Illumination Structure Solves Multitudes of Applications

Introduction

By now, much of the machine vision industry has learned that lighting plays an important role in creating high quality image data that leads to robust vision system performance. Like any signal processing application, machine vision success is dictated by simple signal-to-noise (S/N) constraints, and proper front-end design (illumination and optics) is the absolute key to maximizing the S/N ratio.

What is less known, however, are how just a few simple principles tend to solve more than 80% of all applications. It is unfortunate that the application of illumination to vision applications is still viewed predominantly with a “trial and error”, “black art”, and/or “poke and hope” methodology. Because of these beliefs a tremendous amount of time, energy and money are spent each year reinventing the wheel for many of today’s machine vision applications.

The purpose of this article is to demonstrate how the knowledge of a few simple principles will guide any user to select appropriate lighting that will increase contrast and improve overall vision system performance. These principles will be discussed for generic applications, and will then be extended to cover an interesting set of special cases that use spherical illumination techniques.

Interaction of Light with Matter

The principle for how information is transferred from an object (device under inspection) to a detector (CCD camera) is based on how photons interact with the material within the object. If the device under inspection modifies the incoming light in such a way that the outgoing rays are different from the incoming rays, then we say that the object has created contrast. This is the basic principle of all machine vision applications, as well as all visually related techniques. If the object cannot modify the incoming beam in some discernable fashion, then the device cannot be visible to either a camera or the human eye.

The goal for the machine vision lighting engineer is then quite simple. Provide incoming illumination in such a fashion that the naturally occurring features and properties of the device under test can be exploited to maximize contrast. To maximize the effect that these modifying features exert, we must first understand basic light interaction properties and how they can be enhanced by intelligent illumination design.

When incident light impinges on the surface of any part, there are three basic properties that govern the interaction - reflection, absorption and transmission (see Figure 1). These properties allow the part to modify the incoming light and thereby pass information about itself to some detector (our eyes or a CCD camera). To understand how each of these mechanisms work for any generic object, we can analyze each of these mechanisms as they interact with model surfaces.

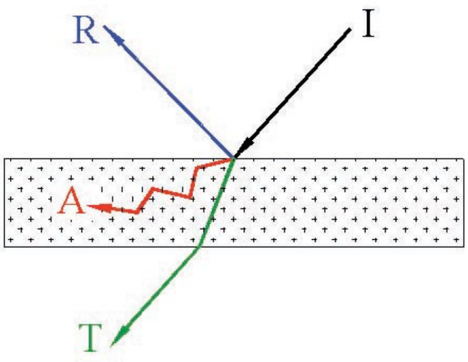

Figure 1: Basic Interaction of Incident Light on a Model Surface. Light is either Reflected (R), Absorbed (A) or Transmitted (T).

Reflection (R) occurs for most surfaces, because some amount of Incident radiation (I) falling on a surface will almost invariably be reflected. In fact, it is difficult to reduce this effect to a negligible value, even in many applications where this is desirable. The reflectivity of the surface is given as the ratio of reflected light to that of the incident radiation:

Reflectance = R/I

Reflectance is a number between zero and one, and is sometimes given as a percentage. A surface with high reflectivity, such as a mirror or a flat shiny surface, may have reflectivity values in the high 90% range but rarely approach 100% due to less than perfect surface interactions and imperfections. In these high reflectivity cases, the angle of reflection (β) equals the angle of incidence (α), as measured from the normal vector (perpendicular to the surface). This reflection condition is commonly referred to as Specular Reflection (see Figure 2).

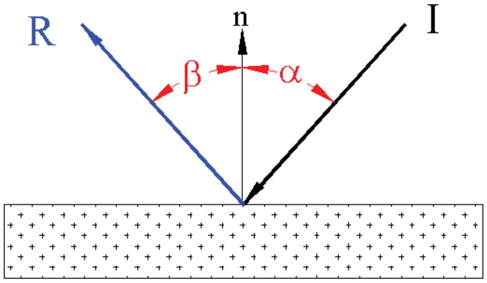

Figure 2: Specular Reflection occurs when light reflects from the surface at the same angle (with respect to normal) as the Incident ray.

