To be or not to be in Color:

A 10 year study of the benefits and pitfalls of including color information in AOI systems

Pamela Lipson
Landrex Technologies and Imagen Incorporated
Cambridge, MA 02138
lipson@imagen-inc.com

Abstract
Imagine that you were choosing a camera for your AOI system. Should you opt for a black and white camera or a color camera? This seems like a no-brainer. Most of us would go for a color camera. As humans, we take for granted that our world is in color. Most people see a brilliant color world effortlessly and without much thought to how the brain creates these color images. Cheap consumer color megapixel cameras are popping up everywhere from digital cameras, cell phones, and web cams. In the consumer space, color adds significant benefits to our everyday life from online interactions to saving our memories**.**

However, there are in fact hidden complexities in the choice of using color in the industrial space. Color cameras are typically slower than their monochrome counterparts. Color adds three times the information, which can slow down back-end processing. The quality of the image is highly dependent on the spectral characteristics of the illumination source. Furthermore, most AOI systems need to be able to resolve and measure small features accurately and reliably, and this resolution requirement can be at odds with how color cameras measure color attributes from the world. Even with the inclusion of a color camera in an AOI system, one must address how color can and should be used in AOI machines so that the system can identify defects and measure process attributes. Because of these and other considerations, it is not entirely obvious as to which camera system is best for a given AOI scenario.

There are significant benefits and pitfalls to using color in automated inspection of printed circuit board and the use of color must be thought through carefully. Over the last 10 years we have studied the science, the technology, and the application of color in AOI. In this paper we bring to light important issues about color that are surprisingly little known except to a select community of color scientists. We also show how these issues aid and detract from AOI inspection. Finally we discuss how even with the adoption of color hardware, software can make or break its overall benefits to AOI inspection.

1. Introduction
Choosing a color image capture system for an industrial optical inspection system is not a trivial task. The choice of a camera, color processing algorithms, camera parameters, lighting, and, lens depends highly on what one wants to accomplish with the color image. The choice must be made upfront to decide why color will be used in the system. The camera, lighting, and other choices can be tailored to these needs. In most cases a color image capture system cannot satisfy all constraints at once and usually the constraints need to be prioritized.

The main question to ask for an AOI system is – "Are color images 'vitamins' or 'painkillers'".

The answer to this question makes a significant difference in how to implement a color system, if at all.

Throughout the past 10 years, we have surveyed the customer's perceived need for color processing in AOI systems and we have also looked at, from an architectural point of view, the AOI designer's need for color in an AOI system. Surprisingly, the two needs are completely different.

1.1 Color from The customer's point of view
From the customers' point of view, they, naturally, want to find all of their defects and to make sure the images that are sent to a review or repair station easy to understand. From our 10 year long and continuing study of customers' defects and preferences, we believe there are only two unambiguous reasons for adding color to an AOI system.

- <u>The painkiller:</u> The first is that there are color values on components which need to be checked in order to determine whether the right part has been placed or whether it has been oriented correctly. While this is a necessary and important use of color this type of color verification has usually been lower down in the list of priorities of defects and accounts for a small percentage of the current inspection requirements. Figure 1 shows an example of a part that requires color based verification.

Figure 1: This is a picture of MELF diode on a reflowed board. The color bands indicate its value. Color processing is important in this example to know if the correct part has been placed. (Note that this part has a tombstone defect which caused a solder defect on the right side)

- <u>The vitamin:</u> The second is that color images make it easier for a human operator to relate the image shown by the AOI system and the actual board under test. While it seems logical that a color image could help operators do their jobs, humans are very flexible and with time can learn the correlation between grey scale images and real life boards. Figure 2 shows an example of a color review station, which customers often ask for when looking at an AOI machine.

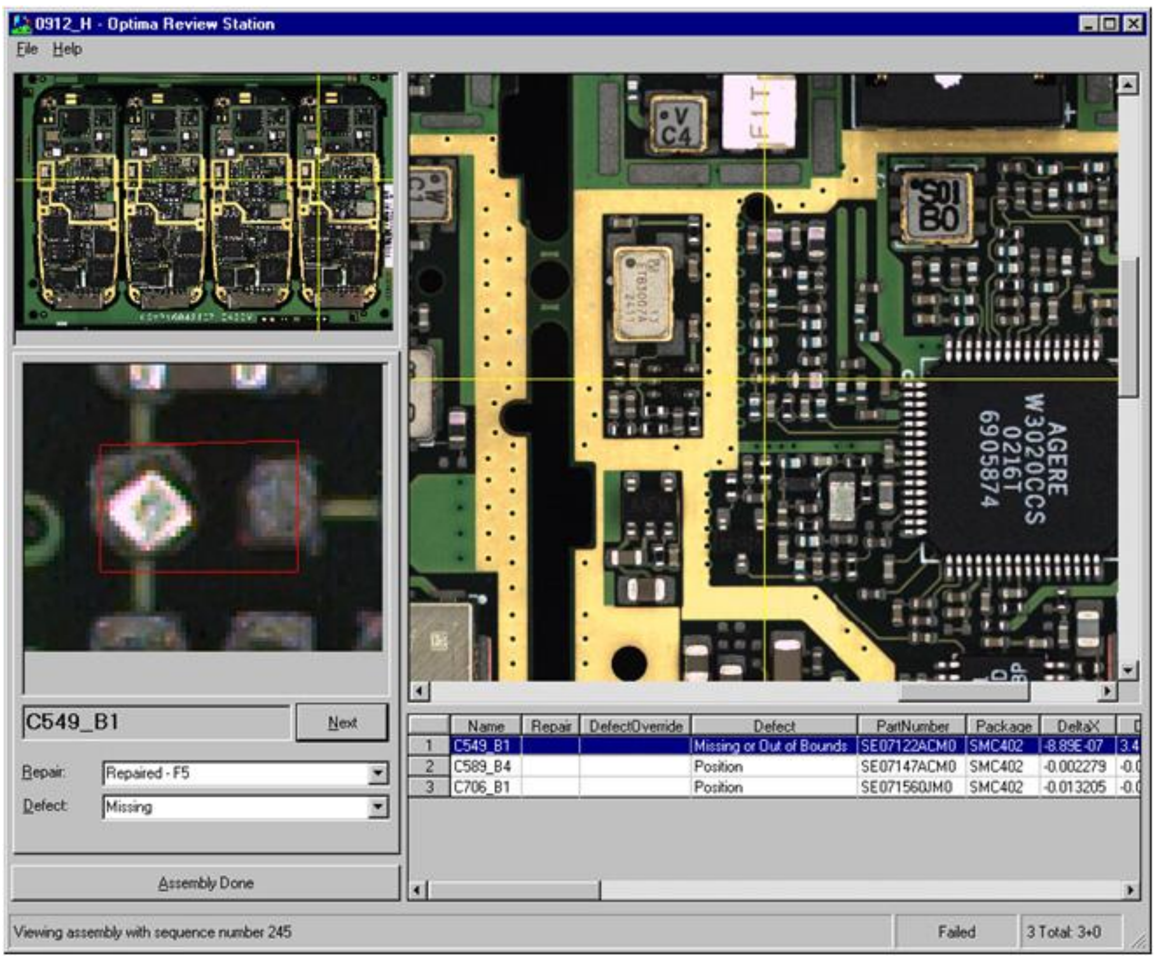

Figure 2- Example of a color review station. This particular review station shows the defects and their position on the board in color. It has been reported that color interfaces make it easier for a human operator to relate these images to the actual board.

1.2 Color from the AOI architects point of view:

From the AOI architects point of view the need for color to inspect PCB boards is not clear cut. The use of color is just one element which makes up the module known as the "image capture system". The image capture system includes one or more cameras, lenses, and lighting modes. Additionally, while the camera and the lighting provide the images, the choice of camera and lighting dictate the type of image analysis algorithms that one can use.

