the extended light source. For this particular example, and as is often the case, the illumination rays defined from the umbra constrain how *close* the light source is to the shadow making object and the illumination rays defined from the penumbra constrain how *far* the light source is from the shadow making object. But for all configurations and shapes, the area between the illumination rays bounds the shape and location of the light source.

In three dimensions the bounding illumination rays need to be generalized to bounding illumination planes. A three dimensional light source is then constrained to lie within a volume bounded by illumination planes.

4.1.4 Coping with Errors

The errors encountered in locating a light source fall into three categories. There are errors localizing features in an image due to sampling, camera noise, and the often ill-defined nature of the features for which we are looking. A second, more serious type of error is mismatches in correspondence. These can be either in the correspondence between a possible shadow making object and a shadow or between features on a shadow making object and features on a shadow. Errors in correspondence can lead to huge errors in locating a light source. Finally, there may be errors in *calculating* the location of a light source even if no errors have occurred on correspondence or in locating features in the image. Calculation errors result from limited precision mathematics. Calculation errors are the least significant of the errors we are likely to encounter and will not be addressed further.

Locate with high precision a point feature such as a corner on a shadow boundary is often extremely difficult. For instance, given a circular light source, the shadow of the corner of a square will be rounded as part of the penumbra. Also, the intensity of the illumination falling at the outer edge of the penumbra of a shadow corner approaches the intensity outside the the shadow. Consequently, the difference in illumination across a shadow boundary approaches zero and definitely falls within the level of camera noise. Even if a shadow boundary provides a sharp change in intensity in an image, edge detection algorithms often suffer from difficulties in localizing edges [Berzins 1984] [Canny 1986]. Finally, the discrete nature of CCD cameras ultimately limits the accuracy with which any feature can be located in an image.

Because the calculation of the location of a light source is an example of triangulation, the solution is particularly sensitive to certain kinds of errors in the data. In particular, if two illumination rays (or illumination planes) are nearly parallel then small errors in their description will produce large errors in the location of their intersection. This is a real concern because in practice the size of a shadow in an image is often small compared to the distance from the camera to the light source.

The solution to dealing with low accuracy in locating features in an image, is to use a large number of features. Since there are only so many features in a single image that we can use, we must rely on features found in a number of images. We must be careful, however, that the errors in locating the light source incurred from the errors in the image features tend to cancel out across a large number of features. In particular, we would like to find illumination planes in images that are taken from distant parts of an environment.

However, we cannot always depend on an environment to remain stable while we move in it nor that we will be able to travel widely in a given terrain. Consequently, we would like to make an estimate of where a light source is from as little as a single image, but to improve our estimate if we can acquire more data. Consequently, we look at bounding the location of a light source and tightening those bounds if more data is available.

Because we intend to use the shadow probe as an aid in determining the location of a light source, the shadow correspondence problem is greatly simplified. The probe is the shadow making object and finding the shadow of the probe has been discussed in Chapter 3. The problem still remains of determining the probe shadow's shape and finding a correspondence between this shape and the shape of the probe. When the shadow cast by the square probe is a quadrilateral and the shadow of the probe arm is also visible, then the correspondence with the probe is easily accomplished. The shadow of the probe arm uniquely identifies one side of the probe shadow and consequently the corners of the probe shadow can be put into their correct correspondence with the corners of the probe.

However, the shadow of a square need not be a quadrilateral. As is shown in Figure 1.1, the shadow of a square cast by a linear light source can be a hexagon. Because of the great range of shapes possible even for the shadow of a square, we have decided not to try to bring individual feature points on the shadow of the probe into correspondence with the the probe except in the case where the shadow is successfully determined to be a quadrilateral. Therefore, we need a more general description of the location of the shadow that will still provide enough information to determine the location of a light source. We also need to take into account the errors in locating the boundaries of a shadows umbra and penumbra.

From Figure 4.4 it is clear that by under-estimating the size of a shadow's umbra while over-estimating the size of the probe we are still guaranteed of having the location of the light source bounded by the umbra illumination rays. Similarly, over-estimating the size of a shadow's penumbra while under-estimating the size of the probe also still guarantees that the penumbra illumination rays bound the location of the light source. For the image of a shadow in the three dimensional world, it is important to under-estimate the size of the umbra and over-estimate the size of the penumbra perpendicular to the direction of the light source projected onto the image plane. This direction is the direction from the image of the shadow towards the image of the probe. If the probe is not visible, its projection onto the image plane can be calculated since the observing agent knows where the shadow probe is.

Under and over estimation of the size of a shadow can be done by a fixed amount, by

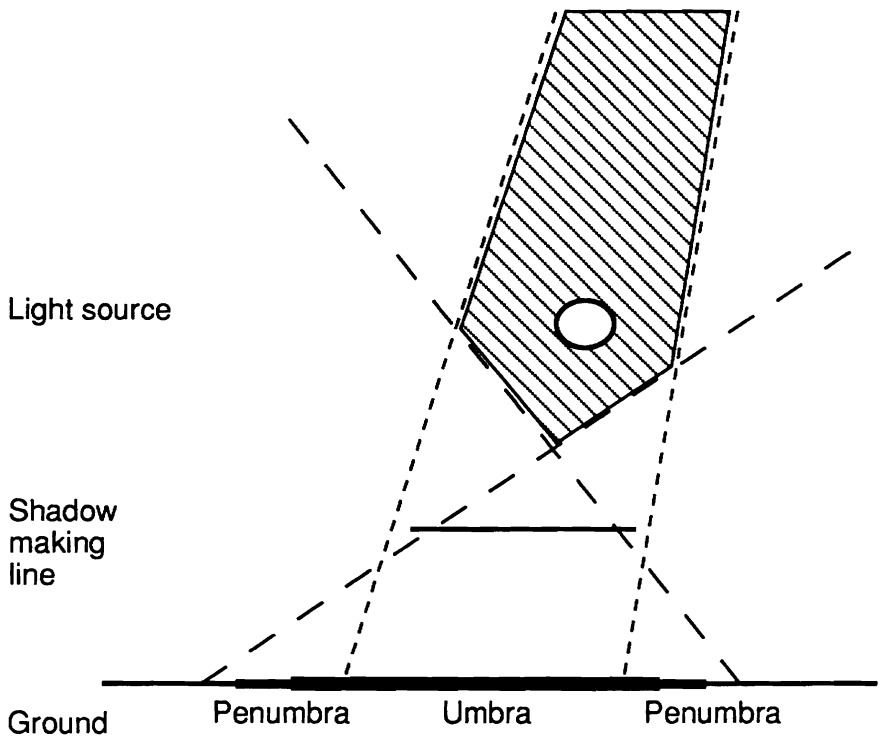

Figure 4.4: Adjusting the location of the shadow boundaries to take into account data errors while still maintaining a bound on the location of the light source. The light source is necessarily located within the shaded region.

an amount relative to the size of the shadow, or based on knowledge of the nature of the errors. What method for setting the amount of under and over estimation will prove most useful will need to be determined by experimentation. Under and over estimating the size of the shadow probe should be done based on the expected errors in the positioning system of the probe. It is assumed that the actual size of the probe is well known.

4.1.5 Strategy Review

In order to locate a light source, an agent places its shadow probe out into the environment and then locates the probe's shadow (umbra and penumbra) in an image of the scene. This part of the process has been discussed in Chapter 3.

Here we discuss how to define bounds on the location of the light source. The bounds from each image position are defined in two sets. First we describe bounds defined by lines in the image plane. We will refer to these as *image bounds.* Later we define bounds based on the probe's position perpendicular to the image plane. We will refer to these as *depth bounds.*

Before continuing, some notation: small letters in italics indicate points $(a, \ldots, z)$, capital letters in italics indicate lines $(A, \ldots, Z)$, bold capital letters indicate planes $(\mathbf{A}, \ldots, \mathbf{Z})$.

First we define the image bounds on the location of the light source. Figure 4.5 provides a schematic of many of the elements necessary to define the image bounds. Initially, a line L is found through the center of mass of the shadow image c_s and the center of mass of the projection of the probe onto the image plane c_p. Let S be a line through c_s perpendicular to L. The orthographic projection of the shadow umbra and penumbra onto S is found. The projection of the umbra is then under-estimated and the projection of penumbra is over-estimated along S. Let P be a line through c_p perpendicular to L. The orthographic projection of the probe onto P is found. An under-estimation and over-estimation of the probe size is made along P. We now use our estimate of the shadow and probe size to define umbra and penumbra illumination rays as in Figure 4.4. Umbra and penumbra bounding planes are defined as passing through the umbra and penumbra illumination rays respectively and the focal point of the camera. Together the umbra and penumbra bounding planes define the image bounds on the location of the light source. The light source must lie within the intersection of the volume between the umbra bounding planes $(\mathbf{U_1}, \mathbf{U_2})$ and the volume between the penumbra bounding planes $(\mathbf{X_1}, \mathbf{X_2})$.

Next, we define the depth bounds on the location of the light source. Figure 4.6 provides a schematic of these depth bounds. We define two planes $(\mathbf{F}, \mathbf{N})$. $\mathbf{F}$ will bound how far

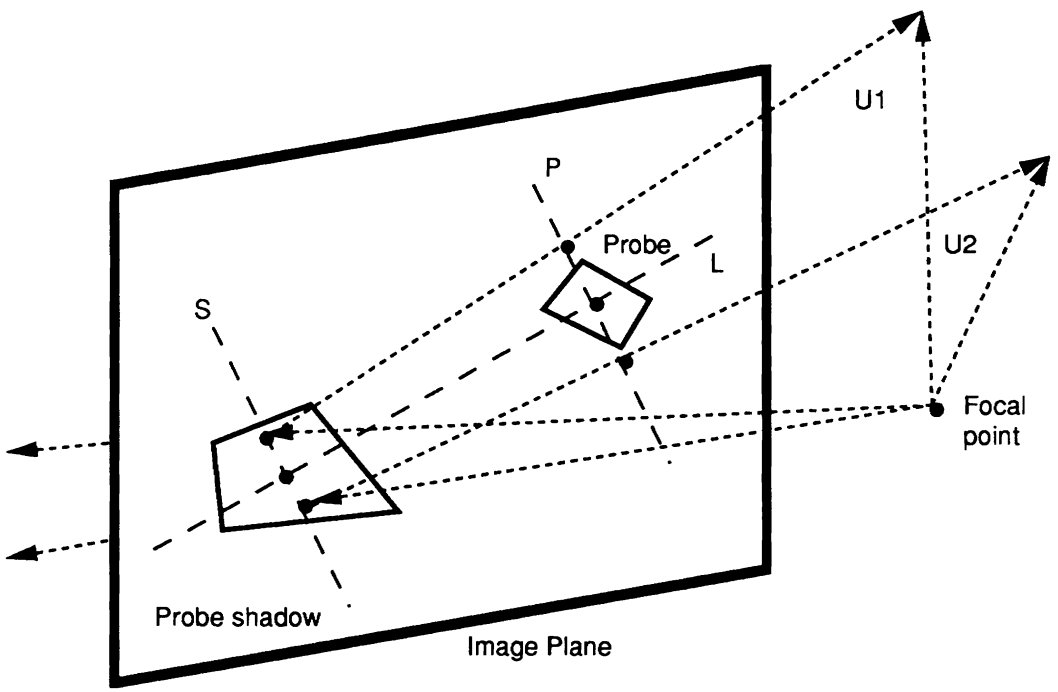

Figure 4.5: The umbra image bounds on the location of the light source in 3D. The light source must lie in the volume bounded by the two plane $U1$ and $U2$ that also includes the line L on the image plane. The umbra image planes extend to the right and away from the focal point of the camera. The penumbra image bounds are defined similarly. See the text for a further explanation.

