We are gradually moving into the third part of the course, which deals with estimation of 3D properties of scenes. Last lecture we discussed how to estimate vanishing points, which directly relate 2D image properties to a 3D property of the scene. Today, we will look at a different kind of relationship between image properties and 3D scene properties, called *shading*.

Shading

The term *shading* typically refers to variations in irradiance along a smooth surface. Recall that if a surface point is illuminated by parallel light source from direction $\mathbf{l}$, then the surface irradiance at the point is proportional to

$$E(\mathbf{X}) = \ L_{src}\ \Omega\ \mathbf{n}(\mathbf{X}) \cdot \mathbf{l}.$$

We will ignore the dependence of $L_{src}\Omega$ since they are constant.

Let's suppose that the surface is Lambertian (recall lecture 5 page 6) and has spatially constant reflectance, that is, the reflectance does not vary with position $\mathbf{x}$ on the surface. An example is a constant color bedsheet (or other drapery surface, or clothing with no pattern printed on it). In this case, the above model implies that the radiance of light reflected from the surface varies only with the surface normal. Everything else in the above equation is assumed to be constant.

Weak perspective

To keep the discussion of shading as simple as possible, we consider an *isolated surface* in the scene. Let XYZ be the usual camera coordinates and assume we can write points on the surface as $(X, Y,\ Z(X, Y))$. We also assume that the range of Z values on this surface is small relative to (say) the minimal depth Z value, which we call Z_0. That is, we assume $|\frac{Z-Z_0}{Z_0}|$ is near zero for all Z values of points on the surface. For example, suppose the surface is a person's shirt (covering their torso) which is seen from a distance of at least a few meters. The distance between the farthest part of the shirt might only be 10 cm greater than the distance to the nearest part of the shirt.

The above assumptions lead to the following *weak perspective* model. For each point (X, Y, Z) on the surface, we have

$$x = f\frac{X}{Z} \approx f\frac{X}{Z_0}$$

$$y = f\frac{Y}{Z} \approx f\frac{Y}{Z_0}.$$

We can thus write surface irradiance $E(X, Y, Z)$ as a function of (X, Y) only (since Z is a function of (X, Y)),

$$E(X, Y) \approx E(\frac{Z_0}{f}x, \frac{Z_0}{f}y) \tag{1}$$

[ASIDE: This is sometimes called *scaled orthographic projection*, namely we are projecting the points parallel to the Z axis, and then scaling the image plane by a factor $\frac{Z_0}{f}$.]

The reason this last equation is important is that, because the surface is Lambertian, surface irradiance at $\mathbf{X}$ is proportional to the radiance of light reflected from $\mathbf{X}$ and so surface irradiance is proportional to image irradiance at (x, y). Thus there is a direct relationship between image irradiance variations and surface irradiance variations. In what follows, we work with surface irradiance only, but keep in mind that this is basically the same as image irradiance for the conditions via Eq. (1).

Surface normal

The surface normal vector is perpendicular to the surface and so, not surprisingly, it is determined by the *depth gradient* $(\frac{\partial Z}{\partial Y}, \frac{\partial Z}{\partial X})$. This can be seen as follows. Consider a step on the surface from (X, Y, Z) to some other nearby point $(X + \Delta X, Y + \Delta Y, Z + \Delta Z)$ on the surface. Then

$$\Delta Z \approx \frac{\partial Z}{\partial X}\Delta X + \frac{\partial Z}{\partial Y}\Delta Y$$

or

$$(\Delta X, \Delta Y, \Delta Z) \cdot (\frac{\partial Z}{\partial X}, \frac{\partial Z}{\partial Y}, -1) \approx 0 \ .$$

The latter relationship holds for any step $(\Delta X, \Delta Y, \Delta Z)$ along the surface. It follows that the vector $(\frac{\partial Z}{\partial X}, \frac{\partial Z}{\partial Y}, -1)$ is perpendicular to the surface. Hence this vector is in the direction of the *surface normal*. If we rescale this vector to unit length, then we get the *unit normal vector*

$$\mathbf{n} \equiv \frac{1}{\sqrt{(\frac{\partial Z}{\partial X})^2 + (\frac{\partial Z}{\partial Y})^2 + 1}}(\frac{\partial Z}{\partial X}, \frac{\partial Z}{\partial Y}, -1) \ .$$

Note that the Z component of $\mathbf{n}$ is negative, since positive Z goes away from the observer, and so for the surface to be visible the normal must have a negative Z. Here we are talking about the *outward normal*, i.e. out of the object.

