



Getting the Most from Your IR Camera

Our thanks to FLIR for allowing us to reprint the following article.

Understanding IR camera calibration and corrections help ensure accurate temperature measurements and thermographic mapping.

Quantitative Measurements with IR Cameras

For best results, IR camera users need to think carefully about the type of measurements they need to make, and then be proactive in the camera's calibration process. Of course, the first step is selecting a camera with the appropriate features and software for the application. An understanding of the differences between thermographic and radiometric measurements is very helpful in this regard.

Thermography is a type of infrared imaging in which IR cameras detect radiation in the electromagnetic spectrum with wavelengths from roughly 900 to 14,000 nanometers (0.9–14 μm) and produce images of that radiation. Typically, this imaging is used to measure temperature variations across an object or scene, which can be expressed in degrees Celsius, Fahrenheit, or Kelvin.

Radiometry is the measurement of radiant electromagnetic energy, especially that associated with the IR spectrum. It can be more simply defined as an absolute measurement of radiant flux. The typical unit of measure for imaging radiometry is radiance, which is expressed in units of Watts/(sr-cm²). (The abbreviation "sr" stands for steradian; a non-dimensional geometric ratio expressing the solid (conical) angle that encloses a portion of the surface of a sphere equivalent to the square of the radius.)

In simple terms, one can think of thermography as "how hot" an object is, whereas radiometry is "how much energy" the object is giving off. Although these two concepts are related, they are not the same thing. IR cameras inherently measure irradiance not temperature, but thermography does stem from radiance. When you thermographically calibrate an IR system you are calibrating / measuring based on effective blackbody radiance and temperature. Therefore, the emissivity of the target object you are measuring is vital to achieving accurate temperatures. (Emissivity or emittance is the radiative property of an object relative to a perfect blackbody.)

Entry level IR cameras with microbolometer detectors operate according to non-quantum principles. The detectors respond to radiant energy in a way that causes a change of state in the bulk material (e.g., resistance or capacitance). Calibration software in these cameras is oriented toward thermographic imaging and temperature measurements. High-end IR cameras with photon detectors operate according to quantum physics principles. Although they also provide high quality images, their software is typically more sophisticated, allowing accurate measurements of both radiance and temperature.

Some reasons why radiance measurements are important include:

- Given a linear sensor, measured radiance is linear with incident energy. Temperature is non-linear with raw digital image counts, even with a linear sensor.
- Given the radiance and area of an object, radiant intensity can be calculated. Knowing total radiant intensity of a target gives a radiometric analyst the ability to model the irradiance generated by the target over various geometric and atmospheric conditions.
- The relationship between spectral bands of interest can be much easier to determine if you are working within radiometric units.
- The comparison between different objects in radiometric terms tends to have less uncertainty because emissivity is not a concern. (One still needs to consider atmospheric and spectral bandpass effects.)
- One can typically convert a radiometric signature from radiance to effective blackbody temperature given a few assumptions or ancillary measurement data. It tends to be more difficult to go from temperature to radiance.

Key Physical Relationships in Camera Operation

There are five basic steps in producing radiometric and thermographic measurements with an IR camera system:

1. The target object has a certain energy signature that is collected by the IR camera through its lens.

2. This involves the collection of photons in the case of a photon detector, or collection of heat energy with a thermal detector, such as a microbolometer.
3. The collected energy causes the detector to produce a signal voltage that results in a digital count through the system's A/D converter. (For example, a FLIR ThermoVisionR SC6000 IR camera has a 14-bit dynamic range in its A/D converter, which creates count values ranging from 0–16,383. The more IR energy incident on the camera's detector (within its spectral band), the higher the digital count.)
4. When the camera is properly calibrated, digital counts are transformed into radiance values.
5. Finally, the calibrated camera's electronics convert radiance values to temperature using the known or measured emissivity of the target object.

Expanding on Steps 4 and 5, an effective blackbody temperature measurement can be derived from a radiance measurement by applying a radiometric calibration, temperature vs. radiance model, and emissivity of the target object or scene. Every IR camera designed for serious measurements is calibrated at the factory. In the calibration lab, the camera takes a number of blackbody measurements at known temperatures, radiance levels, emissivities, and distances. This creates a table of values based on the A/D counts from the temperature/radiance measurements.