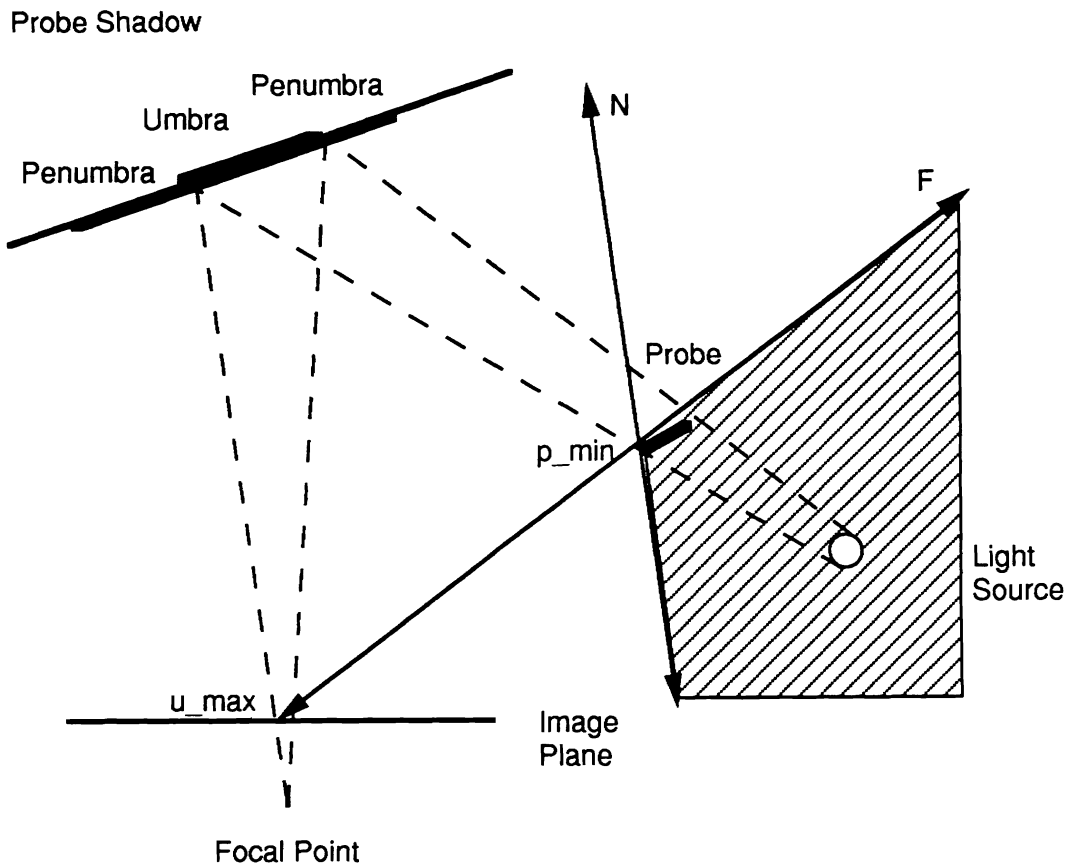

Figure 4.6: The depth bounds on the location of the light source in 2D. The light source must lie in the shaded region defined by the lines F and N. See the text for an explanation.

in depth the light source can be. **N** will bound how near in depth the light source can be. Each plane will go through a the point on the probe p_{min} closest to the visual cone defined by the probe's shadow. In addition, the intersection of both **F** and **N** with the image plane will be perpendicular to L. **F** will be defined to go through a point u_{max} on line L that is the maximum distance of a point on the boundary of the shadow umbra from the probe (c_p). **N** will be parallel to the line through the focal point of the camera and point u_{max}.

The intersection of the volumes enclosed by the image bounds and the depth bounds define the area where the light source is located. As the observer moves and does more experiments with the shadow probe new bounds on the location of the light source can be found and these can be intersected with the previous bounds. In order to combine the results from multiple experiments it is important that the observer know its relative

motion. Errors in the estimate of the relative motion can be somewhat compensated for by further loosening the bounds on the location of the light source for each experiment before combining the results from multiple experiments.

4.2 Locating Where a Shadow is Cast

Besides determining the location of a light source, we want to be able to determine the location of the surface or surfaces on which a shadow is cast. As with finding the location of a light source, triangulation is the method for finding the location of the surface on which a shadow is cast. The basic idea is to look at the intersection of a line from the light source through the shadow making object and a line from the focal point through the image of the shadow. The intersection of these two lines gives the location of the shadow in 3-space. See Figure 4.7. We know the location of the shadow making object since it is the shadow probe. We can determine bounds on the location of the light source as described above. We have already discussed locating the probe's shadow in an image. The difficult part of locating the shadow in the world, is taking into account the limitations in our knowledge about the location of the light source, probe, and shadow image. As with locating the light source, we will depend on bounds to limit the area in 3-space in which the shadow can lie.

The triangulation to locate where a shadow is cast depends on finding the intersection of two cones with quadrilateral cross-sections. The first cone is the *shadow visual cone.* This is defined with its apex at the focal point of the camera and one cross-section defined by a bounding box around the outer edge of the shadow in the image. The *illumination cone* has the shadow probe as one cross-section. We define the four planes that bound the illumination cone from the shadow probe cross-section in the following way. Consider a plane $\mathbf{P_1}$ that initially aligns with the plane of the shadow probe. Fix one side of the shadow probe as an axis for $\mathbf{P_1}$. Rotate $\mathbf{P_1}$ away from the shadow probe and towards the location of the light source. Let $\mathbf{P_1}$ come to rest when it first contacts the polyhedra that bounds the light source location. Note that this is akin to one step of the package wrapping algorithm for the convex hull in 3-space [Sedgewick 1983]. A plane will first contact a polyhedra at one or more vertices. Use one of these vertices and the axis through one side of the probe to define one plane of the illumination cone. Do this for each side of the shadow probe in order to fully define the illumination cone.

In conclusion we take as bounds on the location in 3-space of the shadow the intersection of the shadow visual cone and the illumination cone. If the observer can move while holding

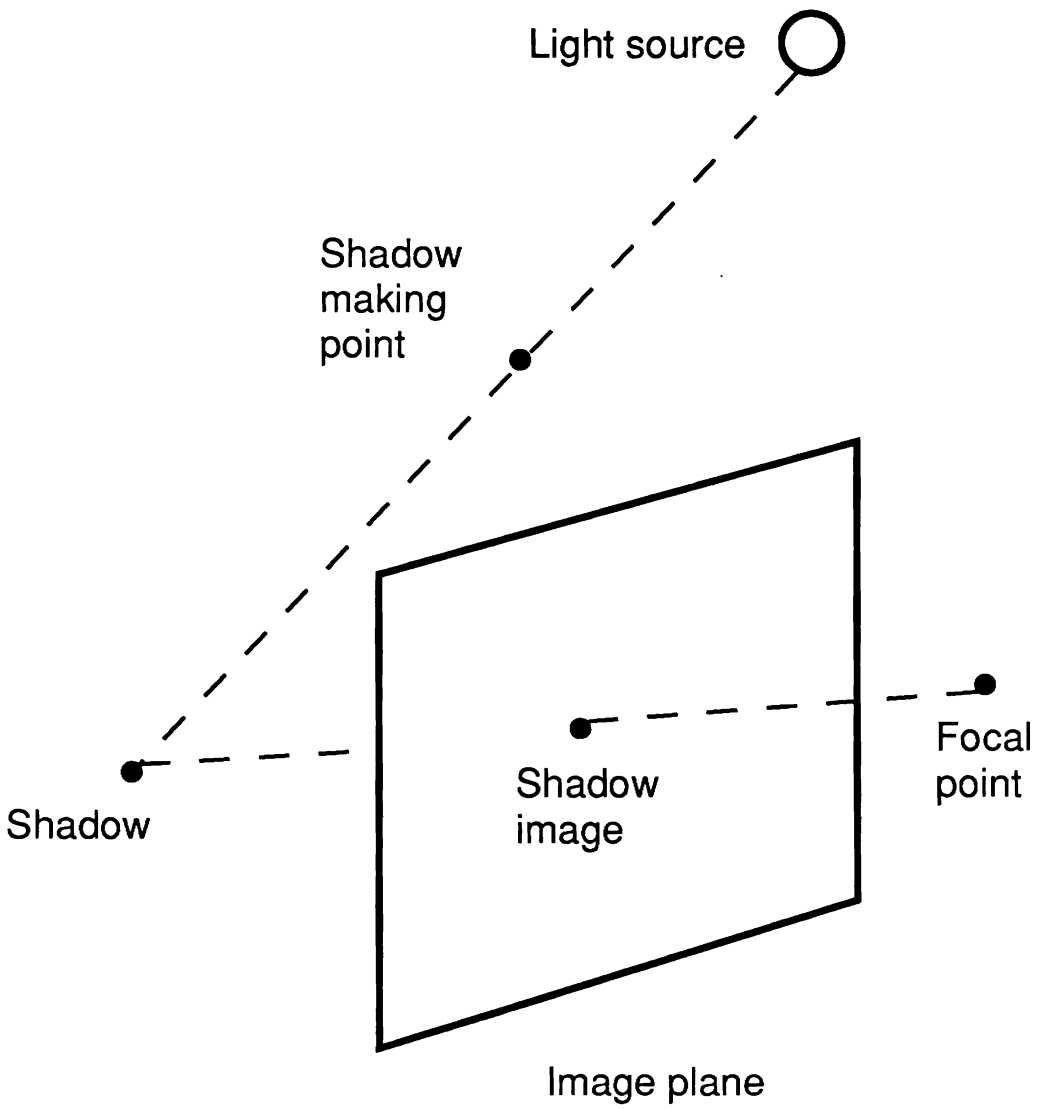

Figure 4.7: Locating a shadow point in 3-space based on the image of the shadow point and the location of the light source and the location of the shadow making point.

the shadow probe stationary relative to the environment then the location of the light source can be further constrained by new shadow visual cones. Alternatively, the shadow probe can be moved while the observer is stationary in order to map out the location of a surface or surfaces on which the shadow is cast.

4.3 Conclusion

In this chapter we have presented algorithms for determining bounds on the location in 3-space of a light source and on the location in 3-space of a shadow. We have done this in a way that takes into account errors in our measurements. We have also taken special advantage of the benefits of having an active observer. As the observer moves in its environment its estimates of the location of the light source and scene objects can be augmented.

At present the algorithms mentioned in this chapter have not been implemented. We propose to implement and test these algorithms on images taken as a mobile camera is moved through our lab.

CHAPTER 5

Shadow Candidates and Color

The shadow cast by the probe provides the agent with a sample of a single location in space directly lit and in shadow. The agent can move the probe in order to get multiple such samples. But it is rarely practical and often impossible to cast a shadow into all parts of a scene. Consequently, we need a strategy for analyzing the surfaces in an image not effected by the probe's shadow.