Absorption also occurs for most surfaces, because some amount of incident radiation (I) falling on a surface will almost invariably be absorbed. The absorption of the material is given as the ratio of absorbed light to that of the incident radiation:

Absorbance = A/I

Absorbance is also a number between zero and one, and is sometimes given as a percentage. A material with high absorption, such as a piece of black felt or flock paper, may have Absorbance values below 5% but rarely approach an absolute zero value. As a special case of absorption, the photons can be re-emitted at another wavelength. This process is called Fluorescence and is rather special case that is outside the scope of the basic mechanisms that are important for the discussion of structure.

Transmission is the last basic interaction mechanism that occurs for objects that are transparent, translucent, or even opaque with physical holes. In these types of materials, some amount of the incident radiation (I) falling on a surface will be transmitted through the material. The transmittance of the material is given as the ratio of transmitted light to that of the incident radiation:

$$\boxed{\textbf{Transmittance} = T/I}$$

As before, Transmittance is a number between zero and one, and is sometimes given as a percentage. A material with high Transmittance, such as a piece of optical quality glass, may have Transmittance values in the high 90% range, but rarely approach 100% due to imperfections and impurities within the material.

The RAT formula

A real world part to be inspected will generally have a combination of properties that have element of reflectivity, absorption and transmission. Conservation of energy tells us that the sum of all forms of outgoing radiation must equal the amount of incoming radiation impinging on the material under inspection. This is commonly referred to as the RAT formula:

$$\boxed{I = R + A + T}$$

From this formula it can now be seen that every material has some characteristic ratio of reflectivity, absorption and transmission that allow it to modify incident light and pass information to a detector. The unique balance of these three parameters determine which interaction can most likely be exploited for enhancing contrast with respect to an adjacent feature which may have a different set of ratios.

If two materials have nearly identical ratios for the above set of parameters, it will be nearly impossible to generate a high contrast signal that will allow them to be visually differentiated. A typical application that demonstrates this would be looking for a shiny black label on a very smooth black rubber sheet. In this example, there is very little difference between the three parameters for the two different materials, and anyone that has tried to solve this one using simple optical techniques is probably still working on it!

So far, these formulas have been presented in a simple manner and treated as constants to gain a better understanding of the three basic mechanisms. In reality, there are two additional factors that make the concept more multifaceted. Each of the parameters of Reflectance, Absorbance and Transmission can vary as a function of the wavelength of incident light, and as a function of the angle of incidence of the light. That means that for many materials, the ratio of R, A, and T will change as a function of the color of light that impinges upon them, and will also change as a function of the angle of incidence of that light. Both of these factors give the illumination engineer additional tools that can be used to exploit differences between similar materials to produce higher contrast signals.

Where there is strong wavelength dependence for any one of these three parameters, the illumination engineer may use the spectral dependence to tailor the incoming spectrum of light to either increase or decrease the contrast with surrounding features. This spectral manipulation can be accomplished using filters in combination with white lightsources (Tungsten Halogen, Fluorescent, Metal Halide, etc.), or by the proper selection of wavelength (or wavelengths) from the more monochromatic sources (LED, laser, specialized fluorescent, etc.) This enhancement technique focuses on the color of light being used, and can be very helpful in improving feature contrast, even if the resulting inspection is done using monochrome cameras and gray scale algorithms.

The variation of these parameters as a function of the incident angle of the incoming light is based solely on geometry. Contrast that is driven solely on geometrical issues is commonly referred to as lighting structure. Increasing feature contrast through the use of controlled lighting structure solves the majority of machine vision applications today. The basic theory put forward for Reflectance, Absorbance and Transmission centers on single element surfaces that are conceptually simple and flat. In reality, most parts have more than one material, and are rarely flat or at the same elevation. The illumination engineer may use these variations to tailor the incoming structure of light to either increase or decrease the contrast of a feature with its surroundings.

The Lighting Hemisphere

Lighting structure is solely a function of where the incident light emanates before impinging on the part under inspection. In the case of front lighting, which is our focus here, the total range of incident light possibilities forms a hemisphere which resides above (and is centered on) the part to be inspected (see Figure 3). All possible illumination solutions are provided by projecting light toward the part from one or more (or all) points on this hemisphere.