Once the counts for each blackbody temperature measurement are entered into the calibration software, the data are then passed through an in-band radiance curve fit algorithm to produce the appropriate in-band radiance vs. count values given the camera system's normalized spectral response function. This produces a radiometric calibration of in-band radiance [W/(sr-cm²)] versus the digital counts obtained while viewing a blackbody over a

range of temperatures. The result is a series of calibration curves. An example of how calibration points are captured is shown in Figure 1.

The calibration curves are stored in the camera system's memory as a series of numeric curve-fit tables that relate radiance values to blackbody temperatures. When the system makes a measurement, it takes the digital value of the signal at a given moment, goes into the appropriate calibration table, and calculates temperature. Due consideration is given to other factors like atmospheric attenuation, reflected ambient temperature, and the camera's ambient temperature drift before the final result is presented.

Ambient Drift Compensation (ADC). Another important consideration in the calibration process is the radiation caused by the heating and cooling of the camera itself. Any swings in camera internal temperature caused by changes in environment or the heating and cooling of camera electronics will affect the radiation intensity at the detector. The radiation that results directly from the camera is called parasitic radiation and can cause inaccuracies in camera measurement output, especially with thermographically calibrated cameras. Certain IR cameras (like the FLIR ThermoVisionR product line), have internal sensors that monitor changes in camera temperature. As part of the calibration process, these cameras are placed in an environmental chamber and focused at a black body reference. The temperature of the chamber and black body are then varied and data is collected from the internal sensors. Correction factors are then created and stored in the camera. In real-time operation, the camera sensors continually monitor internal temperature and send feedback to the camera processor. The camera output is then corrected for any parasitic radiation influences. This functionality is commonly referred to as ambient drift compensation.




Figure 1. Example of camera measurements and corresponding in-band radiance values for given black body temperatures with resulting radiance vs. measurement curve.

Ultimately, the camera must calculate an object's temperature based on its emission, reflected emission from ambient sources, and emission from the atmosphere using the Total Radiation Law. The total radiation power received by the camera can be expressed as:

$$W_{tot} = \epsilon \cdot \tau \cdot W_{obj} + (1 - \epsilon) \cdot \tau \cdot W_{amb} + (1 - \tau) \cdot W_{atm},$$

where ϵ is the object emissivity, τ is the transmission through the atmosphere, T_{amb} is the (effective) temperature of the object surroundings, or the *reflected ambient (background) temperature*, and T_{atm} is the temperature of the atmosphere.

The best results are obtained when a user is diligent in entering known values for all the pertinent variables into the camera software. Emissivity tables are available for a wide variety of common substances. However, when in doubt, measurements should be made to obtain the correct values.

Calibration and analysis software tools available to users are not always contained onboard the camera. While high-end cameras have many built-in software functions, others

rely on external software that runs on a PC. Even high-end cameras are connected to PCs to expand their internal calibration, correction, and analysis capabilities. For example, FLIR's ThermoCAMR RTools™ software can serve a wide variety of functions from real-time image acquisition to post-acquisition analysis.

Whether the software is on the camera or an external PC, the most useful packages allow a user to easily modify calibration variables. For instance, FLIR's ThermoCAM RTools provides the ability to enter and modify emissivity, atmospheric conditions, distances, and other ancillary data needed to calculate and represent the exact temperature of the object, both live and through saved data. This software provides a post-measurement capability to further modify atmospheric conditions, spectral responsivity, atmospheric transmission changes, internal and external filters, and other important criteria as needed.

The discussions that follow below are intended to represent both onboard and external camera firmware and software functions. Where these functions reside depends on the camera.

Typical Camera Measurement Functions

IR cameras have various operating modes to assure correct temperature measurements under different application conditions. Typical measurement functions include:

- Spotmeter
- Area
- Profile
- Isotherm
- Temperature range
- Color or gray scale settings

Cursor functions allow easy selection of an area of interest, such as the crosshairs of the spot readings in Figure 2. In addition, the cursor may be able to select circle, square, and irregularly shaped polygon areas, or create a line for a temperature profile. Once an area is selected, it can be “frozen” so that the camera can take a snapshot of that area. Alternatively, the camera image can remain live for observation of changes in temperature.




Figure 2. IR image of a printed circuit board indicating three spot temperature readings. Image colors correspond to the temperature scale on the right.