In this chapter we investigate the use of color to analyze shadows. As we saw in Section 1.4.2, distinguishing shadows from other reflection factors in general is very difficult. What we propose here is to segment an image into regions such that if shadows are present, a uniformly colored surface directly lit and in shadow is very likely to be represented by a single region or that a cross section of the penumbra of such a shadow will be represented as a single region. Some of the segmented regions will be shadow candidate regions. The shadow candidate regions will be further investigated for evidence to support or refute the hypothesis that a shadow is present. In the latter part of this chapter a further use of color will be made to analyze the shadow candidate regions. In the next chapter the other shadow cues will be sought for the shadow candidate regions.

In Section 1.4.1 we showed that the light measured by the camera from a single surface material lit and in shadow is a line in color space if other reflectance factors do not apply. See Equation 1.2. However, shading, strong local inter-reflections, and other illumination effects complicate the detection of shadows. Consequently, we will make the following assumption:

> **The Linear Color Cluster Assumption for Penumbrae:**
> We assume that the light irradiating a penumbra, with the exception of the the partially obstructed light, does not vary or varies insignificantly.

Consequently, the variation in reflection in a penumbra on a uniformly colored surface is

due entirely to the obstruction of a direct light source. And, in order to find regions of an image that could represent the same surface lit and in shadow across a penumbra we present a scheme in which an image is segmented into line-like or uniform color clusters.

Note that the observer can test the validity of the Linear Color Cluster Assumption for Penumbrae for some shadows in its environment: namely those that it casts with its probe. We will assume that the conditions that hold for the probe's shadow will apply for other shadows in the scene. As the observer explores its environment and examines more shadows cast by its probe, this test of the Linear Color Cluster Assumption becomes more sound.

However under our assumption, line-like color clusters can still originate from physical phenomenon other than shadows. Shading, inter-reflections, highlights, or material changes may also produce line-like color clusters [Lee 1991]. As has been discussed, multiple cues are necessary before a shadow can be recognized with any confidence – color alone is not sufficient.

The analysis of Section 1.4 was done strictly in color space and ignores image or scene locality. Because all the pixels in an image of a complicated scene taken together may result in many line-like color clusters, we introduce local image continuity as a constraint in our color image segmentation. So, we will only be looking for contiguous sets of pixels in an image that form line-like color clusters. This restricts our image interpretation at this point to those shadows for which the same surface can be seen directly lit and adjacently in shadow.

5.1 Color Image Segmentation

Our color image segmentation is founded on three ideas. First, to use line-like color models to take into account shadow candidate regions. Second, to dove-tail the processing between color-space and image-space in order to take into account aspects of each. And finally, the realization that segmentation should be the search for the best description of an image in terms of primitive models [Leonardis et al. 1990].

The image segmentation begins by finding strong color samples in the image. This is accomplished by examining the histogram of the image pixels in color space. We have used a two dimensional color space for the histogramming. The 2D color spaces used include $(\frac{\sin\lambda}{intensity}, \frac{\cos\lambda}{intensity})$ where λ ranges over the visible wavelengths [Lee 1991] and the 2D color space $(\frac{green}{red}, \frac{blue}{red})$. We believe that any 2D color space that tends to de-emphasize intensity is a suitable choice for the initial histogramming. The two 2D color spaces mentioned were

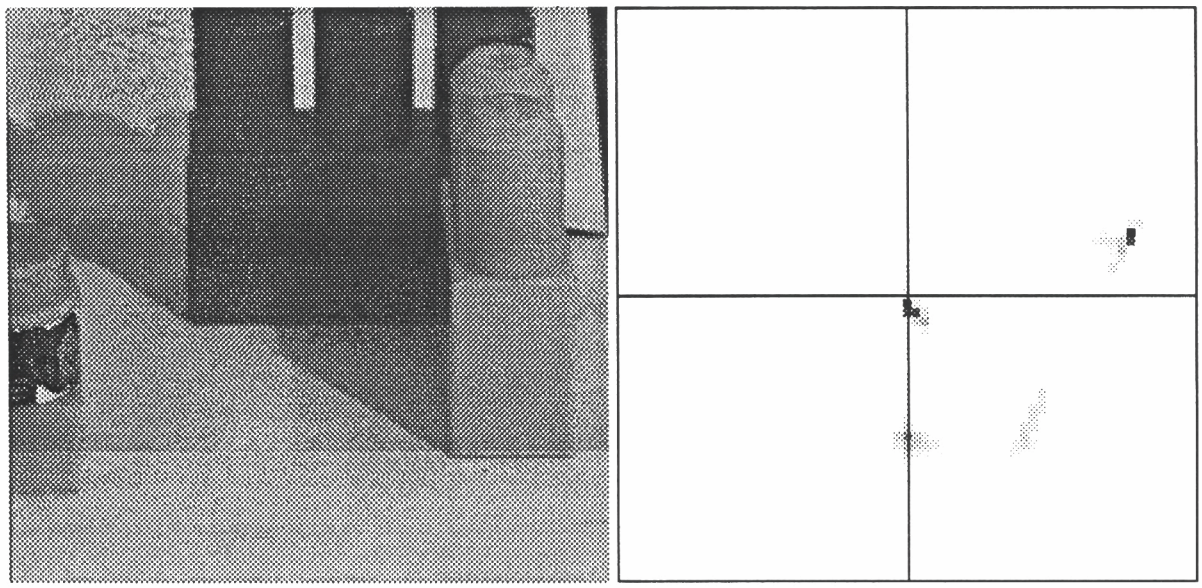

Figure 5.1: Color Histogram. At left is the image. This is the same image as used in Figure 3.4. At the right is a color histogram of the image with the exception of the portion of the image explained by the probe in Figure 3.4. The center of the 2D histogram is the origin of the coordinate system $(\frac{\sin\lambda}{intensity}, \frac{\cos\lambda}{intensity})$. Unsaturated colors are near the origin and saturated colors are at the periphery. Red is to the right, green at the bottom, and blue is at the upper left corner. Strong responses can be seen for the white background (the spot near the center of the histogram) and for the red block (the spot near the right of the histogram). There is a weaker response for the green block not in shadow (the spot at the bottom center of the histogram). The green block in shadow is a line from green to red in the histogram.

convenient for us to implement. See Figure 5.1 for a sample 2D color histogram of an image.

In the 2D color histogram we look for strong peaks. Strong peaks in the histogram will correspond to dominant colors in the original image. From the dominant colors we will try to find, through the segmentation described below, distinctly colored regions in the image consistent with possible shadows. The strong peaks in the 2D color histogram $H(c_1, c_2)$ are found by the following algorithm:

1. LOOP

2. Find the maximum value $H(c_{1max}, c_{2max})$ in the histogram.

3. IF $H(c_{1max}, c_{2max}) < dist * the\_previous\_peak$ THEN EXIT the LOOP.
4. In the histogram find all the elements adjacent to (c_{1max}, c_{2max}) with a value $\geq thresh * H(c_{1max}, c_{2max})$.
5. Record all the found values as a peak.
6. Delete the peak from the histogram.
7. END LOOP

The *dist* parameter was introduced as part of our strategy to inter-leave the analysis between color and image space. By having the *dist* criteria we can find just a few strong peaks in color space, then go back to the image and try to explain parts of the image. If the image cannot be fully explained then we histogram the unexplained parts of the image and again look for peaks in the 2D color space. This enables us to explain those areas of an image for which the color information is weak or ambiguous only after the more more uniformly colored portions have been explained. This is important because the segmentation algorithm involves growing regions with a tolerance based on the variance of the color peak. Consequently, regions of high tolerance can grow easily unless we stop the growth at portions of the image for which we already have a good description with a lower tolerance.

The *dist* parameter was set to 0.5 for our experiments. The *thresh* parameter was set at 0.5 for all the experiments we have done so far. In addition, in our experiments the results were not found to vary for *thresh* values of between 0.7 and 0.3. Both the *dist* and *thresh* parameters depend on the noise in the image and on the non-uniformity in color of the scene. For large amounts of noise or scene variability they both should be set higher. Note that setting the parameters does not depend on an estimate of image noise independent of image variance.

The peaks found in the color histogram are used to find seed regions for the image segmentation. Each peak is used to label pixels in the original image with the peak color. Then, contiguous pixels with the same color label are taken as seed regions for segmentation. The segmentation follows the algorithm of [Leonardis et al. 1990] in which each iteration of the algorithm consists of the following steps:

1. WHILE change DO LOOP
2. Grow current regions based on extrapolating the region model and the use of a tolerance criteria.
3. Update each region to fit the new and old data.
4. Prune away regions based on the overlap, size, model goodness of fit, and model order.
5. END LOOP

Unlike [Leonardis et al. 1990] our region models are not bivariate polynomials functions of pixel location $P(x, y)$. Instead, our region models are uniform or linear functions in color space (q_1, q_2, q_3): either

$$[a_1, a_2, a_3]^T \tag{5.1}$$

or

$$\begin{bmatrix} a_1 \\ a_2 \\ a_3 \end{bmatrix} + \alpha \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix}. \tag{5.2}$$

The seed regions are all uniform. If a uniform region does not grow during an iteration then a linear model is tried. The linear model is accepted if the error is relatively small and the region can be grown by a considerable amount. See [Leonardis et al. 1990] for details of how the models are updated and the regions are pruned.

The tolerance criteria for region growing is determined by the variance of the peak found in the color histogram and the tolerance is allowed to vary uniformly with pixel intensity. The latter is necessary because greater color variation is possible for brighter image pixels.

See Figures 5.2 through 5.4 for results of the color image segmentation algorithm.

Our color image segmentation decouples the region models from the individual pixels. The models apply to region pixels en masse. This simplifies the models and hence the algorithm but also limits the model's expressiveness. However, the decoupling insures that a a single material illuminated by one light source on which numerous distinct shadows are cast is still recognized as one image region. Similarly, a single material illuminated by one light source with a complex shading pattern will also be recognized as a single image region. The post-segmentation processing will concentrate on analyzing individual regions and will not need to compare regions that are not adjacent.