The simple addition of a color camera can complicate the other hardware in the image capture system (the lighting and lens), the software (the image processing), and may compromise the resolution of the system. However, in some cases, the addition of color, when implemented properly, can actually simplify the system architecture, reduce the cost of the AOI machine and increase the ease of use of the system.

<u>Color is one way, of many, to achieve AOI objectives:</u> The majority of AOI systems are used to determine solder defects and part presence and absence. For both parts and solder joints, it is predominantly the geometry of the joint and the part that causes a defect, rather than its color or absolute appearance. In fact, without the proper image processing, color and absolute appearance can be confusing to an AOI system. The geometry of solder joints and part presence/absence can be inspected by a variety of image capture setups. The most common, up until now, include systems based on multiple grayscale cameras and multiple lighting modes. More recently AOI systems with single color cameras and one or more lighting modes have gained in popularity.

The main point is if it is the geometry of the object that one wants to inspect, color is one option of many to achieve that goal.

In the rest of the paper, we'll talk about the mechanics of taking a digital picture and achieving a color output. We'll discuss the impact of lighting and lenses on a color image. We'll also discuss how to return to grayscale values from a color image. Finally, we'll discuss how color can help or deter AOI systems can from detecting defects.

2. Commercial color cameras versus industrial color camera

Today in the consumer world it is difficult to find an electronic device that does not take a color image. Color cameras exist as standalone devices. They are also embedded in our cell phones, pdas, laptops, webcams, and numerous other devices. We can share, upload, and print our images literally with the snap of a few buttons.

Given the prevalence of color digital images in our society, to the general user of industrial equipment, an obvious question is "Why wouldn't my AOI system have a color camera?"

However, the requirements for a consumer color camera and an industrial color camera are quite different. In the consumer sector, companies like Kodak and Sony spend millions, if not billions of dollars, in color science research to make images that are pleasing to the human eye given the variety of light sources that are found in most social situations, (e.g. daylight, nighttime, indoor lighting). They have focused on making capture devices that highlight important colors in the world such as skin tone, sky blue, and green grass. In addition, they have created images with color gamut's that can translate well to popular printer and screen gamut's (meaning images that look good on the screen as well as the printer.) Further work has gone into ease and convenience of taking a color picture and software to reduce artifacts, such as red-eye.

The requirements for a color camera in the industrial setting are completely different. A pleasing image to the human eye is not really a priority. A color image which can illuminate the defects and a color capture system which can meet the beat rate of the inspection task are the true priorities in the industrial space.

Prior to the year 2000, there were surprisingly few options for industrial color cameras. There were two reasons for this. The first was due to the speed requirements of industrial cameras. Images needed to be taken quickly. One needed to flood the image with a lot of light so that the camera can have a short exposure time. At the time, it was very difficult to provide enough white light to take a picture at the required speed. Thus, people mostly employed red LEDs which were bright, strobable, and reasonably inexpensive. However, red LEDs didn't allow for a good full spectrum color image and, therefore, a color camera was unnecessary. The second reason is that transfer bandwidths, computer speeds, and computer memories were not adequate to collect, store and process large amounts of color images. On the processing end, most image processing in industrial systems was actually done on binary images to reduce memory and computational costs.

In 2000, camera makers started to see that both strobable cheap white light sources might become a reality and also that computers were becoming faster and endowed with more memory. Thus, they started to adapt their current grayscale cameras to provide color. The first versions of these were expensive, slow, and did not give very good quality pictures. The reason is that adding color to a good grayscale camera is not an easy proposition and turned out to take much more trial and error engineering to meet the customers' needs.

Most of these color solutions also required special processing beyond the image capture. Thus, cooperation with either frame-grabber companies or stand-alone hardware was required to make real color images a reality in the industrial sector.

Fast forward to now, color cameras have advanced significantly. At about 60 frames per second, they can meet the speed of monochrome cameras. Connection to the pc via protocols such as cameralink, Ethernet, and 300baseT make it possible to transfer the images quickly and facilitate the integration of the camera into the system. However, even with these advances, in the industrial world, monochrome cameras still outsell color cameras. In the rest of the paper we will discuss some of the outstanding issues with integrating a color camera into an industrial system.

3. The basics of grayscale cameras

Prior to entering the discussion on color cameras, this section reviews the basics of grayscale cameras. Most grayscale cameras are 1 sensor cameras. There are two types of grayscale cameras – CCD cameras and CMOS cameras. While CMOS provides some advantages, CCD cameras are still more popular due primarily to their superior image quality. CCD cameras come both in a line scan version and also an area array. The most popular are area array cameras. Therefore, we'll focus on the CCD array camera for this discussion.

A CCD camera uses a grid array of silicon photo-sensors to receive incoming light. These photo-sensors are charge-coupled devices. Each sensor is one element of a picture. In the grayscale camera, these elements translate directly into the pixels of a digital image. The CCD sensors release electrons when hit with a photon of light. During the exposure time, the sensors accumulate electrons as they are being hit by the photons.

When the exposure time is over, the camera reads out values from the CCD and reports the voltage for each photo-sensor. (The voltages are digitized at a later stage either on the camera or framegrabber.) Once the information has been sent, the sensors are "reset" to take another picture. Thus, the more photons that hit a particular sensor, the lighter that pixel will appear in the resulting digital image.

While this sounds quite simple, industrial camera makers faced a lot of challenges. The first challenge was to get enough light into the individual sensors so that one can take a picture in a short period of time. Some solutions to this problem were to add micro lenses or micro mirrors around each sensor. The second challenge was to read (transfer) the signal from the sensor so that it could be sent to the camera or the framegrabber for digitization. Companies spent a lot of time to create electronic circuits that could do this quickly, robustly, with a small form factor, and low power consumption and heat emission. There are several different types of architectures for collecting and transmitting these signals. While the differences between these architectures were important early on, most of the current CCD architectures can perform the required tasks. The third challenge was to get the sensors, lenses, and circuits in a small enough form factors to make the cameras practical. The fourth challenge was to be able to make the electronics cheap enough so that the camera could be offered at a reasonable price. The fifth challenge was how to get the image collected by the camera to the host computer. This challenge, at the time, could not be solved alone, but, up until now, required a frame-grabber, which would interface with the camera, collect the images (and possibly digitize them), and transmit them to the host computer. Just the simple act of getting a camera, frame-grabber and computer to work together was a daunting task on its own and deserves a separate discussion. With the emergence of new transmission standards, such as Ethernet and USB3, the frame-grabber plays a lesser role and may eventually become unnecessary.

Irrespective of the architecture of the camera, there are several parameters of the quality of the image that are important to measure. These parameters are important both for grayscale and color images. These include signal to noise ratio, the total amount of noise often called dark noise or noise when no light enters the camera, dynamic range, ease of saturation, uniformity of the image when given a flat field, number of bad pixels (pixels that simply don't work), and the depth of the data coming out of the camera (8 or more bits).In addition, there are some other technical issues to consider including the transmission protocol from the camera to the frame-grabber or computer, the amount of storage on the camera, and any on-board computational ability in the camera.

4. The Science of Color

Let's now move on to the color part of the discussion. The science of color is a very interesting topic in its own right:

Most people assume that if one can make a camera that can perceive light that it can also perceive the color (or the individual wavelengths) of the light. The reality is that we need to trick grayscale cameras into seeing colors

The main point about light is that the visible spectrum of white light contains a range or spectrum of colors. This range of colors is usually not obvious to the human observer unless we are looking at an artifact like rainbow which demonstrates the multifaceted colors of white light.