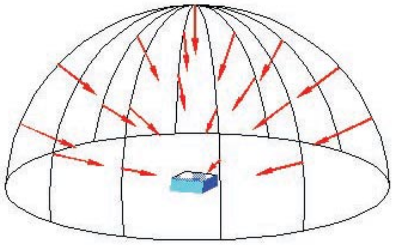

Figure 3: The Lighting Hemisphere. Incident front light comes from any possible combination of points on this hemisphere.

The job of the illumination engineer is to select the proper range of incident light emanation points that maximizes the contrast for features of interest, while minimizing the contrast for non-critical features that may render algorithms unreliable. The infinite range of possible solutions makes this appear to be a daunting task. However, in practice the majority of applications can be reduced to one of three broad categories that are commonly described as Brightfield, Darkfield and Dome (or Cloudy Day) illumination structures.

Brightfield

Let us assume that the part under inspection is flat and has a Reflectance value that is non-zero. The incident lighting hemisphere can then easily be broken into two well-defined zones (see Figure 4). Since the angle of reflection equals the angle of incidence for specular reflection, all light emanating from points within this angle will be reflected by the part back into the camera. Since the background will tend to be bright for a majority of materials, lighting modes which consist primarily of structure which emanates from this part of the lighting hemisphere is called Brightfield. Lighting modes which consist primarily of structure that emanates from the remainder of the lighting hemisphere are called Darkfield because the majority of the light generated is reflected off into space and never makes it to the camera.

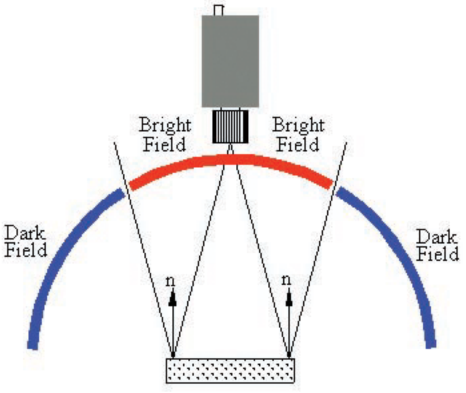

Figure 4: Brightfield and Darkfield define the two basic illumination modes. Brightfield modes reflect back to the camera, Darkfield modes cannot.

Brightfield lighting modes come in a variety of styles, sizes and shapes and provide varying degrees of contrast enhancement depending on the nature of the part under inspection. The purest and most interesting of the Brightfield modes is that produced by what is commonly called a Coaxial Illumination mode, available in many sizes and shapes and is sold under a variety of different brand names.

Coaxial Illuminators produce light which appears to emanate from the detector, bounce off the part, and then return back upon itself to the detector. To accomplish this lighting mode, a beamsplitter is oriented at 45 degrees so as to allow half the light impinging on it to pass through, and the other half to be specularly reflected (see Figure 5). For reflective surfaces, the background signal is very high and uniform over the field of view.

Any variation on the specular surface (reflective, transmissive or absorptive) will result in a reduction in the amount of light making it back into the sensor, causing the image of this area to appear darker than that of the surrounding bright background. This is an excellent lighting mode for flat parts that have a mixture of highly reflective areas surrounded by non-specular absorptive areas. Flat gold pads on a fiberglass circuit board provides an excellent example of the high contrast capability of the Coaxial Illumination mode.

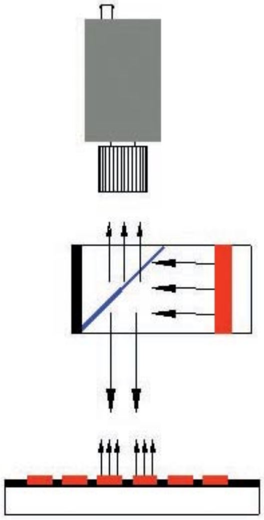

Figure 5: Coaxial Illumination is one of the most common forms of brightfield illumination. It is extremely useful for flat objects with both reflective and absorbing features.