The surface irradiance is then

$$E(X, Y) \ = \ L_{src}\Omega \frac{1}{\sqrt{(\frac{\partial Z}{\partial X})^2 + (\frac{\partial Z}{\partial Y})^2 + 1}} \ (\frac{\partial Z}{\partial X}, \frac{\partial Z}{\partial Y}, -1) \cdot (l_X, l_Y, l_Z). \quad (2)$$

where $\mathbf{l} = (l_X, l_Y, l_Z)$ is the unit vector pointing toward the light source.

This model holds only when $\mathbf{n}(X, Y) \cdot \mathbf{l} \ \geq 0$, since it is meaningless to have negative intensities. If the inner product of $\mathbf{n}$ and $\mathbf{l}$ is less than zero, this implies would imply that the surface is facing away from the light source at that point. In this case, the surface would not be illuminated by the source. It would be in shadow, and its illuminance component from the source would be zero (not negative). To keep the notation down, we are not considering this case. But you should understand it is there.

Bas relief

Let's look at the case of a *bas relief* surface, namely a surface that is nearly flat (planar) except for small hills and valleys. An example is a wrinked shirt or stucco, a coin, the bumps on the bricks on the painted wall in our classroom. In particular, we restrict ourselves to the case that the surface slopes $\frac{\partial Z}{\partial X}$ and $\frac{\partial Z}{\partial Y}$ are sufficiently small in magnitude, that one can approximate Eq. (2) by a Taylor series expansion around $(\frac{\partial Z}{\partial X}, \frac{\partial Z}{\partial X}) = (0, 0)$, and we keep terms up to second order.

Recalled that $(1+u)^{-\frac{1}{2}} \approx 1 - \frac{1}{2}u +$ higher order terms, we get

$$\frac{1}{\sqrt{1+(\frac{\partial Z}{\partial X})^2+(\frac{\partial Z}{\partial Y})^2}} = 1 - \frac{1}{2}\{(\frac{\partial Z}{\partial X})^2 + (\frac{\partial Z}{\partial Y})^2\} + \ H.O.T.$$

We then substitute into Eq. (2):

$$E(X,Y) \quad = \quad \{1 - \frac{1}{2}\{(\frac{\partial Z}{\partial X})^2 + (\frac{\partial Z}{\partial Y})^2\} + \ H.O.T.\} \ (\frac{\partial Z}{\partial X} l_X + \frac{\partial Z}{\partial Y} l_Y - l_Z)$$

If $|\frac{\partial Z}{\partial X}|$ and $|\frac{\partial Z}{\partial Y}|$ are small enough, then we can ignore terms that are higher than second order in these partial derivatives. This gives:

$$\begin{aligned} E(X,Y) \quad &\approx \quad -l_Z \quad + \quad (\frac{\partial Z}{\partial X} l_X + \frac{\partial Z}{\partial Y} l_Y \ ) \ + \frac{l_Z}{2}((\frac{\partial Z}{\partial X})^2 + (\frac{\partial Z}{\partial Y})^2) \\ &= \quad \text{constant} \quad + \quad \text{linear} \quad + \quad \text{quadratic} \end{aligned}$$

The linear component depends on l_X, l_Y. The quadratic component depends on l_Z. The linear term tends to dominate over the quadratic term if $\mid l_Z \mid \ll |(l_X, l_Y)|$, whereas the quadratic term tends to dominate when $\mid l_Z \mid \gg |(l_X, l_Y)|$.