Figure 3 - Visible Light is made up of a multitude of colors. Generally, we are aware of this only when looking at a rainbow.

It turns out that the range of colors present in the visible spectrum of light can be approximated by only a few dominant colors. These colors can be combined together to make other colors. Because cameras by their nature, as described in the previous section, collect light emitted from a source, they belong to something called an "additive color system". An additive system adds colors together on top of black to produce other colors. The primary colors in an additive system are Red, Green and Blue. In combination of sets of twos, they can produce the secondary colors Cyan, Magenta, and Yellow. Adding the three primary colors together produces grey or white. Figure 4 illustrates these phenomena.

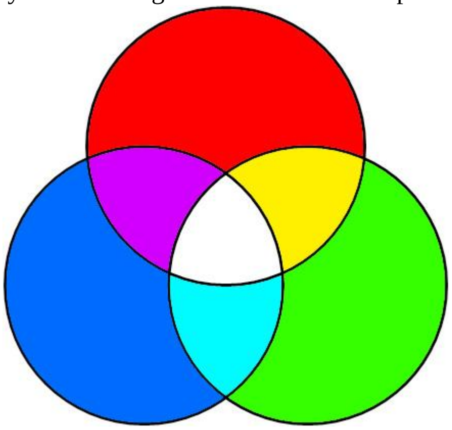

Figure 4- Additive Color Mixing. Three primary colors (red, green, and blue) can be combined to create a variety of other colors.

(Note this is quite different from a printing process, which is a subtractive process. This process subtracts colors from white to create other colors. The primary colors in this system are Cyan, Magenta, Yellow and Black.)

Thus, in order to get color pictures of the world using a camera and a broad band light source, one must get a minimum of a red, green and blue measurement for every pixel in order to try to reconstruct that pixel's "true" color.

5. Color Cameras

There are two types of color cameras that are generally available. Each attempts to gather a red, green and blue measurement of the world. Both types are built upon the traditional CCD grayscale camera that we discussed earlier.

5.1 Three CCD color cameras

This type of color camera uses three CCDs; each CCD is devoted to one of red, green or blue. An image comes into the camera and is sent through a trichroic prism assembly. This prism assembly converts the light into the three different colors and each color is directed to a separate CCD. The resulting images from the three "color channels" are processed separately and then combined into a final image. Figure 5 illustrates how a three CCD camera works.

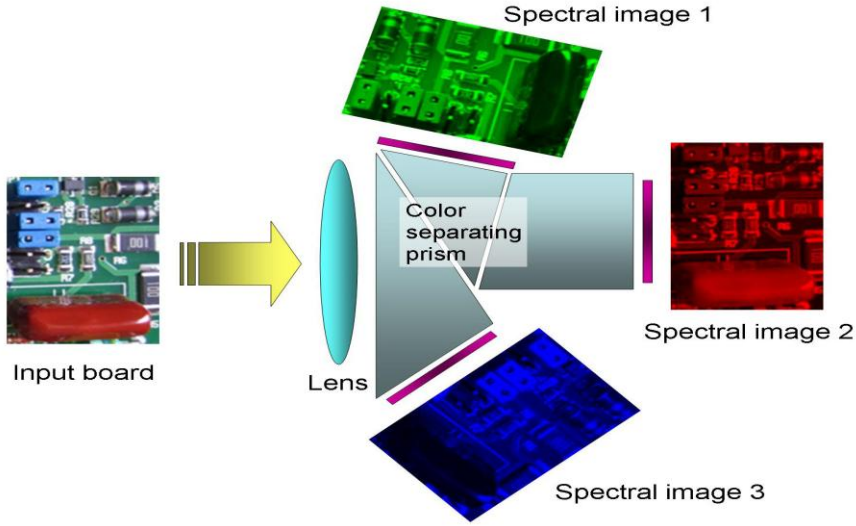

Figure 5- Architecture of a 3-CCD color system

Three CCD cameras have not gained significant traction in the industrial market. Early iterations of the cameras had issues with robustness and integration of the cameras suffered from insufficient bandwidth to quickly transmit the three color values for each pixel. On the robustness issue, the CCDs and the prism have to be exactly aligned in order to properly split light for processing by the three CCDs. Thus, the setup has to be perfect and the alignment must stay that way through the lifetime of the camera. While these design problems have been primarily eliminated and bandwidths have increased, there are two main problems preventing broad adoption of the three CCD cameras. The first is cost. Three CCD cameras can be five times more expensive than their alternatives. Additionally, they require special lenses that are tailored to the optics of the camera, which is not a trivial issue.

However, the advantage of a three CCD camera is that the image resolution and color fidelity is quite good because three separate measurements of color are being made for each individual pixel.

5.2 One CCD color-camera

The second solution adapts a regular single CCD camera into a color camera by using the single CCD to capture the red, green, and blue components of the stimulus. A popular method employed is to put an array of color filters over the sensor elements. The most common arrangement of these color filters is called the Bayer pattern. The Bayer pattern uses one red, one blue and two green pixels to form an imaging block of four elements. This pattern of four elements is repeated over the entire CCD array, as illustrated in figure 6.

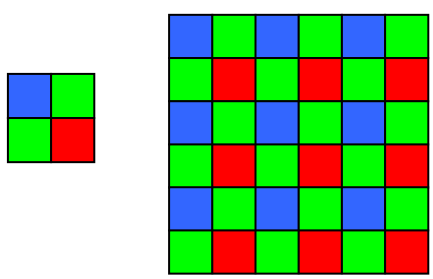

Figure 6- The Bayer Filter is made of blocks of 1 red, 1 blue, and two green elements. This pattern is repeated to cover the entire CCD.

While this arrangement allows one to get a color picture with only one CCD, the camera is not capturing an independent color value for red, green and blue at each sensor element. Thus, in order to get a full color picture, we must derive two color values at each pixel from the neighboring pixels. This process of filling in the color information is called "color recovery".

In summary, the advantages of a one CCD color camera are cost, robustness and speed. The disadvantage is that one is not getting a true color measurement at each pixel. In addition, additional computation resources are required to fill in this missing information.

5.3 Are all CCDs made equal? Our evaluations:
Surprisingly, not all CCDs with Bayer patterns provide good color images. In our tests of several brands, we have found that only a small proportion actually yields a good range of colors for our inspection needs.

Our criterion for evaluating color was the following:

1) We wanted to see a broad range of colors expressed by the camera. One good way to check this is to look at color targets such as the Macbeth color chart which has a wide array of calibrated colors (see figure 7). This chart is commonly used to evaluate color reproduction systems. A second good way is to look at PCB boards and PCB parts that have a large range of colors. Figure 8 and 9 show examples of these respectively.

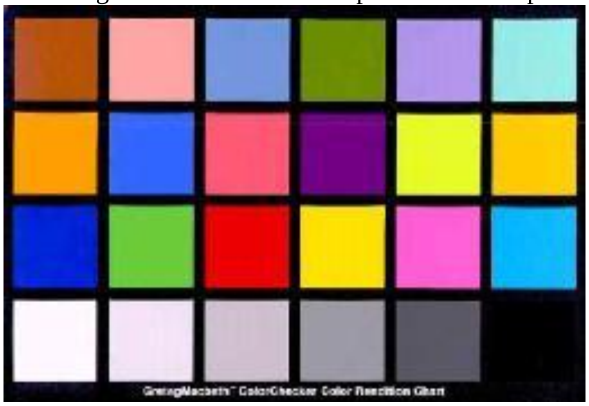

Figure 7- Macbeth color checker chart.