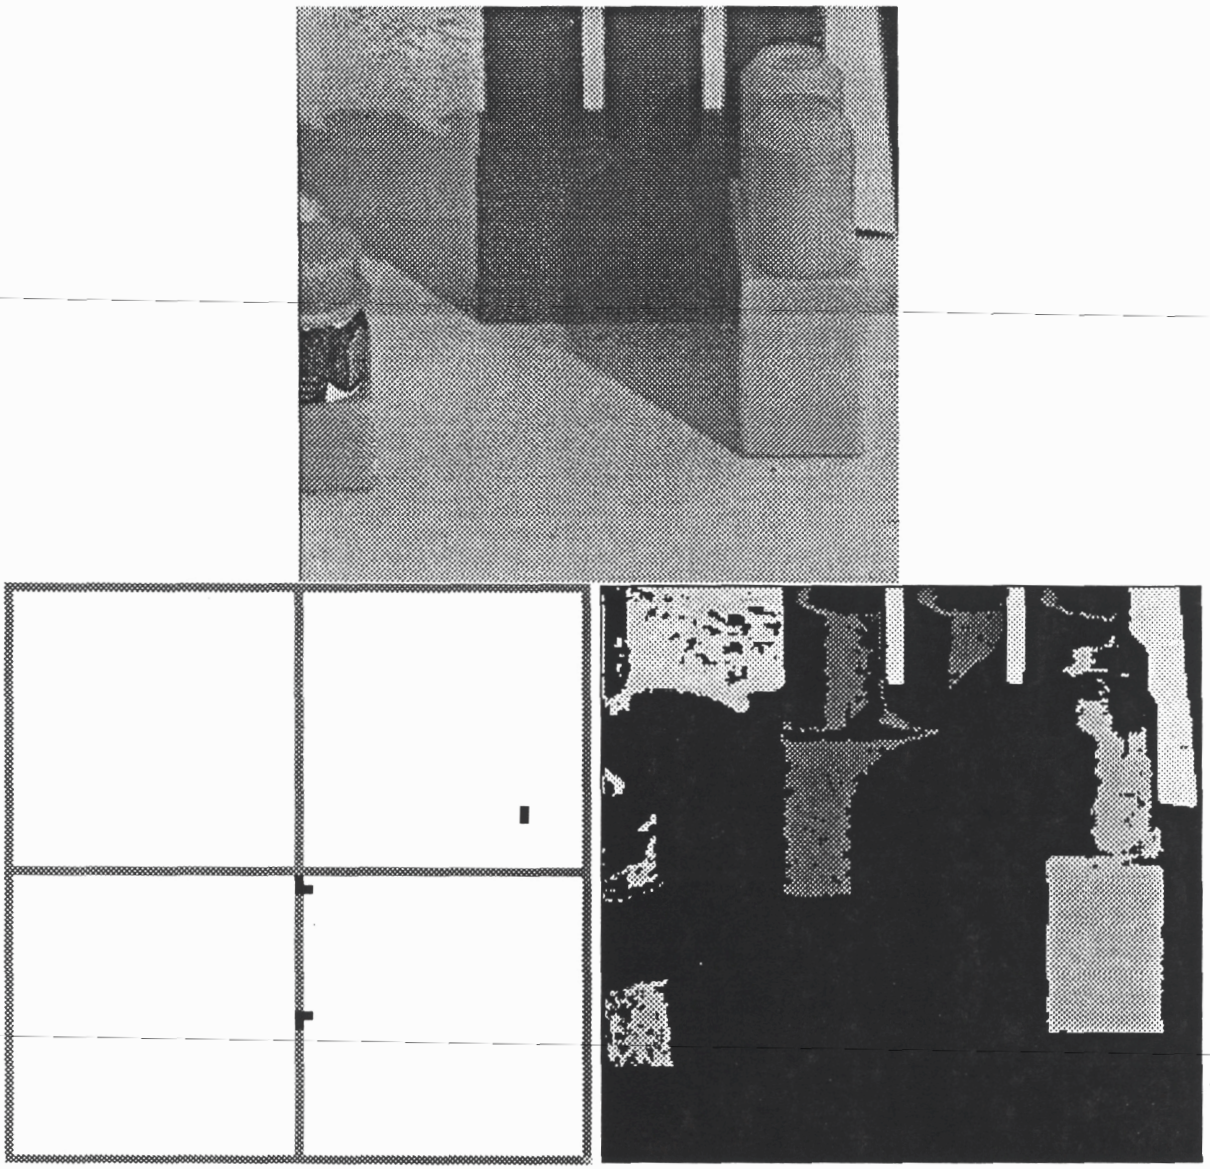

Figure 5.2: Color Image Segmentation: Scene 1. Top : The original image. Bottom Left: Strong peaks found in the color histogram of the image. Compare this figure with Figure 5.1. Bottom Right: The seed regions as found from the strong peaks in the color histogram.

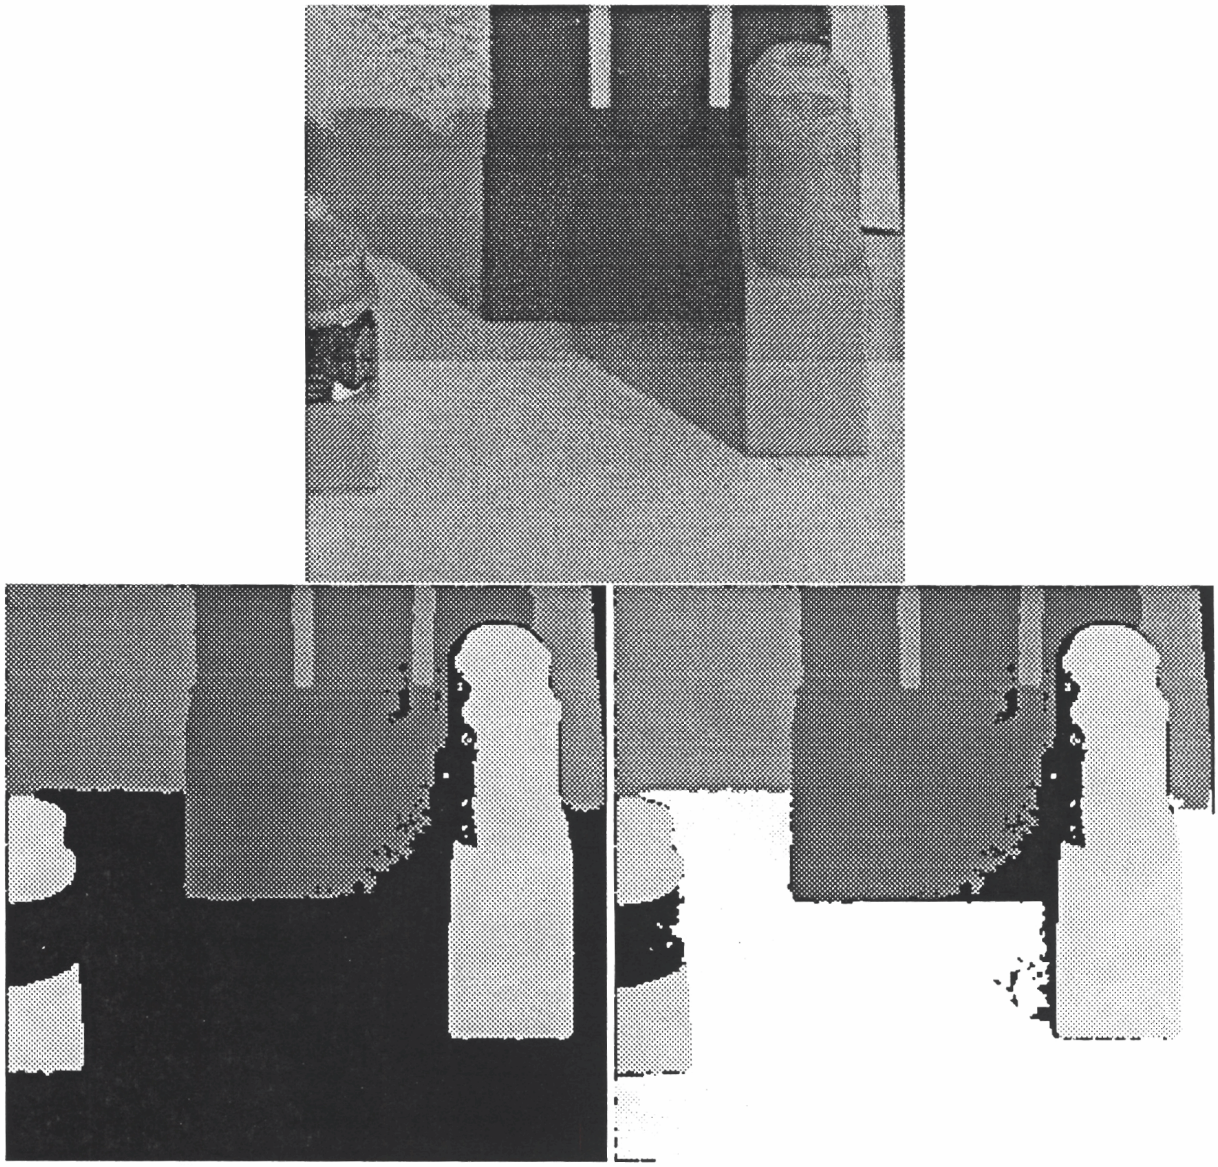

Figure 5.3: Color Image Segmentation: Scene 1 continued. Top : The original image. Bottom Left: The segmentation of the image not including the region found directly from the probe's shadow (see Figure 3.4). Bottom Right: The full description of the major regions of the image. Note that strong reflections from the blocks have made some pixels difficult to label.