Darkfield

Similarly, if the part under inspection is flat and has a Reflectance value that is non-zero, all light emanating from points below the Brightfield angle will be reflected off the part away from the detector. Since the background will tend to be dark for a majority of materials, lighting modes which consist primarily of structure which emanates from this part of the lighting hemisphere are called Darkfield.

Darkfield illuminators come in a variety of styles, sizes and shapes and are sold under a variety of brand names. They provide varying degrees of contrast enhancement depending on the nature of the part under inspection. For reflective surfaces, the background signal is generally very low and uniform over the field of view.

Any variation on the specular surface (predominantly reflective) will result in an increase in the amount of light making it back into the sensor, causing the image of this area to appear lighter than that of the surrounding dark background (see Figure 6). This is an excellent lighting mode for flat parts that have surface variations or imperfections that deviate from the perfect flat background. Application examples include surface flaw detection (scratches, pits, etc.) as well as OCR on stamped characters.

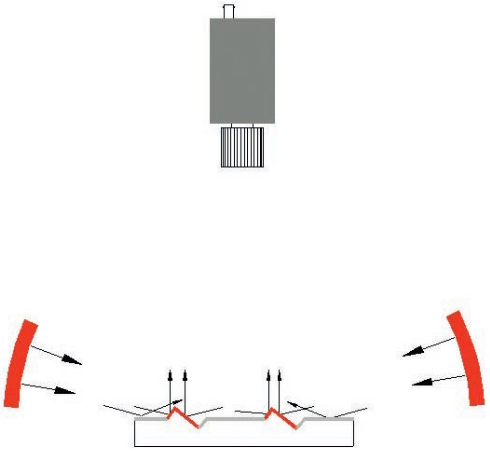

Figure 6: Darkfield Illumination provides low angle incident lighting that highlights any deviations from a perfectly flat surface. Most of the light generated never makes it to the camera.

Dome or Cloudy Day

The Dome or Cloudy Day illumination approach is simply the combination of both Brightfield and Darkfield lighting modes. The intent here is to provide light from all (or most all) points on the hemisphere, which creates perfect shadow free illumination. The result is similar to the effect observed during a perfectly overcast day, where the direction of the sun is completely indiscernible. Light appears to come equally from all angles, and all objects are bathed in light equally from all directions.

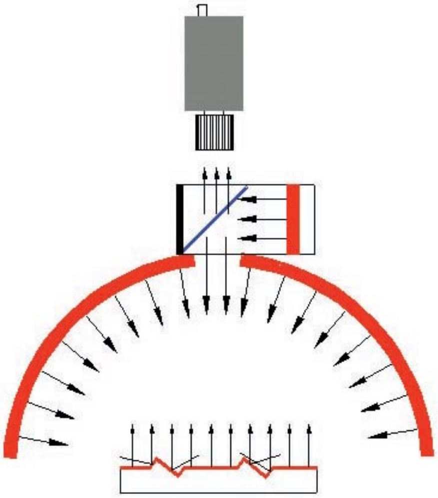

Figure 7: Dome or Cloudy Day Illumination reduces the contrast for many surface variations that may add unwanted contrast into the image.

Depending of the application and the inspection to be performed, the range of possible solutions may extend from simple two mode combinations to the full dome combined with the coaxial illuminator to fill the hold that the detector peers though (see Figure 7). Since many structural details within the scene tend to be "washed out", this lighting mode is excellent for removing unwanted feature contrast that could be problematic for machine vision algorithms. Blister pack inspection provides an example of how this lighting mode removes unwanted contrast created by the clear plastic packaging and waffle pressed foil, while delivering information about the object packed within.

The Basic Assumption

To this point we have made a key assumption that is not always valid. Most of the application work and theory presented assumes that the part under inspection is fairly flat. While this is a reasonable assumption for a great deal of the parts that are being inspected today, there are several groups of parts that violate these assumptions. Highly spherical, cylindrical or domed parts change the general assumption that defines the transition region between Brightfield and Darkfield regions.

In the previous examples we used a flat part to define the transition region between Brightfield and Darkfield illumination zones. To understand how this definition changes for curved surfaces, we only need to realize that for specular reflections the angle of reflection equals the angle of incidence, as measured from the normal vector. In the flat part, the normal is always pointing up and parallel at all points on the part under inspection.