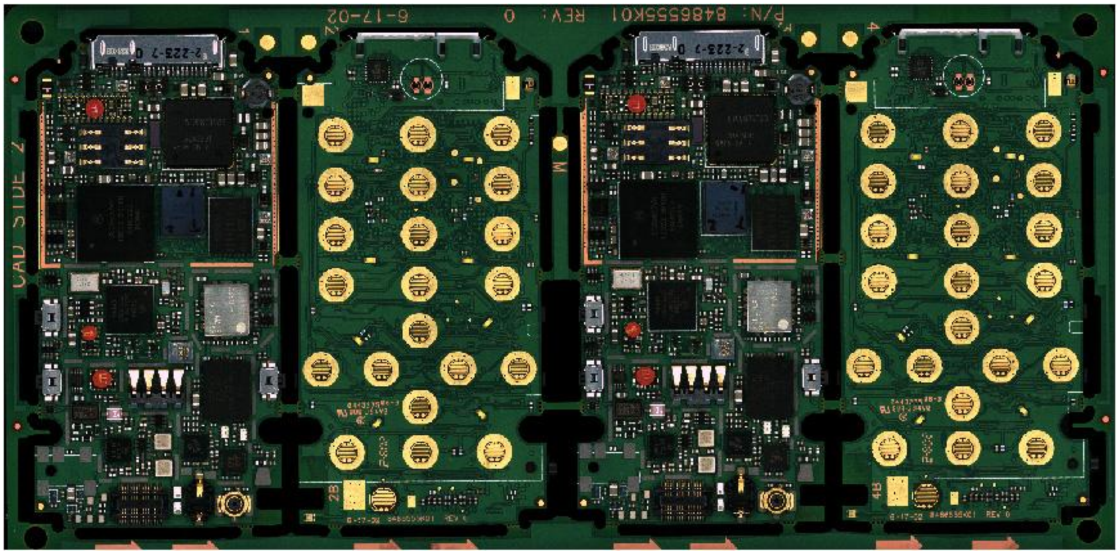

Figure 8- Example of a PCB board with a range of typical colors.

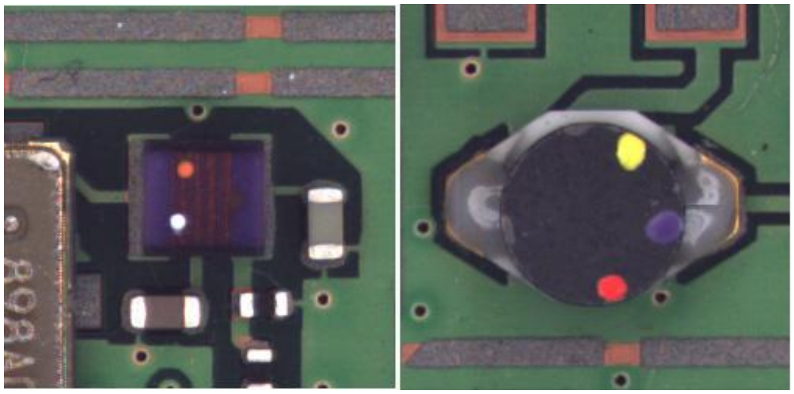

Figure 9- Example of components with a variety of colors that are either intrinsic to the part or are due to markings that are added on in a later process.

2) We wanted to see a good separation of board colors and part colors. We found that the most difficult case in the PCB domain was to distinguish dark (black) parts from a dark green board. Thus, we focused on color cameras and color processing that could boost this signal difference as much as possible. Figure 10 shows examples of dark parts with adjacent dark backgrounds. Without good color separation of the board and the part, it is hard to distinguish the body of the part and the background and the edge transition between the part body and the board.

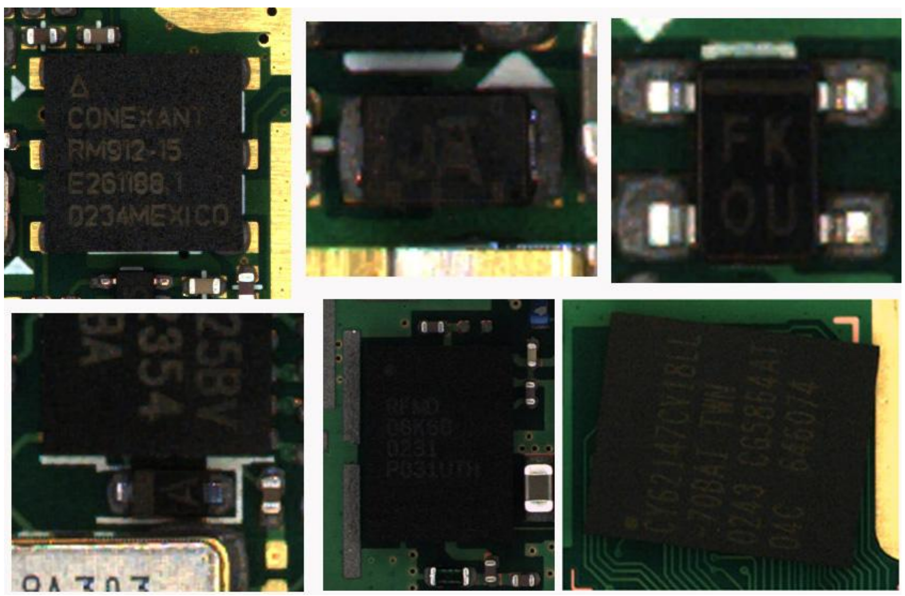

Figure 10- Examples of dark parts on dark backgrounds. We required a color system that it could distinguish the transition between the part and the board.

3) We wanted to make sure we could tell the difference between red and yellow so that we could see the stripes on tantalum capacitors in order to inspect for correct orientation. (See figure 11 for an example.)

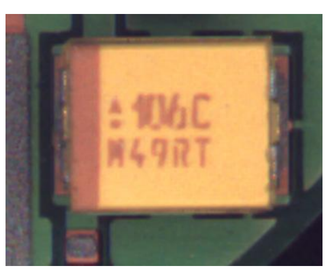

Figure 11- Example of good color discrimination for a yellow tantalum capacitor with a red orientation mark.

4) We needed good dynamic range in order to see the difference between end cap luminance and paste luminance. In addition, we were looking for separation between end cap color and pad color. Figure 12 shows several capacitors and resistors in addition to bare pads. In this picture end caps are very bright and easily distinguishable from the paste, which is dark grey, and the pads, which are gold.

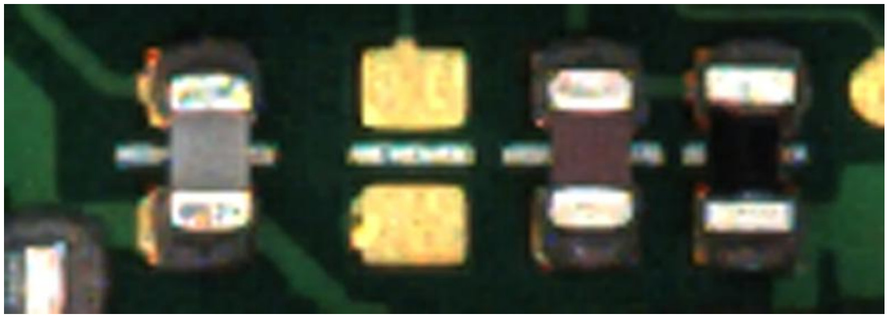

Figure 12: Example of good color discrimination between the end caps, paste and pads.

At the time we evaluated the cameras there were two dominant 1 megapixel CCD sensors in the market. Contrary to our initial expectation that a CCD is a CCD is a CCD, we found that significant differences exist between different brands. We found that one company's sensor performed significantly better across all of our criteria and also provided more pleasing images to the human eye. We reevaluated CCDs in 2001 and again in 2005. In our most recent tests, we found that while the gap between the two dominant sensors has become closer, the original one still had advantages over the other.
The moral of this story for us was to give up on the idea that all CCDs are created equal. Before making a purchase/deployment decision, each brand needs to be rigorously evaluated on its own merits.