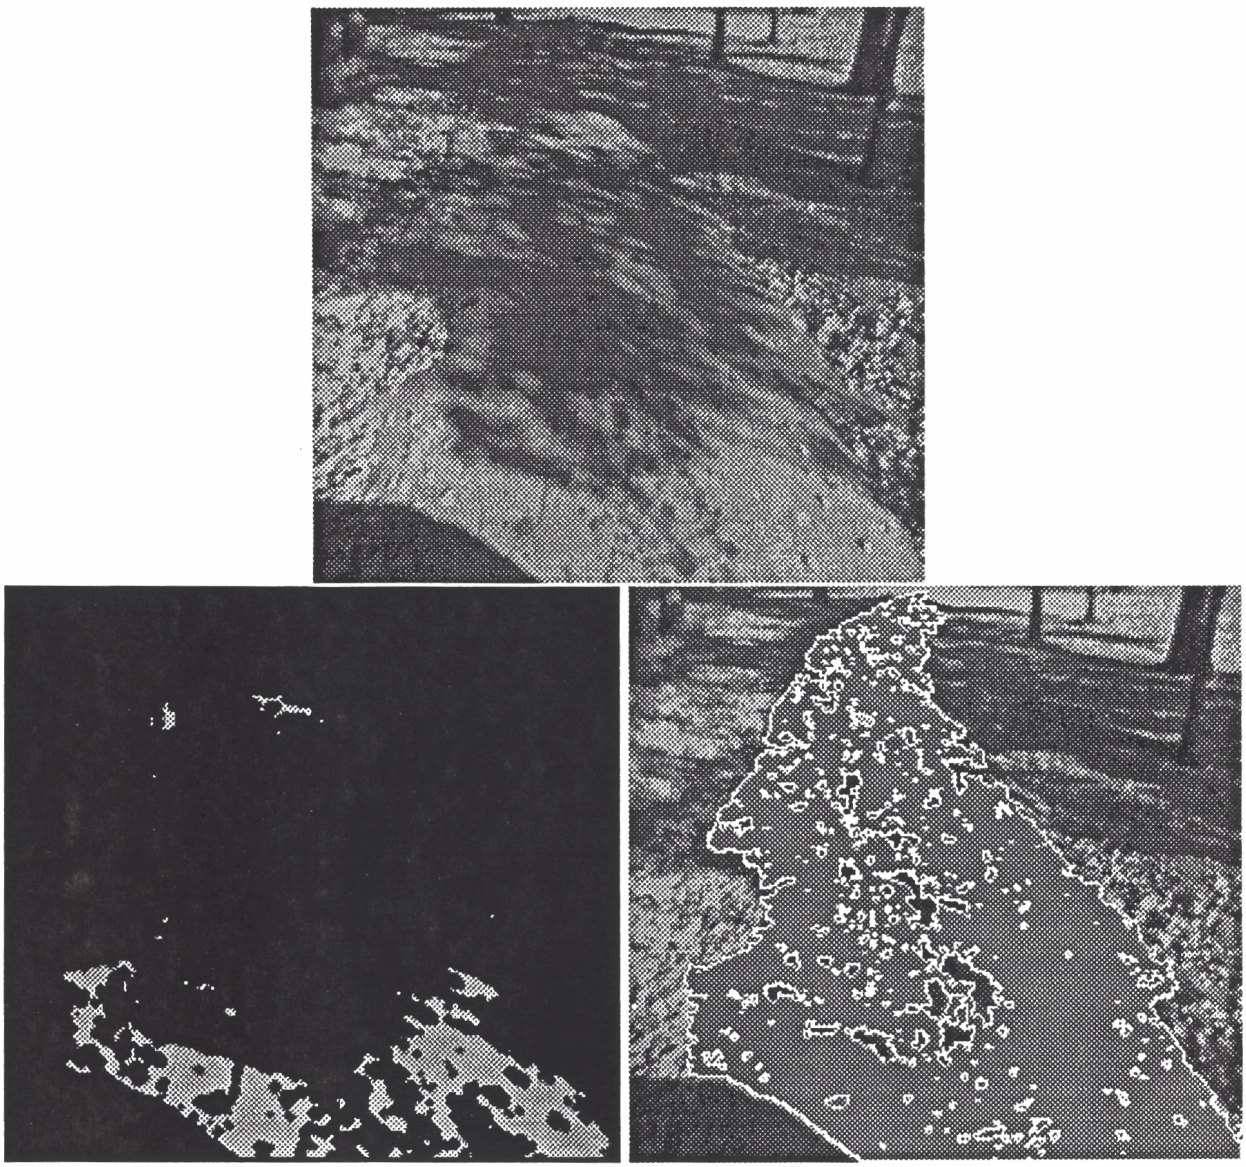

Figure 5.4: Color Image Segmentation: Scene 2. Top: An image taken in the autumn of a road running under some trees, Bottom Left: The seed region manually chosen from the color histogram. Bottom Right: The results of running the color segmentation algorithm super-imposed on the original image. Note that the leaves and oil spots on the road have not been labeled as part of the road material. Image courtesy of the Carnegie Mellon University Navlab project.

5.2 Color Region Analysis

There is a frequently encountered scenario in which a uniformly colored surface directly lit and in shadow will correspond to a single line-like color cluster. This is the scenario where there is one strong extended direct source of illumination and inter-reflections are a local phenomenon: any sunny day in an environment without strongly specular materials. Based on our probe's shadow we can determine that such a scenario holds and make use of that in our color shadow analysis.

For the sunny day scenario, if a segmented image region has a linear rather than a uniform model then it could be one material seen lit and in shadow. Image regions such as these become our shadow candidate regions. If a region does contain a shadow then the umbra of the shadow will correspond to those pixels in the region whose intensity is lowest. The fully lit portion of the region will correspond to those pixels in the region whose intensity is highest. For a shadow, pixels in a region falling between the extremes of intensity will belong to the penumbra.

However, in the typical indoor environment, rooms are illuminated by multiple light sources. In addition the color of these lights tends to vary. Outdoor lighting through a window is differently colored than incandescent lighting which is differently colored than fluorescent lighting. Often at least two of these light sources are present simultaneously in indoor environments. Consequently, we must expect complex shading of even uniformly colored surfaces. Consequently, one surface lit and in shadow may not be represented as a single line-like color cluster. However, we still assume the linear color cluster assumption for penumbrae holds. So, across the width of a penumbra we expect a single image region from segmentation. However, along the perimeter of a penumbra we could have multiple regions due to various non-shadow illumination effects. Under our color segmentation algorithm we have a piece-wise linear representation of a penumbra under conditions of multiple differently colored illuminations (or varying surface albedo).

From what an observer learns about the illumination conditions in a scene from the casting his own shadow, the observer can determine whether shadows cast on a uniformly colored surface are likely to be found within a single image region or across several image regions. If shadows are expected strictly within image regions, then the grouping of neighboring regions together is not necessary in order to find a complete shadow cast on a single material.

5.2.1 Shadow Color Bias

Because line-like color clusters can result from physical events besides shadow, our observer must further analyze the shadow candidate regions found by our color segmentation procedure. Here we present further tests of the color clusters to support or refute their origin from shadows.

The simplest criteria for the linear color clusters follows from the fact that shadows are darker than the same surface directly lit. Consequently, for a color space whose bases are band-limited functions (such as red, green, blue), the color cluster for a shadow must not get brighter along any of the bases.

In addition, if all the shadows in a scene are illuminated by the same light, we expect the shadows to show a similar relative change in intensity, hue and saturation. We propose to judge any trends in the ambient light illuminating the shadows based on the results of the shadows cast by the probe. For instance, if all the probe shadows cast show a bias towards blue along a measure of hue then we will expect all shadows to follow this rule under the present lighting conditions. We propose to look for trends along the criteria of hue and saturation. We also propose to examine if the ratio of a surface in total shadow to the same surface directly lit for the probe shadows is constant. (See Section 3.1.1 for a discussion of this ratio.) If the ratio is bounded for a variety of surfaces in shadow, then the observer can use this ratio as a color criteria for detecting shadows. What the most practical means of measuring color trends in the probe's shadows will be is still to be determined by our further investigations.

5.3 Conclusion

In this chapter we have presented a general algorithm for segmenting color images into regions that form uniform or linear color clusters in color space. Based on this representation of an image we propose to begin our recognition of shadows. Shadows will be represented as piece-wise linear color clusters under our linear color cluster assumption for penumbrae. Under this assumption, the width of a penumbra will always be a single segmented region. Consequently, all linear color clusters become shadow candidate regions. Many of these regions can be discounted as shadows because they do not show a darkening simultaneously along each of red, green, and blue. Other regions can be discounted because they show color trends not compatible with the results from the probe's shadow. Additional cues will be brought to bear on the remaining shadow candidate regions in the next chapter.

CHAPTER 6

Shadow Candidates and Other Cues

In this chapter we discuss how the shadow cues other than color can be used to analyze the results of the color image segmentation discussed in the previous chapter. We do not propose to use texture continuation or an analysis of object and shadow silhouettes in our work but we do discuss how these cues could be incorporated into our system.

6.1 Some Object Must Cast the Shadow

In Chapter 4 we showed how to use shadows cast by an active observer to locate a source of illumination and to locate where the shadows were cast. Here we make use of this information and knowledge about a scene to discount the possibility that some regions of an image could be shadows. The key idea is to determine that no object lies between the light source and a portion of the scene visible in an image and hence that no object could be casting a shadow in that portion of the scene.

We assume that an image has been segmented into labeled regions, some of which may be shadows or contain shadows cast on a surface or surfaces. Consider a region labeled R. If R contains a shadow, then the projection onto the image plane of the shadow making object would lie between R and the location of the light source projected onto the image plane. Consequently, if all the image area between R and the image of the light source can be discounted as an obstruction to the light, then we know that R cannot contain a shadow. If the light source is not visible in the same image as R, we assume that the observer can pan in the known direction of the light source in order to acquire a sequence of images in which the area between R and the light source is visible. If multiple light source exist in a scene, then image area between R and each of the light sources must be examined.

We assume that the location of the light source is known (or that bounds are known on its location). We also assume that some parts of the image are known not to obstruct light from a source from reaching any other visible surface. This will always be true for the sky and is often true for the ground.

We will often assume that the ground can be recognized in a scene. Many observers must be supported by some surface such as the ground, And they are likely to be capable of examining their support in order to recognize it in images of their surroundings. Alternatively, many autonomous vehicles are provided with information about the appearance of roads over which they can move (see for example [Funka-Lea and Bajcsy 1992]).

If the image area between a candidate shadow region R and the projection of the light source can be completely explained by the sky, ground, and perhaps other scene elements which could not cast a shadow, then R cannot be a shadow. However, if any image area between R and the light source remains unaccounted for then R may contain a shadow.

6.2 Shadows as the Projection of a Silhouette

The shape of a shadow will be the projection of a silhouette of an object obstructing the light emitted from a light source. This is an obvious cue to the presence of a shadow and has been used in a simplistic way in [Lowe and Binford 1985]. In their work an attempt is made to put the corners in a line drawing into correspondence (see Figure 2.1). However, as discussed in Section 1.3, this is generally a difficult cue to use because of the possible complexity of the correspondence between a shadow and the shadow making object. It is important to know when as well as how to try to use this cue to help in the recognition of shadows.

The image of the candidates for the shadow making object that casts a particular shadow must lie between the image of the shadow and the projection of the light source onto the image plane. The recognition of such candidates was discussed in Section 6.1. So, we may assume that some candidates for shadow making object have been found and the problem is now one of attempting to find a correspondence between the candidate shadow making objects and the shadow.

If we do not know the complete three-dimensional shape of an object, the only information about silhouettes of the object available to an observer are those silhouettes that are directly observed. However, the observed silhouette can only be casting a shadow in the same image if the light source is behind the silhouette. Consequently, correspondence

between an object silhouette and a shadow should not be attempted unless this geometry is confirmed. Earlier we showed how to use the shadow probe to locate a light source. In addition, the nature of the correspondence cannot be determined unless the shape of the light source is known. In our work the shape of a light source is approximated by tightly constraining the location of the light source. Only for a point light source will the properties of perspective projection govern the correspondence between the shadow making object and the shadow. For instance, the shadow of a conic silhouette can only be guaranteed to be a conic for a point light source. Once perspective projection has been found to hold, invariant descriptors [Duda and Hart 1973] [Forsyth et al. 1991] can be used to test correspondence or the corner matching technique of [Lowe and Binford 1985] can be used.

We do not propose to use this cue in our system, only to recognize when it may be applicable.

6.3 Penumbra and Umbra Structure

In a scene with an extended light source, shadows can be expected to have a penumbra and umbra structure. Using the shadow of the probe and the techniques of Chapter 4 we can determine if a light source is well approximated by a point source.

In [Jiang and Ward 1992] a detected penumbra and umbra structure is used as a cue to detect shadows. However, they detect penumbrae based on relative image brightness while we base our detection on a linear model in color space discussed in Section 1.4.1 and Chapter 5. In addition, if we assume that the objects that cast shadows contain no holes and do not overlap relative to the light source, and that the shadows are cast onto contiguous locations in space, then the shadows will have a penumbra darkening towards a central point, line, or umbra. The case of a point or line apply when there is no umbra, only a penumbra.

Objects with holes can have shadows with very complex shading patterns. Consider, for instance, a series of irregular wire meshes between a light source and a uniformly colored surface on which a shadow is cast. The wires nearest the surface will cast shadows with relatively narrow penumbrae while the wires farthest from the surface will cast shadows with wide penumbrae. Because of the irregular nature of the meshing, the amount of obstruction in the shadow will vary in a complex but smooth (differentiable) way. The situation is further complicated if the different levels of the meshing are allowed to move relative to each other. In this case, there may be no stable shading pattern in the shadow. This scenario may seem contrived but it is not too far from the reality of shadows beneath trees

on a sunny, windy day.