For a rounded part, the normal remains perpendicular to the surface, but because the surface is no longer flat, the direction of the normal varies across the field of view and is no longer parallel at all points. For a surface with a slight convex curvature, this phenomenon effectively increases the Brightfield portion of the hemisphere, and reduces the Darkfield portion (see Figure 8).

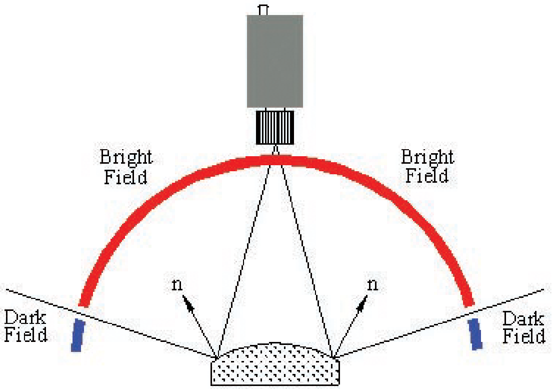

Figure 8: Curved surfaces alter the effective Brightfield and Darkfield regions because the normal vectors are no longer parallel.

As the curvature continues to increase towards a spherical ball, the Brightfield region continues to grow until such a point as the entire hemisphere can now only provide the Brightfield lighting mode. The implication is that any light we provide from anywhere will be reflected into the camera in the form of localized glints and glare which is generally not desirable. A uniform, high quality Cloudy Day type illuminator may be used to reduce or eliminate this highly non-uniform glare, but many times at the expense of washing out the very details that are to be inspected. It is also interesting to note that our ability to provide darkfield lighting modes has been severely diminished, if not completely eliminated, for many of these types of applications.

Spherical Brightfield Illuminators

In the case where brightfield illumination is the method necessary to properly inspect features on a curved surface (spherical or cylindrical), we can use a novel optical approach that is sometimes referred to as spherical illumination using a device generically called a Spherical Brightfield Illuminator (**SBI)**. The goal for the SBI is to provide light in a fashion that all incident rays impinge upon the surface parallel with the normal vector.

In practice, there are several variants for this basic concept. In the case where the part has a spherically reflective dome, a collimated coaxial lighting device may be combined with a convex spherical lens to create a Convex Spherical Illuminator. If the focal point of the convex spherical lens is congruent with the center point of the part's curvature, this device will put a much larger portion of the curved surface into a uniform brightfield mode for inspection.

Schematically, light travels from the collimator to the beamsplitter, where 50% of the energy reflects toward the part under inspection (see Figure 9). The collimated light then passes through the convex lens which has been positioned

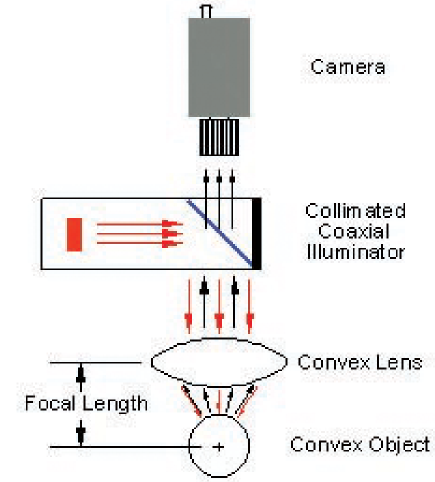

Figure 9: In a Spherical Brightfield Illuminator (SBI) collimated coaxial light passes through a lens element that projects incident light normal to the curved surface.

such that all rays will impinge parallel to the normal at all points upon the curved surface. Since the angle of reflection equals the angle of incidence for all points on the curved surface, the light reflects back upon itself, is re-collimated by the convex lens, and then passes through the beamsplitter where it is then imaged by the detector.

In the real world, there are perhaps more cylindrical applications than spherical ones, but the same principle may be applied with exactly the same results. The only modification to this approach is that the convex spherical lens is replaced with a convex cylindrical lens that is aligned parallel to the axis of the reflective cylindrical part to be inspected.