6. The Science of Color Recovery

Assuming that one chooses a one CCD Bayer filter camera to take the color images, one must choose both an algorithm to implement the color recovery and the mechanism that will actually do the computational work (e.g. frame-grabber, onboard camera resources, host CPU).

Color recovery is the real "gotcha" in the process of choosing a color camera because the color recovery, for each pixel, must create 2 colored values to fill-in the missing data.

When we started our investigation, the process seemed like a fairly simple task. However, when we started evaluating the different methods we saw that the color recovery process is actually quite complex. Additionally, because we were developing a metrology machine which performed subpixel measurement, we required good color image quality at every pixel.

The simplest way to do color recovery is to treat the whole imaging block (1 red pixel, 2 green pixels, and 1 blue pixel) as one pixel, and average the 2 green channels. However, this provides much poorer resolution than that of the equivalent monochrome CCD.

If one cannot afford to reduce the number of pixels from the CCD by 4, various color-recovery algorithms attempt to restore the missing resolution by interpolating between the pixels of the same color. This provides RGB data at every pixel location. However, the color recovery is inherently difficult because it is trying to infer values that do not exist. Color recovery algorithms vary in their computational cost, complexity, and ability to faithfully reproduce colors and edge detail.

There is no perfect color-recovery algorithm. Irrespective of the complexity of the algorithm, there will always be artifacts in the image. Saying this another way, there will always be errors in how the system infers the values of the missing pixels. Figure 13 shows one example of a common artifact from a one CCD Bayer pattern camera, called color fringing. Other common problems are zipper effects, where a single row of pixels alternates between two colors or gray levels from one pixel to the next and rounded corners, where a false color replaces the real color at the junction of two lines. Artifacts can be magnified when the edges in the image are not perfectly horizontal or vertical. Diagonal edges, from synthetic patterns, and non-sharp edges, as found in real world images, can magnify the errors from a color recovery algorithm.

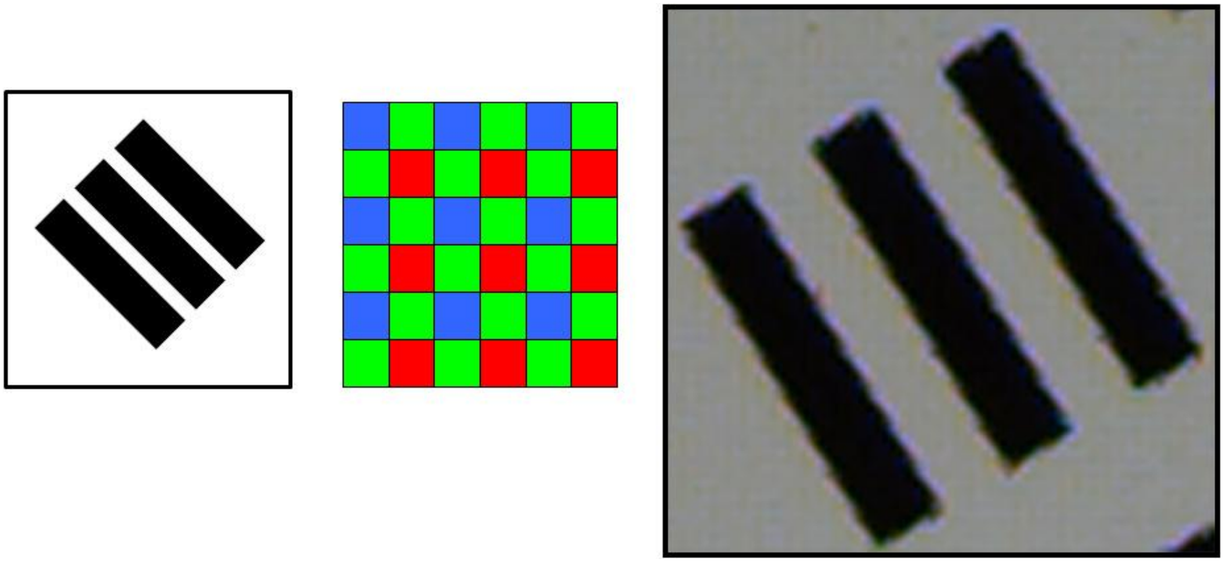

Figure 13- The left image shows the original target (3 diagonal bars). When imaged through a Bayer filter one CCD camera and processing with a color recovery algorithm, the resulting image can have significant artifacts. The purple lines around the diagonal bars in the image on the right are one type of typical error that results from this process.

In summary, it there is a significant tradeoff between trying to get good color rendition and good resolution. It is very difficult to achieve good performance in both aspects.

6.1 Details of how to evaluate color recovery algorithms

Spotting issues with color recovery algorithms requires careful evaluation of a target with multiple orientations, multiple sizes and multiple different spacings between elements. We used the USAF-1951 target as our first mechanism for evaluating the color recovery algorithms. Figure 14 shows an example of this target. This high contrast, chrome on glass target contains groups of 3-bar patterns of different sizes, each with equal bar width and spacings, from 500 microns bar width to 2.2 micron bar width. We took images of the target both when aligned with the camera sensor and also at a 45degree rotation offset from the camera sensor orientation. From these images we can see where the artifacts are and at what resolution does the image quality break down, termed "usable resolution". We also took images of the same target in small subpixel and pixel offsets from the camera array to see the effect of where the target elements hit the array.

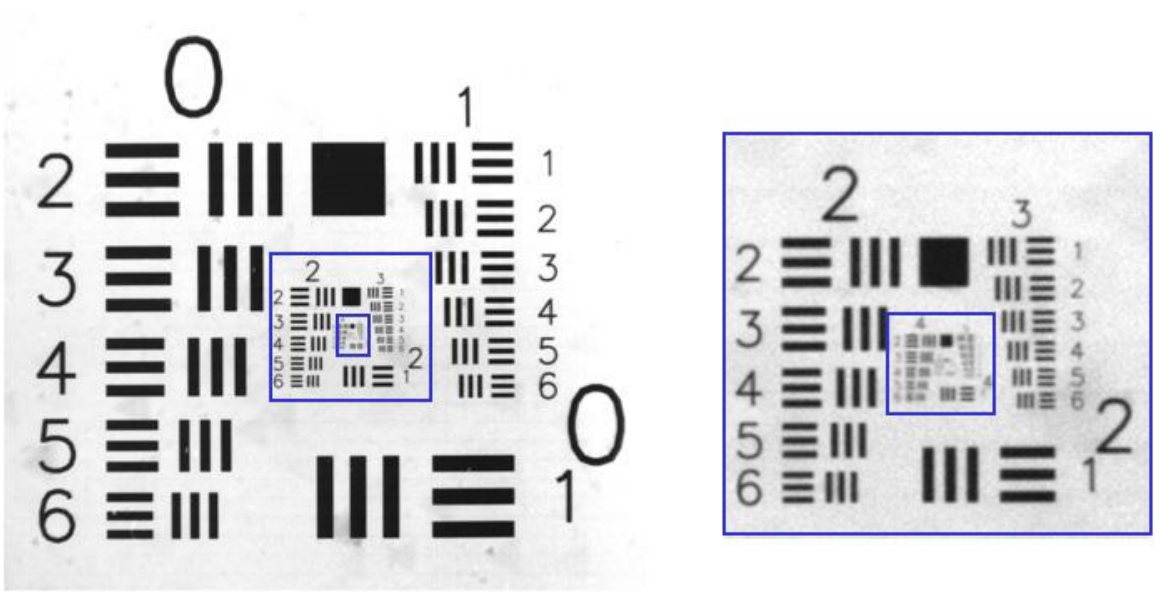

Figure 14- Example of the USAF-1951 target which is typically used to measure a variety of optical instruments. The left image shows the full target. The right image shows the area used in the subsequent analysis of the color recovery algorithms. The middle section of the right image proved the most useful in investigating the color recovery algorithms for our pixel size.