Shadows that are cast onto non-contiguous surfaces, such as a table and the floor, will not necessarily have their umbra surrounded by a penumbra. At a geometric discontinuity, say the edge of the table-top, the umbra of a shadow may appear to be at the contour of the shadow.

Because of the potential complexity of the structure of a shadow, we propose to use the penumbra-umbra structure as a limited cue for shadows. If we find that a potential shadow has a compact umbra completely surrounded by penumbra then we will consider this evidence in favor of a shadow identification. If such a penumbra-umbra structure is not found, this is considered inconclusive evidence and no judgment is made.

6.4 Shadows as Apparent Surface Marks

As an observer moves in a fixed environment, shadows remain stationary relative to the surfaces on which they are cast. As such, shadows resemble surface marks. We propose to test this shadow cue for only simple surface geometries. Namely, we will only test that the boundary of a shadow cast on a plane lies in that plane.

We plan to test for planarity based on a small set of feature points found in two images. Whether or not a set of five scene points lies on a plane can be determined by imaging the points from two distinct views. This follows from the invariance to perspective projection of what [Duda and Hart 1973] calls *two-dimensional projective coordinates.* Two-dimensional projective coordinates are basically an extension of the *cross-ratio* of four points on a line. The cross ratio is also invariant under perspective projection. See Appendix B for the definition of two-dimensional projective coordinates and the cross-ratio. The important result is that we can test whether or not a small set of points seen in two images comes from a planar surface without having to recover the parameters of the plane.

When a sufficient number of feature points are available, we propose to first test that the image region hypothesized to be the surface directly lit is tested for planarity. For this we need five feature points on the lit region. If the lit surface is not planar, then we do not test for co-planarity of the shadow. If the lit surface is planar, we test two feature points on the hypothesized shadow region with three points on the lit surface for planarity. If these five points are planar, we record this as further evidence of a shadow. If the five points are not planar, we record this as evidence against a shadow.

If five points are not available on the lit region, but three are, then we proceed directly

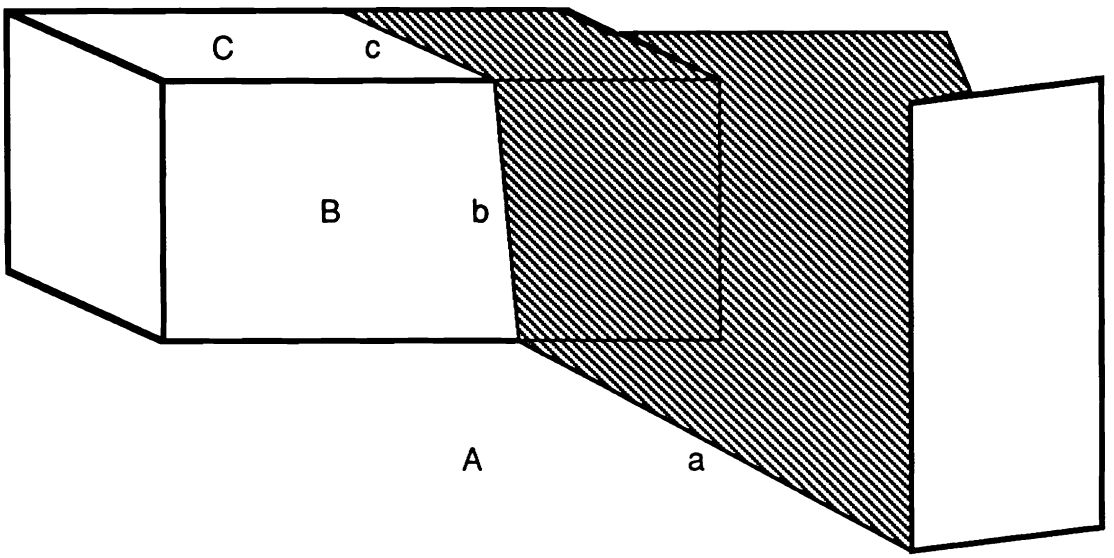

Figure 6.1: A shadow of a rectangle is cast on three different surfaces: A, B, and C. Three segments of the shadow contour are labeled: a, b, and c. See the text for a discussion.

to test for co-planarity across the shadow boundary. In this case, however, we note that the evidence for a shadow is weaker because we do not know if the underlying surface is really planar and hence if the assumption of our test is valid. If the five points are not planar, we still record this as evidence against a shadow although the evidence is again weaker.

6.5 Shadows and Surface Discontinuities

If the surfaces within a scene are not oriented relative to the direction of illumination, then shadow boundaries will change direction as they cross surface discontinuities. See Figure 1.2 or 6.1. Consequently, if we know a surface discontinuity exists in a scene then we can expect a shadow cast across this discontinuity to show a change in direction in the image at the surface discontinuity.

If we know the geometry of the surfaces involved and their orientation relative to the direction of obstructed illumination, then we can determine quantitatively what the change in the shadow boundary should look like. For instance, consider the scenario in Figure 6.1. Assume orthographic projection with a viewer centered coordinate system with z giving the distance from the viewer and x and y are in the image plane. Let $G_C = (p_C, q_C)$ be the orientation of surface C. Let $V_c = (\Delta x_c, \Delta y_c, \Delta z_c)$ be the direction of the shadow boundary

c on surface C. The slope of the image of c is $\frac{\Delta x_c}{\Delta y_c}$. Let $G_I = (p_I, q_I)$ be the orientation of the illumination plane through the shadow boundary and the light source. Note that we will only try to use this cue if the shadow boundary appears to be a straight line on the continuous portion of a surface near a surface discontinuity and hence that G_I is a plane. If we know G_C and G_I we can determine the slope of c since

$$\begin{gathered}
\Delta x_c * p_C + \Delta y_c * q_C + \Delta z_c = 0 \\
G_C \cdot c = -\Delta z_c \\
\Delta x_c * p_I + \Delta y_c * q_I + \Delta z_c = 0 \\
G_I \cdot c = -\Delta z_c \\
(G_C - G_I) \cdot c = 0 \\
(p_C - p_I) * \Delta x_c = (q_I - q_C) * \Delta y_c \\
\frac{\Delta x_c}{\Delta y_c} = \frac{q_I - q_C}{p_C - p_I}.
\end{gathered}$$

If we do not know G_I but we do know the orientation G_A of surface A and the orientation G_B of surface B then we can derive G_I. From similar arguments as those used above, it follows that

$$\begin{gathered}
G_A \cdot a = -\Delta z_a = G_I \cdot a \\
G_B \cdot b = -\Delta z_b = G_I \cdot b
\end{gathered}$$

and hence that

$$\begin{gathered}
\begin{bmatrix} a \\ b \end{bmatrix} G_I^T = \begin{bmatrix} G_A \cdot a \\ G_B \cdot b \end{bmatrix} \\
G_I^T = \begin{bmatrix} a \\ b \end{bmatrix}^{-1} \begin{bmatrix} G_A \cdot a \\ G_B \cdot b \end{bmatrix}.
\end{gathered}$$

The superscript T stands for array transpose.

However, we cannot assume that an observer will always know the orientation or location of the surfaces on which a shadow is cast. Nor can we assume that an observer has recognized all surface discontinuities. Instead we will assume that strong, isolated step edges correspond to surface discontinuities. Obviously, surface discontinuities need not produce such image discontinuities and alternatively, that albedo changes can produce such image discontinuities. But, when a proposed shadow boundary changes direction across a strong image discontinuity, we will take it as supporting evidence for the shadow labeling. However, when a proposed shadow boundary does not change direction across a strong image discontinuity, we will not consider the evidence conclusive of any hypothesis unless there

is additional evidence that the image discontinuity does in fact correspond to a surface discontinuity. In latter case, the shadow hypothesis is undermined.

6.6 Texture Continuation

Under general viewing conditions it is unlikely that shadow boundaries will align with a change in surface markings, including surface texture. Consequently, the continuation of texture despite a change in intensity has been used to detect shadows [Witkin 1982], [Adjouadi 1986]. Both authors use correlation of image samples taken on either side of a possible shadow boundary to judge if texture continues across the boundary. Adjouadi also compares the power spectra sampled on either side of possible shadow boundary for similarity of form. The exact nature of this comparison is not stated by Adjouadi.

A large body of literature exists concerning various texture measures (see [Haralick 1979] for a review). Any of these techniques may be helpful with detecting shadows if the texture measures recovered do not vary with the types of spectral changes associated with shadows. However, as we have seen, different color changes are possible for differently colored materials when shadowed. In addition, relatively little work has been done on color texture discrimination. Consequently, the texture techniques used by Witkin and Adjouadi are limited in scope to textures that vary only in intensity.

In this work we do not intend to examine what texture measures are best suited to the problem of recognizing texture despite shadows. Currently, no completely general texture recognition scheme exists in the sense that it can discriminate between all classes of visual textures that humans do. Texture remains a difficult problem in automatic image interpretation and is outside the scope of this proposal.

However, given a suitable texture measure, texture continuation would be tested in our system within the shadow candidate regions found from color image segmentation for textures that varied only in intensity. Color textures would need to be examined separately from our color image segmentation or the examination would need to be done on collections of color regions found during segmentation.

It is also extremely important that for any texture measure used across a possible shadow boundary that we have a criteria for determining the variability of the texture measure when the surface is consistently lit. Without knowing the variability of the texture when consistently lit, we cannot judge if the texture continues when the surface irradiance changes. However, with a texture variability measure we have a simple threshold criteria

for judging if the texture continues across the possible shadow boundary. Neither Witkin nor Adjouadi appear to have implemented such a test.

CHAPTER 7

Research Proposal

We have described elements of a system that would enable an active observer to interpret images acquired of its environment in such a way as to take into account and utilize shadows in the scene. We have made special use of a probe at the observer's disposal that is used to generate shadows. From what the observer can recover about the scene's geometric and spectral properties from shadows it casts itself, we have outlined how to analyze the naturally occurring shadows in the scene.

7.1 Control Structure Outline

This proposal stresses what cues an active observer can use to recognize shadows in its environment and how the individual tests can be implemented. Less emphasis is put on efficiency in the use of the cues and on any interplay between the cues. We believe that all the cues can be tested for with relatively little computation time.

We propose to segment color images into regions such that a single surface seen lit and in shadow will be represented by piece-wise linear color clusters. For scenes lit with a single light source, we assume that shadows will have a penumbra and hence that the light source is not a point light source. An analysis of the regions produced by image segmentation, based on knowledge gained from shadows actively cast by the observer will be done based on seven shadow cues. We propose to use all the cues for which we have the prerequisite data. From the cues we will compile evidence supporting, undermining, or refuting the possibility that an image region corresponds to a shadow in the scene. See Table 7.1 for a list of the types of evidence provided by each cue. When discussing the cues we will use the numeric labels found in Table 7.1.