Example: 2d cosine

A specific example is a cosine function with depth modulation in the X direction, similar to a hanging curtain,

$$Z(X,Y) = Z_o + \ a \ \sin(kX)$$

where the constant a is sufficiently small that the high order terms can be ignored. Then,

$$\frac{\partial Z}{\partial X} = ak\cos(kX)$$

and so

$$E(X,Y) \approx -l_Z \ + \ a \ l_X \ k \ \cos(kX) + l_Z \frac{a^2}{2} k^2 \ \cos^2(kX).$$

Notice that l_Y has no explicit effect here, though I does have an implicit effect, since (l_X, l_Y, l_Z) is a unit length vector e.g. large l_Y would automatically imply that l_x and l_Z are close to zero.

In class I gave a several "demos" of these effects. I used the overhead project as a parallel light source and I showed the shading that arises on a piece of paper that has been folded and then stretched out so that it has a sine wave like height function. The shading on the surface was out of phase with the height of the surface.

When the light was from the right and close to the "horizon" of the surface, so $(l_X, l_Y, l_Z) \approx (1, 0, 0)$, the maximum irradiance was on points where the normal was leaning toward the right. (See figure on next page.) This gives the maximum first order shading effect, and almost no second order effect because $l_Z \approx 0$.

If we keep this light source direction and then rotate the surface so that all depth modulation is in the Y direction, then the shading disappears. We saw this in the demo: when the corrugations

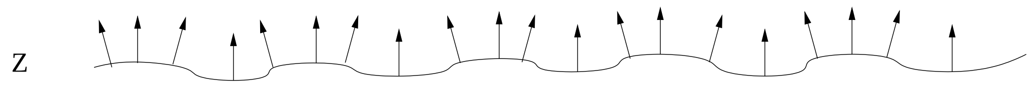

on the paper defined horizontal iso-depth lines and the source light was horizontal, the surface appeared uniform in brightness.

What about the quadratic term. If we illuminate the sine-depth surface from above, so that $(l_X, l_Y, l_Z) \approx (0, 0, 1)$, then only the quadratic term will be there. This term attains its maximum contribution when $\frac{\partial Z}{\partial X} = 0$, which for a cosine surface occurs at the maxima and the minima of depth $Z(X,Y)$. In particular, since $\cos^2\theta = \frac{1}{2}(\cos 2\theta + 1)$, it follows that the frequency of $I(X,Y)$ is twice that of $Z(X,Y)$ – that is, the maxima of $I(x,y)$ occur at both the maxima and minima of $Z(x,y)$, where the gradient vanishes. This is known as the *frequency doubling* effect. The modulations in irradiance that are due to the l_Z component (second order) occur at twice the frequency as those that are due to the l_X and l_Y components (first order).

Linear shape from shading ($L_Z \approx 0$)

Let's reconsider the condition that the light source direction $\mathbf{l}$ is near perpendicular to Z and Y axes of the surface, say $(l_X, l_Y, l_Z) \approx (\sqrt{1-\epsilon^2}, 0, \epsilon)$. In this case, we have

$$E(X,Y) = -l_Z + \frac{\partial Z}{\partial X} l_X$$

For each Y, we can do a 1D integration to get:

$$\begin{aligned} \int_{X_0} E(X,Y)dX &= -l_Z(X - X_0) + l_X \int_{X_0}^{X} \frac{\partial Z}{\partial X} dX \\ &= -l_Z(X - X_0) + l_X(Z(X) - Z(X_0)) \end{aligned}$$

Thus, given the surface irradiance function, we can say something about the Z values of different points on the surface [1].

This example should at least give you a sense of how you might solve the general *shape from shading* problem.[2] In the general problem, one trieds to solve the "non-linear first order partial differential equation", namely Eq. (1). The solution also involves integration of the partial derivatives of Z along an image curve.

[1] see A.P. Pentland, "Linear shape from shading", in International Journal of Computer Vision, 1990

[2] The first person to solve this problem was Berthold Horn at MIT in his Ph.D. thesis in the late 1960's, and there was much followup work on this problem in the 1980s and even some in the 1990s.