We also tested each candidate color recovery algorithm algorithm on images of a real board with components of parts in paste. We made sure to take pictures of boards with a range of different color and luminance components.

Figure 15 show some images of the USAF-1951 target that are taken with a single CCD Bayer filter camera and no color recovery. In these images and subsequent images in figure 16 the figures focus on the innermost pattern in order to see the issues with the color recovery algorithms. The first image in Figure 15 shows the non color recovered Bayer image. The checker pattern is evident. The second picture shows the green channel only. The third picture shows the blue channel. (The red channel, not shown, is similar in appearance to the blue channel.) It is clear the resolution in the green channel is significantly greater and the artifacts significantly less than in the red (and blue) channel.

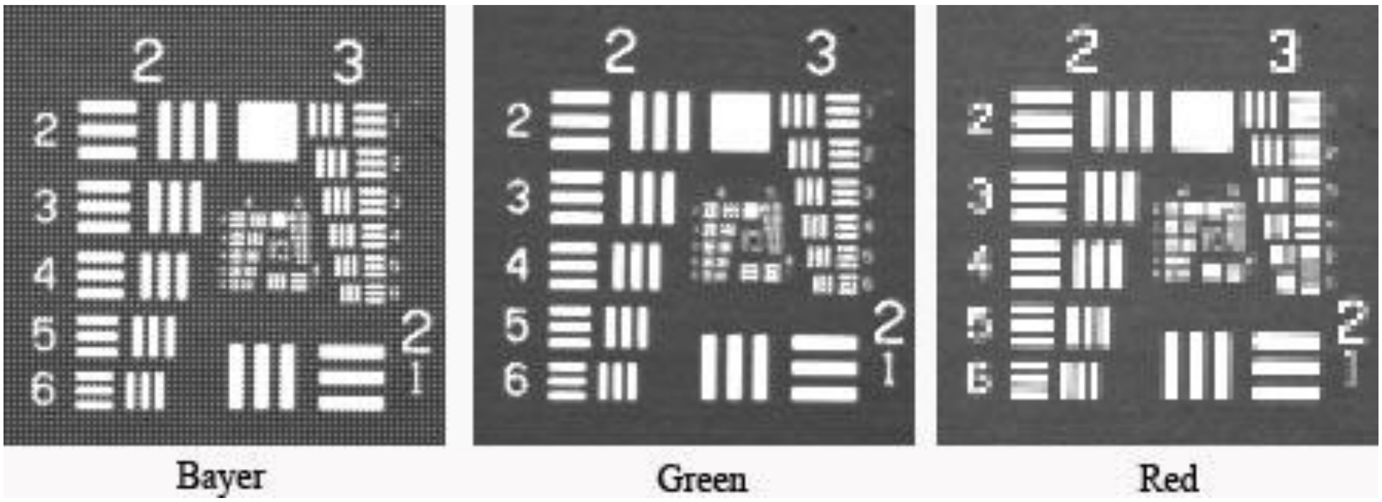

Figure 15- This figure shows the difference in resolution and artifacts between the green channel and the red channel. These differences make color recovery of all three colors an even greater challenge.

Figure 16 shows the color reconstructed image with a variety of different algorithms. The main point to take away from these images is that different algorithms can have very different outcomes using the same input data. The “Base Algorithm” has color fringing on all of the stripes and the “usable” resolution degrades at level “3 2” (the two numbers refer to the column and row indices, respectively, of the subpattern). Using algorithm “Improvement 1”, the color fringing is less, however, the “usable” resolution level is roughly same. With algorithm “Improvement 2”, there is very little color fringing until levels “3 4”. The “usable” resolution also is significantly improved. It should be noted that the more sophisticated the color recovery algorithm, the more computation resources are required.

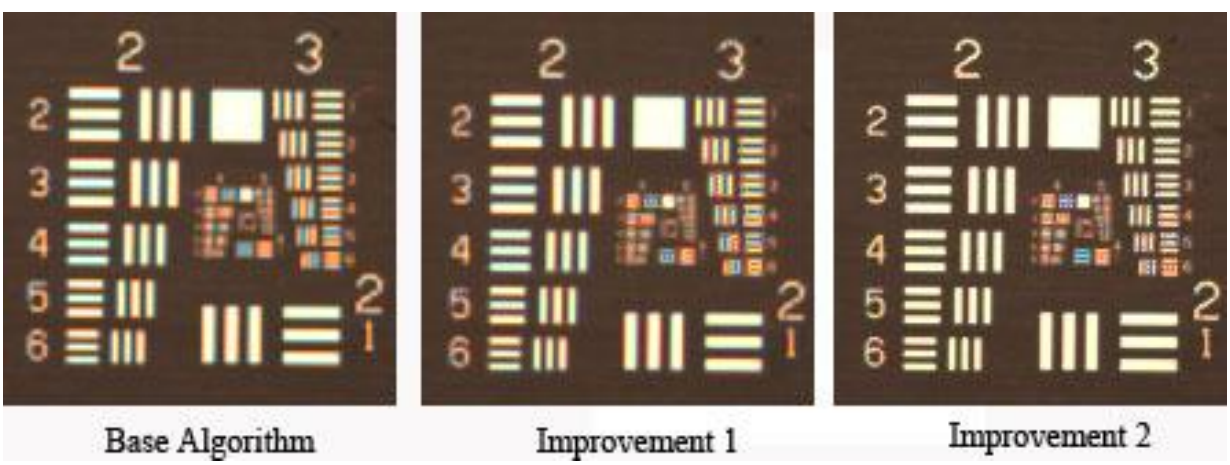

Figure 16- These images demonstrate the difference in color recovery algorithms. The Base algorithm uses a simplistic recovery method, but produces significant color artifacts. Algorithm “Improvement 2” reduces the color artifacts and increases the resolution at the cost of processing time.

Even with the best algorithm, we still see some artifacts in real images of printed circuit boards. Figure 17 shows a four color images with a 33 micron pixel size. The images illustrate some of the unresolved issues with color processing such as false colors, zipper effects, and edge and corner artifacts.