Shadow Cue	Evidence		
	Supporting	**Undermining**	**Refuting**
1. Color	Color trend matches probe shadows.	Color trend bucks probe shadows.	Non-uniform intensity change across RGB.
2. Shadow making object present	Some object in the image between the shadow and the light source.		No object found.
3. Projection of a silhouette	Shadow and object can be put in correspondence.	Shadow and object cannot be put in correspondence.	
4. Penumbra and umbra	Penumbra surrounding solid umbra.		
5. Shadows are on a surface	Surface and shadow coplanar.	Surface and shadow not coplanar.	Planar surface without coplanar shadow.
6. Shadow boudnaries and surface discontinuities	Shadow boundaries change direction when crossing image discontinuities.	Shadow boundaries don't change direction across surface discontinuities.	For a known surface geometry, shadow boundaries aren't as predicted.
7. Texture continuation	Texture continues across a shadow boundary.	Texture ends at a shadow boundary.	

Table 7.1: The nature of the evidence provided by each cue in determining whether or not an image region corresponds to a shadow in the scene.

Some of the cues depend on the results of casting shadows with the probe. Cue 1 relies on the spectral analysis of the probe's shadow. The usefulness of the spectral analysis of the probe's shadows depends on the observer's examination of the color plate on the back of the probe to determine that the ambient lighting conditions have not changed. Cues 2 and 3 rely on a determination of the location of the light source. We propose to locate the light source using the probe and its shadow. These cues cannot be used until the observer has done at least one experiment with the probe. In addition, all shadow hypotheses are suspect until the observer has successfully cast a shadow and hence verified that shadows are a possibility in the current environment. And, the color image segmentation that underlies our shadow analysis becomes suspect if the Linear Color Cluster Assumption is found not to hold for the penumbra of the probe's shadow (and by assumption for shadows throughout the scene).

Most of the techniques we have presented for detecting shadow cues improve in reliability with additional data, additional processing time, or better scene knowledge. Cue 1 improves with the number of different surfaces within a single scene onto which the agent has cast the probe's shadow. Cue 1 also improves with a better analysis of the data from the probe shadows. Both cues 2 and 3 improve with the estimate of the location of a light source. In addition, both cues improve with the number of objects whose shape and ability to produce shadows has been recognized. Cue 4 becomes more reliable if we know that the objects in a scene meet our assumption of containing no holes and of having limited overlap relative to the direction of a light source. Cues 5 and 6 improves with our knowledge of surface gradients and orientation relative to the light sources. Cue 7 improves with the observer's ability to describe and recognize texture. Consequently, we expect an active observer's ability to recognize shadows to improve with the time that the observer has to explore its environment.

We propose in our initial experiments to take any undermining or refuting evidence as clear evidence that an image region is not a shadow. In addition, we will measure the confidence we have that an image region is a shadow based on the number of supporting pieces of evidence. Each cue can be counted again in new views of the same image region as the observer moves in the environment. We do not address here how the observer maintains object identity as new images are acquired. For experimental purposes, we will do this manually when we have sequences of images. Through experiments with images of shadows in a variety of environments we plan to test the adequacy of our proposed control structure. If a refined control structure is needed we hope to base it on what we learn from our

experiments.

7.2 What Needs to Be Done

Algorithms have been implemented and methodologies tested for the work in Chapter 3, *Shadow Probe*, and in Section 5.1, *Color Image Segmentation.* Only preliminary testing has been done on the methods of Section 5.2, *Color Region Analysis.* The methods of Chapter 4, *Shadow Probe Geometry*, and Chapter 6, *Shadow Candidates and Other Cues*, have yet to be implemented. Further testing is need for the work of all chapters.

We propose to test our system on images of a variety of scene types. These will include scenes contrived in our lab to contain the cues we have listed for detecting shadows. These scenes will generally contain objects such as wood and plastic blocks. We will also examine scenes in the lab which we have not arranged. These will be taken as examples of *natural* indoor scenes. Finally, we will examine natural outdoor scenes on sunny and hazy days.

7.3 Contributions

If successful, our system for recognizing shadows would be a great aid to the computer vision community. Various existing visual modules require that there be an accounting for shadows prior to their use. Examples of such visual modules include object recognition, road following for autonomous navigation, and shape from shading. Consequently, there is a real need for efficient shadow identification prior to the completion of surface and object recovery.

To date, methods of identifying shadows have been overly simplistic – generally relying on shadows to be the darkest parts of an image. In this work we make use of the spectral and geometric properties of shadows in order to devise a set of cues that strongly suggest the existence of a shadow. These cues work on image regions and hence, we only require that an image be segmented into regions of related color. However, if geometric information is available for the scene, then the observer's ability to successfully recognize shadows will improve under our system.

In the course of identifying shadows, we also present a new modification on an existing image segmentation algorithm [Leonardis et al. 1990]. Our modification provides a general description of color images in terms of regions that is particularly amenable to the analysis of shadows.

We also present methods by which an observer can learn about its environment from shadows. These are shadows that the observer actively casts using a shadow probe. These shadows allow the observer to experimentally determine the number and location of light sources in the scene, to locate the cast shadows, and to gain information about the likely spectral changes due to shadows. The method for locating a light source and the surface on which it is cast is new. It takes into account errors in imaging and image processing and it takes special advantage of the benefits of an active observer. The information gained from the probe is of particular importance in effectively using the various shadow cues.

APPENDIX A

Penumbra Width

Here we will determine the width of a penumbra for a shadow in 2D. Let P be the width of the penumbra. Let W be the width of the outer envelope of the light source as "seen" from one end of shadow making line. See Figure A.1. Let S be the distance from the penumbra to the shadow making line and let L be the distance from the shadow making line to the light source. The definition of the various angles can be seen from the figure. From the law of sines we know that

$$\frac{P}{\sin B} = \frac{S}{\sin A}$$

and that

$$\frac{W}{\sin B} = \frac{L}{\sin C}.$$

So, the width of the penumbra is

$$P = (\frac{S}{L})(\frac{W \sin C}{\sin A}). \tag{A.1}$$

Often the envelope of the light source is nearly parallel to the ground and in that case

$$\sin A \approx \sin C \Longrightarrow P \approx \frac{S\,W}{L}.$$

In addition, W is generally fixed and we are only interested in the case where the distance between the shadow making object and the ground varies. In this case, we have

$$P \propto \frac{S}{L}. \tag{A.2}$$

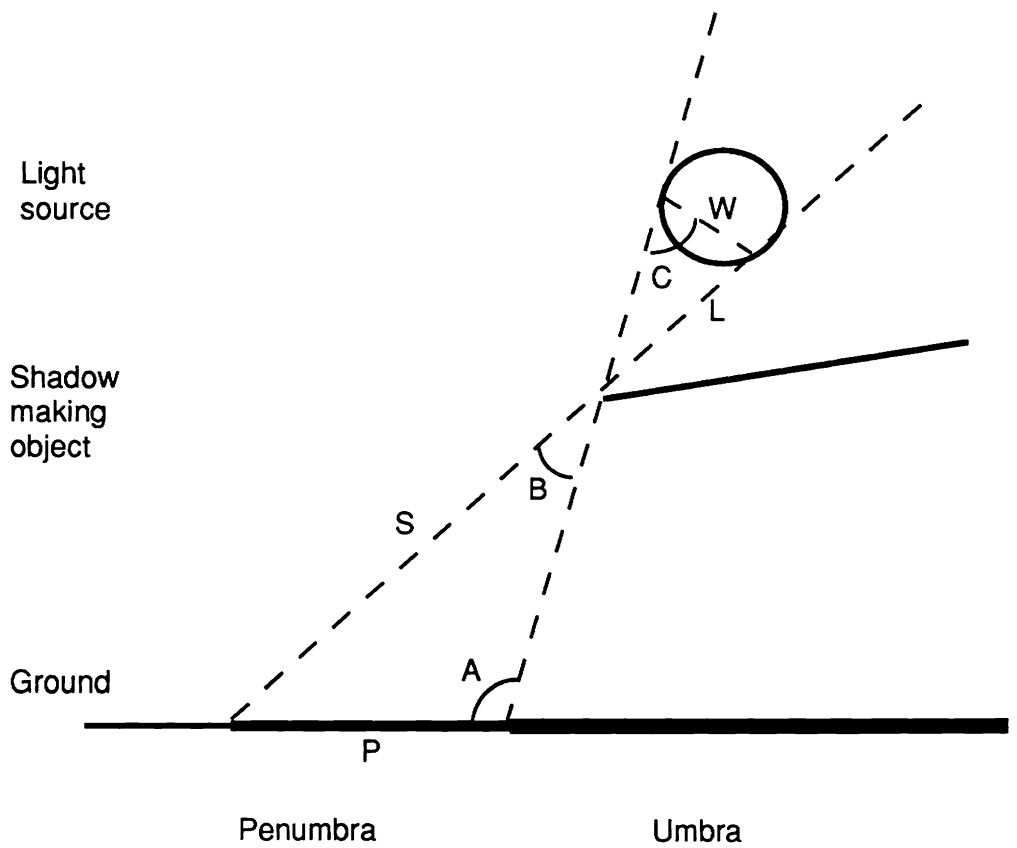

Figure A.1: 2D Geometry of a Shadow Penumbra.

APPENDIX B

Determining Planarity Based on Points in Two Images

B.1 Cross Ratio

The cross ratio is a description of the relationship between four points that lie on a line that is invariant to perspective projection. We define the cross ratio as

$$CR(x_1, x_2, x_3, x_4) = \frac{(x_3 - x_1)(x_2 - x_4)}{(x_2 - x_1)(x_3 - x_4)}. \tag{B.1}$$

For Figure B.1 the cross ratio projective invariance is

$$CR(x_1, x_2, x_3, x_4) = CR(z_1, z_2, z_3, z_4) \tag{B.2}$$

$$CR(y_1, y_2, y_3, y_4) = CR(z_1, z_2, z_3, z_4). \tag{B.3}$$

See [Duda and Hart 1973] for the outline of a proof for the above. Alternate definitions of the cross ratio can be realized by permuting the labels of the four points. There are, however, only six distinct possibilities.

B.2 Two-Dimensional Projective Coordinates

Two-dimensional projective coordinates are a description of the relationship between five points that lie on a plane that is invariant to perspective projection. Call the five points in question (A, B, C, U, P). See Figure B.2 for an illustration of the points and one of their projective coordinates. Projective coordinates are defined relative to the triangle defined

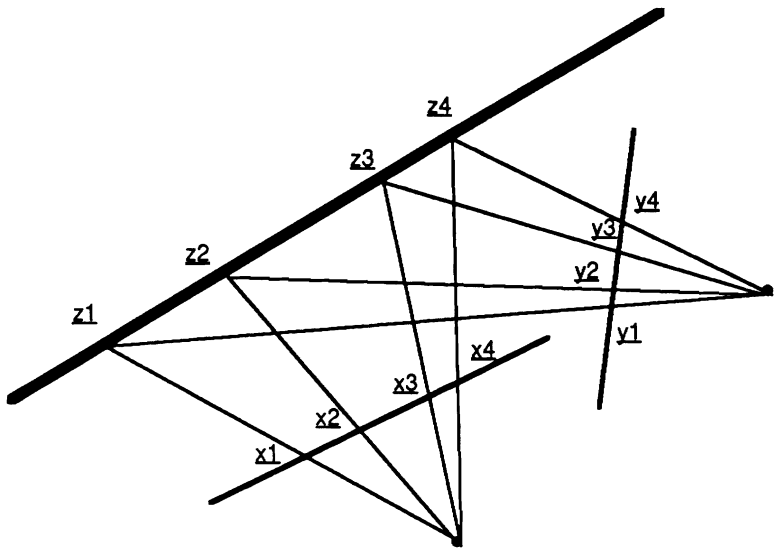

Figure B.1: The Cross Ratio.

by the points (A, B, C) and hence there are three projective coordinates for the five points. For five points we need only two projective coordinates to uniquely specify the invariant relationship. However, in the case where P lies on the side of the triangle defined by (A, B, C) the coordinate on that side must be used. The projective coordinate on the AC axis is defined as $CR(A, X, Y, C)$, where the differences measured in Equation B.1 are now signed distances between points in 2-space. A proof of the invariance of two-dimensional projective coordinates based on the cross ratio is given in [Duda and Hart 1973].

