Image formation - color

August 31, 2017

We have seen *where* in the 3D world a pixel on the image corresponds to. Now the question is, what determines the color of the pixel?

1 Radiance

Consider a piece of paper left outside at midday. We want to ask how much radiant energy, or rather radiant power (energy per unit time) falls on the paper?

Clearly, the larger the piece of paper, the more the power that falls on it. So if the piece of paper has an area A, then we might say that the paper gets radiant power $\Phi = LA$ for some constant L.

But this is not right. What if the light rays are not perpendicular to the piece of paper, but oblique? As can be seen from Figure 1, the piece of paper will receive less light. In fact, it can be seen that the amount of light received only depends on the projected area of the paper, projected onto a direction perpendicular to the light. If light is coming in at a direction that makes an angle θ with the surface normal, then the projected area is $A \cos\theta$. So we might claim that the radiant power incident on the piece of paper is $\Phi = LA\cos\theta$.

But this assumes the entire piece of paper gets the same amount of light. What if the different parts of the paper get different amounts of light? This

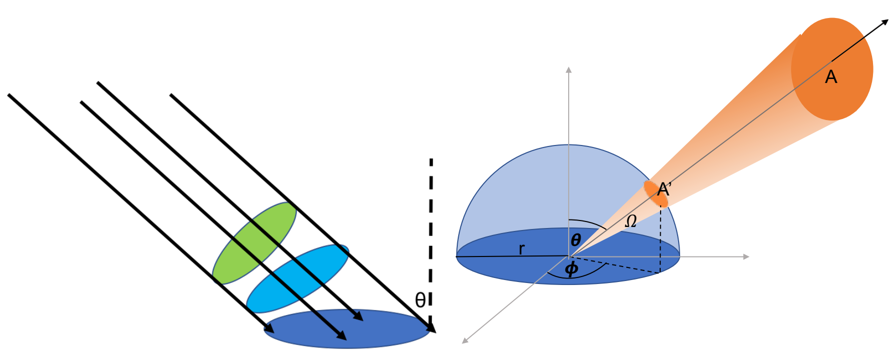

Figure 1: Setup. Ω is the solid angle, and equals the area A projected onto the sphere, A', divided by the squared radius

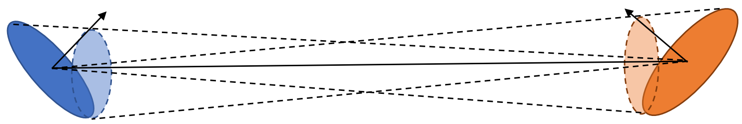

Figure 2: Conservation of radiance

would mean that L is no longer constant. What we can do is to divide the piece of paper into infinitesimal pieces of area dA. We will assume that the piece centered at $\mathbf{x}$ with area dA gets *power per unit projected area* $L(\mathbf{x})$. In this case, total power would add up the contributions of all infinitesimal pieces: $\int_A L(\mathbf{x}) dA \cos\theta$.

However, till now we have assumed that we are getting a constant amount of light from a fixed direction. But in general, the paper will receive different amounts of light from different directions. A direction in 2D is represented by an angle, but in 3D it is represented by a *solid angle*. Just as the angle subtended by an arc is defined as arc length divided by the radius, solid angles are subtended by *surface patches* on a sphere and are defined as patch area divided by the radius squared. Solid angle varies between 0 and 4π, and the unit is steradians.

As before, we can divide the set of all directions (i.e., all solid angles Ω) at a point into infinitesimal cones of directions, represented by an infinitesimal solid angle $d\Omega$. We can then assume that each infinitesimal patch of paper at location $\mathbf{x}$ receives a constant amount of *power per unit projected area per unit solid angle) from the infinitesimal cone in the direction* Ω denoted by $L(\mathbf{x}, \Omega)$. In this case, adding up the contributions of all infinitesimal pieces again, the total power is $\Phi = \int_S \int_A L(\mathbf{x}, \Omega) d\Omega dA \cos\theta$.

The quantity $L(\mathbf{x}, \Omega)$, which denotes the power per unit area per unit solid angle incident on a surface patch at $\mathbf{x}$ in direction Ω, is called the *radiance*. Recall that a pixel in a pinhole camera is a point on a screen that receives light from a particular direction: such a pixel records the incoming radiance at that point in that direction.

Here we were defining the *incoming radiance*, we can similarly define the *outgoing radiance*.

1.1 Conservation of radiance

Radiance has the property that it is conserved along a ray. To see this, consider a ray of light. This ray of light goes from a point $\mathbf{s}$ on one surface to a point $\mathbf{r}$ on another surface a distance R away. Let us assume that the ray from $\mathbf{s}$ to $\mathbf{r}$ is in direction θ_s, ϕ_s with respect to the first surface patch, and in direction θ_r, ϕ_r with respect to the second surface patch. Consider an infinitesimal surface patch of area dA_s on the first surface, and another of area dA_r on the second surface.