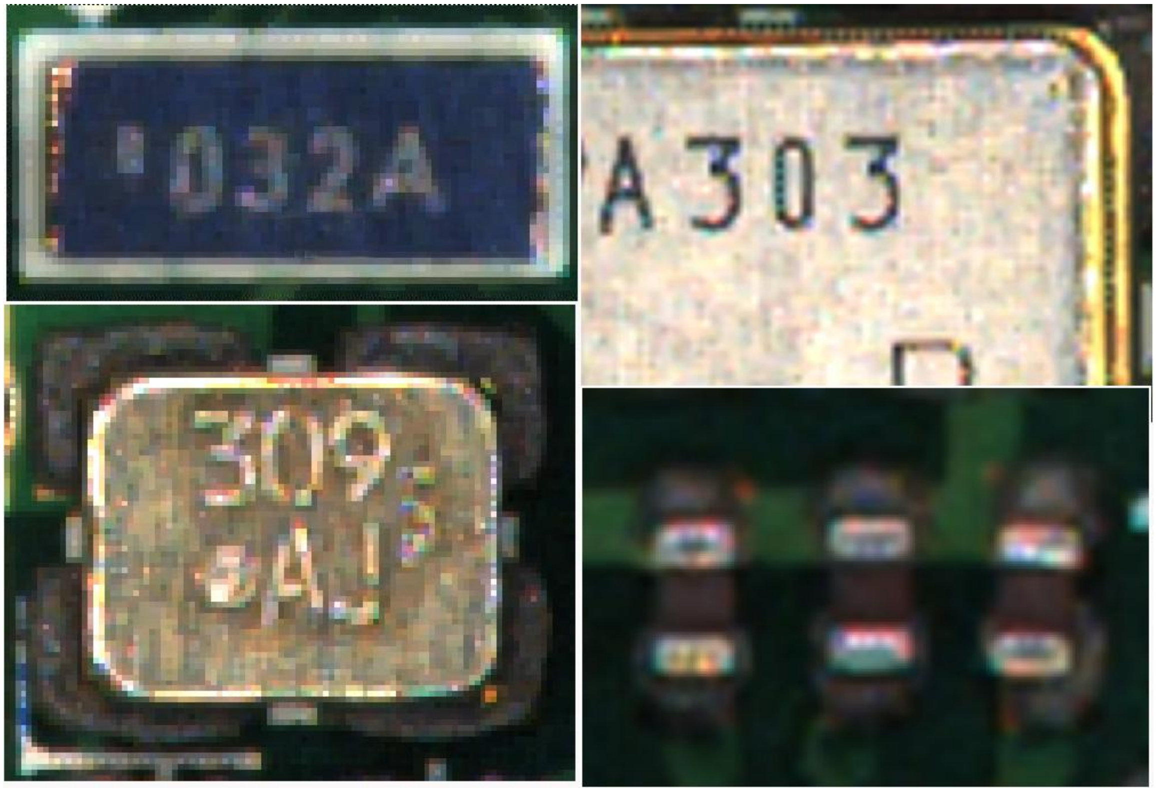

Figure 17- Examples of color artifacts on real parts on PCB boards. These images were created with one of the most sophisticated color recovery algorithms. This demonstrates that color artifacts cannot be avoided and, therefore, must be accounted for in the inspection system's subsequent image processing.

6.2 Where should we implement the color recovery?

In principle, there are three options for where to implement the color recovery. The first is to do the recovery in hardware on the camera itself. The second is on the frame-grabber or programmable card. The third is in the host CPU. When we first started working with color images, option 2 seemed to be the only realistic solution. Now with cameras and computers with more computational power and memory, options 1 and 3 are possible.

If the color recovery is done at capture-time either on the camera or on the frame-grabber, then the operation is a one-shot deal: the same image must be used for all types of image-processing operations and we have to optimize every aspect of it as best we can. However, if the processing is done on the host computer, it is possible to perform color-recovery "just-in-time" or "on-demand. The color-recovery algorithm can be selected to optimize the aspects that are important for a particular use, for instance, pleasing color at low resolution for some cases, high resolution and high fidelity color for other cases. By using on-demand processing, color images could be stored in monochrome format, allowing instant 3x compression. Color recovery that is done "on demand" allows for the system to optimize the color recovery based on what the system is inspecting.

7. What about lighting?

7.1 Lighting color

Most people who consider adding a color camera to an AOI system focus mainly on the technology of the camera and the mechanics of the color recovery. However, a significant piece of the equation is the lighting used to illuminate the board. If one wants to be able to capture colors over the entire range of the visible spectrum, one must have a full-spectrum light source. Lighting itself acts as a filter on the stimulus which the camera is viewing. If the lighting has a narrow spectrum, such as a red LED, then the range of colors one can measure are substantially reduced.

Figures 18 illustrates the spectrum of two different broad spectrum light sources. The top image shows the spectrum of natural light, which is our ideal light source. The bottom image approximates spectrum of one type florescent light. While some florescent lights are similar to the natural light, it is not possible to strobe florescent lights. Thus, they can only be used in "stop and go" applications which significantly increase image capture time.

Up until 2000, the only strobable alternative that approximated natural light was xenon strobe. It was very effective at producing a lot of light in a small area. However, its size, heat and power issues made it an unattractive option.

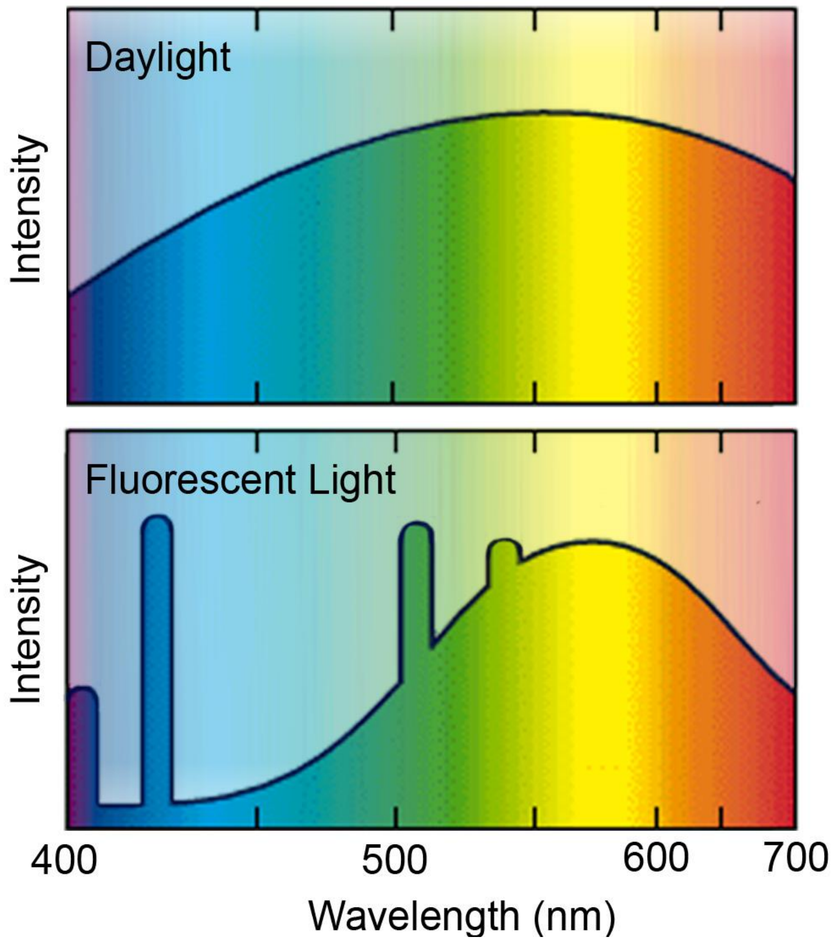

Figure 18- The top image illustrates the broad spectrum of the natural. The lower image illustrates an example of man-made (florescent) light designed to mimic natural light. Ideally, for an industrial color application, we would want our lighting source to have a wide spectral range.

Prior to 2000, the only practical strobable light was the red LED. One curve in Figure 19 shows the very narrow spectrum of the red LED. This light source would highlight anything on the board that had a red component or a reflective nature. Other areas of the board would turn black. If the PCB boards were green and the parts on the board were somewhat reddish, the lighting had the advantage of highlighting the parts while de-emphasizing the boards without any additional computational (software) processing. However, because of the light's narrow spectrum the red LED renders a color camera ineffective at taking a color picture.