An example of the effects of such a technique is shown in Figure 10. This reflective cylindrical part has a diameter of approximately ϕ1.5mm, and has small geometrical depressions that need to inspected. As can be seen from the comparative images, the SBI technique increases the uniform brightfield zone by wrapping the incident light around the curvature and re-imaging the resulting light which has been specularly reflected parallel to the normal of the cylinder. In this manner, the protrusions can now be imaged with high contrast against the uniform brightfield background.

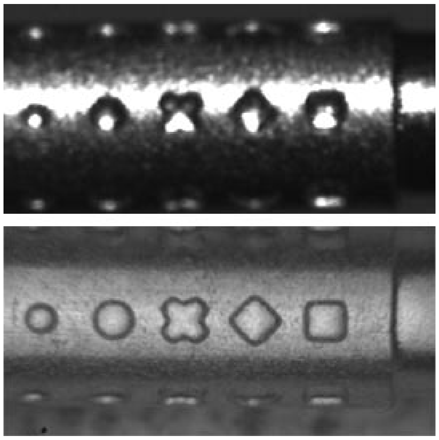

Figure 10: A comparative example of normal Brightfield (upper image) versus the SBI technique (lower image) on a cylindrical part with etched depressions. Note the increased brightfield region and improved uniformity for the SBI.

In a similar manner, reflective surfaces that are inwardly domed or depressed may also be imaged using the same spherical brightfield technique with a slight modification. Here the same basic setup is utilized, but the convex spherical lens is replaced with a concave spherical lens. The concave lens is aligned such that the focal point for the lens and the concave surface are congruent (see Figure 11). Similar to the convex SBI, this configuration puts provides a significantly improved brightfield illumination scheme for the depressed curved surface. Similar to before, the collimated coaxial incident light passes through the concave lens and the rays impinge parallel to the normal at all points upon the concave curved surface.

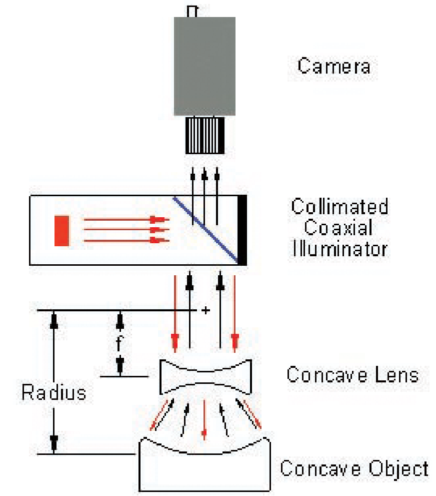

Figure 11: In the concave SBI, collimated coaxial light passes through a concave lens element projecting incident light normal to the concave surface.

The incident light reflects specularly back upon itself (again parallel to the normal for all points on the concave surface), and is re-collimated by the concave lens. The light then passes through the beamsplitter where it is imaged by the detector.

A common application of this technique is for the imaging of soda and beer can bottoms. Here the application requires that the date and lot code is to be read from the bottom of the cans. These codes are generally printed onto the concave bottom of the reflective aluminum can, creating a difficult illumination problem.

Since the inks used tend to be dark in color and are fairly absorbing, this application is one in which brightfield illumination should be the obvious choice (mixture of absorbing and reflective features). The curved surface introduces an additional challenge because normal brightfield techniques fail to produce a uniform illumination field over the entire dome. As can be seen in Figure 12, darkfield illumination provides a uniform background, but fails to create high contrast with the date lot codes. The concave SBI technique increases the uniform brightfield zone such that the black absorbing ink can now be imaged with high contrast against the reflective background.

Figure 12: Concave SBI provides high contrast brightfield illumination for objects that have concave domed or depressed regions.

Conclusion

In the majority of applications, a simple mental review will tend to direct the illumination engineer to use either a brightfield or darkfield approach to create high contrast for the features that are to be inspected. In special cases where surface artifacts or defects are to be ignored, the full dome or cloudy day illumination technique is the technique that may be required.

From the myriad of lighting devices available on the market today, it is surprising that a vast number of applications can be solved utilizing simple brightfield and/or darkfield illumination schemes. These simple principles will help guide any user to select the appropriate lighting scheme to increase contrast and improve overall vision system performance.