If the projective coordinates of 5 points seen in two images are not equivalent then either they are not the same five labeled points or they do not lie on a plane. We will assume that the points have been correctly put into correspondence and consequently that any time the projective coordinates are not equivalent that this is proof that that the five points do not lie on a single plane.

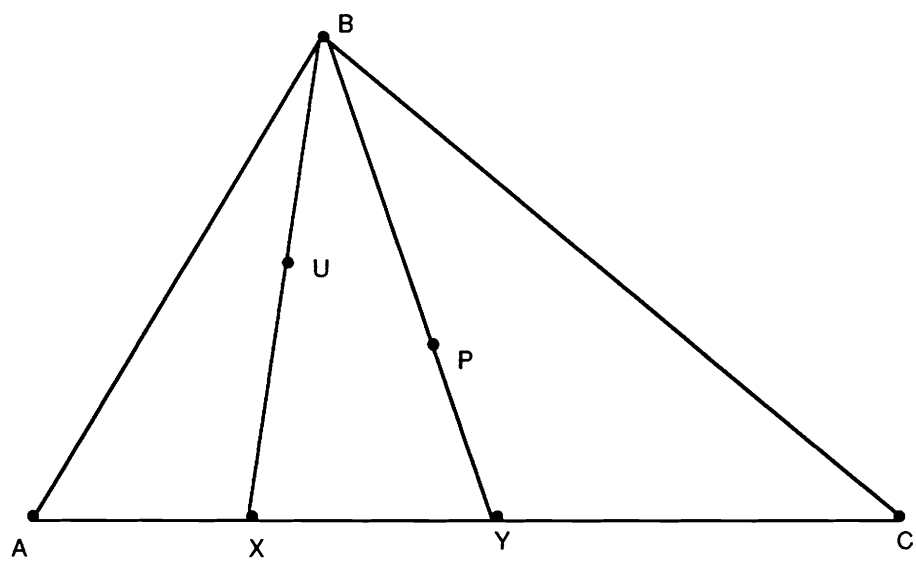

Figure B.2: Two-Dimensional Projective Coordinates.
For the five data points (A, B, C, U, P) we define the projective coordinates on the AC axis as $CR(A, X, Y, C)$. Projective coordinates on the BC or CA axes can be defined similarly.

Bibliography

[Adjouadi 1986] Malek Adjouadi, Image analysis of shadows, depressions, and upright objects in the interpretation of read-world scenes, in *Proceedings of the Eighth International Conference on Pattern Recognition,* pp. 834–838, Paris, France, October 1986.

[Bajcsy et al. 1989] Ruzena Bajcsy, Sang Wook Lee, and Aleš Leonardis, *Computational Aspects of Color Constancy,* Technical Report MS-CIS-89-50, GRASP LAB 189, University of Pennsylvania, Philadelphia, PA, December 1989.

[Berzins 1984] V. Berzins, Accuracy of laplacian edge detection, *Computer Graphics and Image Processing,* 27:195–210, 1984.

[Canny 1986] John Canny, A computational approach to edge detection, *IEEE Transactions on Pattern Analysis and Machine Intelligence,* 8(6):679–698, November 1986.

[Crisman 1990] Jill Dawn Crisman, *Color Vision for the Detection of Unstructured Roads and Intersections,* PhD thesis, Electrical and Computer Engineering Department, Carnegie Mellon University, May 1990.

[Duda and Hart 1973] Richard O. Duda and Peter E. Hart, *Pattern Classificaiton and Scene Analysis,* John Wiley & Sons, Inc., 1973, pages 407-414.

[Forsyth et al. 1991] David Forsyth, Joseph L. Mundy, Andrew Zisserman, Chris Coelho, Aaron Heller, and Charles Rothwell, Invariant descriptors for 3-D object recognition and pose, *IEEE Transactions on Pattern Analysis and Machine Intelligence,* 13(10):971–991, October 1991.

[Funka-Lea and Bajcsy 1992] Gareth Funka-Lea and Ruzena Bajcsy, Vision for vehicle guidance using two road cues, in *Proceedings Intelligent Vehicles '92,* Detroit, Michigan, June 1992.

[Gershon et al. 1986] Ron Gershon, Allan D. Jepson, and John K. Tsotsos, Ambient illumination and the determination of material changes, *Journal of the Optical Society of America A,* 3(10):1700–1707, October 1986.

[Gibson 1966] James J. Gibson, *The Senses Considered as Perceptual Systems,* Houghton Mifflin Co., Boston, Mass., 1966.

[Hambrick et al. 1987] Larry N. Hambrick, Murray H. Loew, and Robert L. Carroll Jr., The entry-exit method of shadow boundary segmentation, *IEEE Transactions on Pattern Analysis and Machine Intelligence,* 9(5):597–607, September 1987.

[Haralick 1979] Robert M. Haralick, Statistical and structural approaches to texture, *Proceedings of the IEEE,* 67(5):786–804, May 1979.

[Huertas and Nevatia 1988] A. Huertas and R. Nevatia, Detecting buildings in aerial images, *Computer Vision, Graphics, and Image Processing,* 41(2):131–152, February 1988.

[Irvin and McKeown 1988] R. B. Irvin and David M. McKeown, Methods for exploiting the relationship between buildings and their shadows in aerial images, *IEEE Transactions on Systems, Man and Cybernetics,* 19(6):1564–1575, November 1988.

[Jiang and Ward 1992] Caixia Jiang and Matthew O. Ward, *Shadow Segmentation and Classification,* Technical Report, Computer Science Dept., Worcester Polytechnic Institute, Worcester, MA, 1992.

[Kender and Smith 1987] John R. Kender and Earl M. Smith, Shape from darkness: deriving surface information from dynamic shadows, in *Proceedings of the First International Conference on Computer Vision,* pp. 539–546, London, UK, June 1987.

[Lee 1991] Sang Wook Lee, *Understanding of Surface Reflections in Computer Vision by Color and Multiple Views,* PhD thesis, GRASP Laboratory, Dept. of Electrical Engineering, University of Pennsylvania, August 1991.

[Leonardis et al. 1990] Aleš Leonardis, Alok Gupta, and Ruzena Bajcsy, Segmentation as the search for the best description of the image in terms of primitives, in *IEEE International Conference on Computer Vision,* Osaka, Japan, December 1990.

[Liow and Pavlidis 1990] Yuh-Tay Liow and Theo Pavlidis, Use of shadows for extracting buildings in aerial images, *Computer Vision, Graphics, and Image Processing,* 49(2):242–277, February 1990.

[Liu and Moore 1990] J. G. Liu and J. M. Moore, Hue image RGB color composition: a simple technique to suppress shadow and enhance spectral signature, *International Journal of Remote Sensing,* 11(8):1521–1530, 1990.

[Lowe and Binford 1985] David G. Lowe and Thomas O. Binford, The recovery of three-dimensional structure from image curves, *IEEE Transactions on Pattern Analysis and Machine Intelligence,* PAMI-7(3):320–326, May 1985.

[Marr 1982] David Marr, *Vision,* W. H. Freeman and Co., New York, NY, 1982, pages 88-89.

[Nagao et al. 1979] Makota Nagao, Takashi Matsuyama, and Yoshio Ikeda, Region extraction and shape analysis in aerial images, *Computer Graphics and Image Processing,* 10(3):195–223, July 1979.

[Nishita et al. 1985] Tomoyuki Nishita, Isao Okamura, and Eihachiro Nakamae, Shading models for point and linear sources, *ACM Transactions on Graphics,* 4(2):124–146, April 1985.

[Ranson and Daughtry 1987] K. Jon Ranson and Craig S. T. Daughtry, Scene shadow effects on multispectral response, *IEEE Transactions on Geoscience and Remote Sensing,* GE-25(4):502–509, July 1987.

[Rubin and Richards 1988] John M. Rubin and W. A. Richards, Color vision: representing material categories, in *Natural Computation,* edited by Whitman Richards, MIT Press, 1988.

[Scanlan et al. 1990] Joseph M. Scanlan, Douglas M. Chabries, and Richard W. Christiansen, A shadow detection and removal algorithm for 2-D images, in *IEEE Proceedings of the International Conference on Acoustics, Speech, and Signal Processing,* pp. 2057–2060, Albuquerque, NM, April 1990.

[Sedgewick 1983] Robert Sedgewick, *Algorithms,* Addison-Wesley Inc., Reading, Mass., 1983.

[Shafer 1985a] Steven A. Shafer, *Shadows and Silhouettes in Computer Vision,* Kluwer Academic Publishers, Boston, Mass., 1985.

[Shafer 1985b] Steven A. Shafer, Using color to separate reflection components, *COLOR Research and Applications,* 10:210–218, 1985.

[Takita et al. 1991] Shinichi Takita, Kazufumi Kaneda, Toshio Akinobu, Haruhiko Iriyama, Eihachiro Nakamae, and Tomoyuki Nishita, A simple method for rendering penumbrae caused by sunlight, *The Visual Computer,* 7(5-6):259–267, September 1991.

[Thirion 1992] Jean-Philippe Thirion, Realistic 3D simulation of shapes and shadows for image processing, *CVGIP: Graphical Models and Image Processing,* 54(1):82–90, Januaray 1992.

[Thompson et al. 1987] William B. Thompson, Michael T. Checky, and William F. Kaemmerer, Shadow stereo – locating object boundaries using shadows, in *Proceedings of the National Conference on AI (AAAI-87),* pp. 761–766, Seattle, WA, July 1987.

[Torrance and Sparrow 1967] K. E. Torrance and E. M. Sparrow, Theory for off-specular reflections from roughened surfaces, *Journal of the Optical Society of America,* 57:1105–1114, 1967.

[Turk et al. 1988] Matthew A. Turk, David G. Morgenthaler, Keith D. Gremban, and Martin Marra, VITS – A vision system for autonomous land vehicle navigation, *IEEE Transactions on Pattern Analysis and Machine Intelligence,* 10(3):342–361, May 1988.

[Waltz 1975] David Waltz, Understanding line drawings of scenes with shadows, in *The Psychology of Computer Vision,* edited by Patrick Henry Winston, McGraw-Hill Book Co., 1975.

[Witkin 1982] Andrew P. Witkin, Intensity-based edge classification, in *Proceedings of the National Conference on AI (AAAI-82),* pp. 36–41, Pittsburgh, PA, August 1982.

[Woo et al. 1990] Andrew Woo, Pierre Poulin, and Alain Fournier, A survey of shadow algorithms, *IEEE Computer Graphics and Applications,* 10(6):13–32, November 1990.