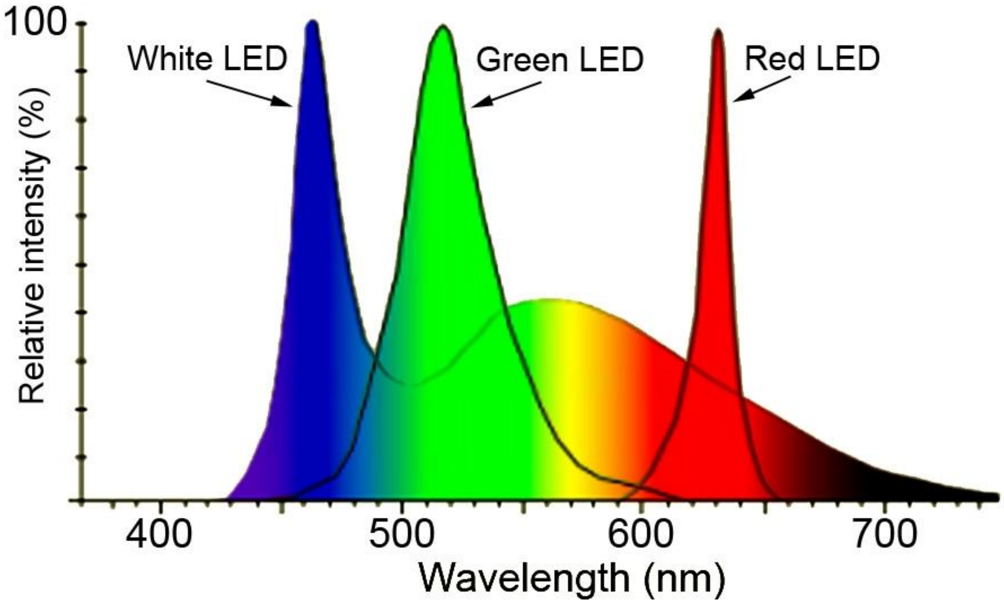

Figure 19- This image shows the spectrum of a typical red LED, green LED and the "white" LED. The red LED and green LED have quite narrow spectrums. The most commonly used white LED is composed of a blue LED coated with a phosphor. The result is far broader spectrum than either the green or red LED, making it a good choice for industrial applications.

In early 2000, the creation of the first bright white LED changed the face of the industrial lighting field. The most common form of white LED really isn't white. It is constructed of a blue LED coated with a phosphor that, when excited by the blue LED light, emits a broad range spectrum in addition to the blue emission (see figure 19). The result is a fairly white light. (One alternative to this type of LED is a white LED that actually has blue, red, and green components in it.) The merits of the white LED are significant. They have a reasonably broad spectrum. They have become reasonably cheap. They are bright, strobable, long lasting and a group of LEDs can be arranged into almost any form factor. Thus the white LED is a very desirable light choice.

At the same time as the white led became available, advances were made in the intensities of the green and blue LEDs. (See Figure 19 for the green LED spectrum). Some AOI lighting designers found that they wanted to be able to control the shape of the light spectrum depending on what they are imaging. Thus, it is possible to create a lighting head composed of red, green and blue LEDs. If one composite picture is desired for processing, adding the spectrums of the red, blue and green LEDs gives a good approximation to the broad spectrum. However, because these three are narrow spectrum sources, there still will be some holes or gaps, which will cause some colors to be under-represented in the resulting images. However, this architecture allows one the flexibility to weigh these lights differently depending on, for instance, the color of the PCB board. (Also one can take separate pictures of the board under test using separate light colors and use a grayscale camera to create a false color image. The latter gives the combination advantage of the resolution of a grayscale camera and a color picture.)

We are in the midst of this LED revolution. LEDs are still evolving and changing in a number of ways including the design of their semiconductors, thermal packaging, and collection optics. The future for LEDs looks very exciting and we expect their wide scale adoption in almost every lighting arena.

7.2 Lighting direction

The spectrum of the light source is not the only factor that affects the quality of the image. The direction of that light source also has a large effect on the content of the image. Some elements of a PCB board image well under direct on-axis light. Other elements image better with light that comes in from the side. In order to make images that reflect the full range of colors on the board and the components, we found the best light source is a "cloudy day" illuminator, which is similar to daylight on a cloudy day. This type of light source has uniform diffuse light from all angles. While easy to describe, in practice this was hard to design and implement in an efficient form-factor, with a reasonable price tag, and sufficient intensity.

8. How do lenses affect color?

The goal of the lens is to transmit and focus the light reflecting from the stimulus (e.g. PCB board) onto the sensor array of the camera. One might expect that a lens that works with a monochrome camera should work equally well with a color camera. However, this is not the case. Red, green, and blue wavelengths focus at different locations due to their different wavelengths. This is known as chromatic aberration. Special lenses such as *Achromatic* lenses and *Apochromatic* lenses are specially corrected to bring two or three colors into focus in the same plane. Naturally, these add cost to the system and the color issues, even with the introduction of these lenses, are not completely resolved. Beyond this problem, lenses introduce other a host of other types of errors. Thus, the lens choice (and software to correct for residual lens errors) is critical to the success of adding color to a PCB inspection system. Finally, the choice of camera determines what type of lens is appropriate. For instance, three CCD cameras have different lens requirements than one CCD cameras.

9. How do I go from color to grayscale?

Going from color to grayscale is actually an easy function. One simply takes a linear weighting of the red, green and blue channels and combines them together. A very common weighting of the channels to produce a grayscale image is shown below.

$$Y = 0.3*R + 0.59*G + 0.11*B$$

This process assumes that the image that has come from camera has already been converted to color and is being output in an RGB format. However, if one uses a Bayer filter, one CCD camera, it may be possible to get to the raw data from the camera and use these data directly or with some simple processing to get the gray values. For instance, one could use only the green pixels to create a grayscale image. Simple interpolation would need to be done to fill in the missing green values. (Naturally, this assumes, however, that the green channel adequately represents the stimulus that one is looking at.) The main idea is that full color recovery may not be required for the majority of the image. Simple algorithms could be used on areas that require grayscale processing. More complex algorithms could be used in areas that require full color processing. Additionally, if color is only required for the human operator, the underlying algorithms may do all their processing in grayscale. The color image could be constructed only for viewing purposes.

10. Summary

To summarize, color offers many potential advantages to an AOI system. However, it is important to remember that many of these advantages come at some cost, requiring a careful assessment of the tradeoffs. Furthermore, color is not always essential for an AOI system. Our own experience bears this out.

We can compare two very different systems that we built and modified over the last 10-15 years, first, the Landrex Optima™ 7300 series and, second, the Landrex Optima™ 7200 series. The first is based on grayscale analysis and the other color analysis. They both achieve their goals of finding defects. However, they perform their tasks based on different types of information. The first system was designed primarily as a post-reflow defect analysis system. It has multiple, angled monochrome cameras, red LED lighting, and multiple lighting modes. With this combination of cameras and lighting modes, it is able to capture the shape of solder joints, which are inherently colorless (or spectral), and also the geometry of parts to determine presence and absence. This type of analysis is known as “shape from shading” with elements of “shape from multiple cameras” The second system, in contrast, was designed primarily as a pre-reflow metrology system and has a single camera, Bayer filter color system. It uses color analysis as the basis to distinguish part and paste from the bare board and other features on the board. Sophisticated color processing helps to eliminate color artifacts. This type of analysis is known as “shape from color”. These two machines were architected in two different ways. However, both perform their required tasks.

In summary, many customers today view AOI systems in a binary fashion as either having color or not having color. If the AOI system does not have color capability, it may be deemed to be old-fashioned or unable to perform its necessary functions. The real question the customer must ask is: How do the image capture system and the image analysis algorithms work to find defects? If color is a component of this system, what is it used for and how do the issues associated with color recovery, lighting, and lens artifacts affect its performance?

Acknowledgements: *The author would like to thank Warren Smith of VISICS Corporation (70 Hastings Street Wellesley, MA 02481, warren.smith@visics.com) for his help with the paper and for contributing images for figures 15 and 16.*