From the point of view of the first surface patch, let's calculate the total power it sends to the second surface patch. For this, we need the outgoing radiance in the first patch in the direction of the second, the solid angle subtended by the second patch on the first, and the projected area of the first patch.

To compute the solid angle subtended by second surface patch at $\mathbf{s}$, we want to project this area onto a sphere centered at $\mathbf{s}$. Let us do this projection onto a sphere centered at $\mathbf{s}$ of radius R. Because this patch is infinitesimally small and thus much smaller than the radius R and thus the sphere, instead of projecting onto the sphere, we can project it onto the tangent plane of the sphere. This tangent plane is perpendicular to the radius vector, which in this case is the vector from $\mathbf{s}$ to $\mathbf{r}$. The angle between this vector and the surface normal at dA_r is θ_r. Thus the area of the second patch projected onto the sphere is $dA_r \cos\theta$, and the solid angle subtended is $\frac{dA_r \cos\theta}{R^2}$

The rojected area of first surface patch in direction (θ_s, ϕ_s) is $dA_s \cos\theta_s$. Total power is then $L_o(\mathbf{s}, \theta_s, \phi_s) dA_s \cos\theta_s \frac{dA_r \cos\theta_r}{R^2}$.

From the point of view of the second surface patch, let's calculate the total power it receives from the first patch. Solid angle subtended by the first surface patch at $\mathbf{r}$ is $\frac{dA_s \cos\theta_s}{R^2}$. Projected area of second surface patch in direction (θ_r, ϕ_r) is $dA_r \cos\theta_2$. Total incident power is then $L_i(\mathbf{r}\theta_r, \phi_r) dA_r \cos\theta_r \frac{dA_s \cos\theta_s}{R^2}$.

Equating the two, we have that $L_o(\mathbf{s}, \theta_s, \phi_s) = L_i(\mathbf{r}, \theta_r, \phi_r)$.

This property means that we can track radiance along a ray. A camera pixel records the radiance of the incoming ray, and this property means that this is the same as the *outgoing radiance at the world point the pixel sees, in the direction of the pinhole.*

2 Relating input and output light

In general, the physics of how light incident on a surface is radiated out is very complex. However, there are two simple cases:

Specular surfaces : Specular surfaces act like mirrors. Every incident ray incident at an angle θ is reflected along a fixed direction that makes the same angle with the normal and is co-planar with the normal and the incident ray. Concretely, for a given outgoing direction θ_o, ϕ_o, there is a fixed input direction θ_i, ϕ_i:

$$L_o(\mathbf{x}, \theta_o, \phi_o) = L_i(\mathbf{x}, \theta_i, \phi_i) \tag{1}$$

Lambertian surfaces A lambertian surface is one where the output radiance is the same in all directions. If there is incident light of power I in direction $\hat{\mathbf{l}}$, and the surface normal is $\hat{n}$, then the outgoing radiance in all directions is:

$$L_o(\mathbf{x}, \theta_o, \phi_o) = \rho I \hat{\mathbf{l}} \cdot \hat{\mathbf{n}} \tag{2}$$

ρ here is called the albedo, and is the fraction of incident energy that is reflected by the object (as opposed to absorbed). For example, a black tyre has a low albedo, whereas snow has a high albedo. A lambertian surface looks like dull rubber.

Because a lambertian surface has the same radiance in all directions, this means that different cameras in different locations will see the same color for such a surface. Several computer vision algorithms will assume that objects are lambertian for this reason.

3 From image to shape

The expression in Equation 2 is a product of two parts: the albedo ρ, which essentially corresponds to the intrinsic color of the object, and $I\hat{\mathbf{l}} \cdot \hat{\mathbf{n}}$, which depends on the lighting and the geometry. Since pixels in an image also record the radiance, if we assume all objects in the scene are lambertian, the image itself is a product of two images: the *albedo* or *reflectance* image, which corresponds to the intrinsic color of the object, and the *shading* image, which is only a function of the lighting and geometry.

$$I = I_r \odot I_s \tag{3}$$

I_r and I_s are called intrinsic images, and this decomposition is called an intrinsic image decomposition.

However, in general, this decomposition into reflectance and shading images is hard. Equation 3 implies that the actual color of objects in an image depends not just on the objects intrinsic color, but also on shape and lighting. This makes it hard to identify the same object viewed under different lighting conditions, since the color seen in the image can change quite a bit.

4 Color and wavelength

We have talked till now of grayscale images, and just intensity. However, all the above quantities can be expressed as additionally functions of wavelength / color.