The spotmeter finds the temperature at a particular point. Depending on the camera, this function may allow ten or more movable spots, one or more of which may automatically find the hottest point in the image. The area function isolates a selected area of an object or scene and finds the maximum, minimum, and average temperatures inside that area. The isotherm function makes it possible to portray the temperature distribution of a hot area. Multiple isotherms may be allowed. The line profile is a way to visualize the temperature along some part of the object, which may also be shown as a graph (Figure 3).




Figure 3. Graph of temperature along a selected area of a target object using a camera's profile function

The temperature measurement range typically is selectable by the user. This is a valuable feature when a scene has a temperature range narrower than a camera's full-scale range. Setting a narrower range allows better resolution of the images and higher accuracy in the measured temperatures. Therefore, images will better illustrate smaller temperature differences. On the other hand, a broader scale and/or higher maximum temperature range may be needed to prevent saturation of the portion of the image at the highest temperature.

As an adjunct to the temperature range selection, most cameras allow a user to set up a color scale or gray scale to optimize the camera image. Figure 4 illustrates two gray scale possibilities.




Figure 4. Gray scale images of car engine; left view has white as the hottest temperature; right view shows black as the hottest

In Figure 2 a so-called “iron scale” was used for a color rendering. In a manner similar to the gray scale used in Figure 4, the hottest temperatures can be rendered as either lighter colors or darker colors. Another possibility is rendering images with what is known as a rainbow scale (Figure 5). In some color images, gray is used to indicate areas where the camera detector has become saturated (i.e., temperatures well above the top of the scale).




Figure 5. Rainbow scale showing lower temperatures towards the blue end of the spectrum




Figure 6. Isotherm color scale with each color having an isotherm width of 2°C

While choice of color scale is often a matter of personal preference, there may be times when one type of scale is better than another for illustrating the range of temperatures in a scene. In the case of isotherm measurements, areas with the same thermal radiance are highlighted. If we use a color scale with ten colors, we will in fact get ten isotherms in the image. Such a scale sometimes makes it easier to see the temperature distribution over an object. In Figure 6, the temperature

scale is selected so that each color is an isotherm with a width of 2°C.

Still, it is important to realize that an isothermal temperature scale rendering will not be accurate unless all of the highlighted area has the same emissivity, and the ambient temperatures are the same for all objects within the area. This points out common problems for IR camera users. Often, emissivity varies across an object or scene, along with variations in ambient temperatures, accompanied by atmospheric conditions that don't match a camera's default values. This is why IR cameras include measurement correction and calibration functions.

Emissivity Corrections

In most applications, the emissivity of an object is based on values found in a table. Although camera software may include an emissivity table, users usually have the capability of inputting emissivity values for an object ranging from 0.1 to 1.0. Many cameras also provide automatic corrections based on user input for reflected ambient temperature, viewing distance, relative humidity, atmospheric transmission, and external optics.

As described earlier, the IR camera calculates a temperature based on radiance measurements and the object's emissivity. However, when the emissivity value is unknown or uncertain, the reverse process can be applied. Knowing the object temperature, emissivity can be calculated. This is usually done when exact emissivity values are needed. There are two common methods of doing this.

The first method establishes a known temperature by using an equalization box. This is essentially a tightly controlled temperature chamber with circulating hot air. The length of time in the box must be sufficient to allow the whole object to be at a uniform temperature. In addition, it is absolutely necessary that the object stabilize at a temperature different from the surroundings where the actual measurements will take place. Usually, the object is heated to a temperature at least 10°C above the surroundings to ensure that the thermodynamics of the measurements are valid.

Once the object has reached the set temperature, the lid is drawn off and a thermogram is captured of the object. The camera and/ or software for processing thermograms can be used to get the emissivity value.

Another (“adjacent spot”) method is much simpler, but still gives reasonably exact values of the emissivity. It uses an area of known emissivity. The idea is to determine the temperature of the object with the camera in the usual way. The object is adjusted so that the area with unknown emissivity is very close to an area of known emissivity.

The distance separating these areas must be so small that it can be safely assumed they have the same temperature. From this temperature measurement the unknown emissivity can be calculated.




Figure 7a. PCB heated to a uniform 68.7°C, but digital readouts are incorrect.




Figure 7b. PCB with emissivity correction using the “adjacent spot” technique. Digital readouts now indicate the correct temperatures at all locations.

The problem is illustrated in Figure 7, which is an image of a printed circuit board (PCB) heated to a uniform temperature of 68.7°C. However, areas of different emissivities may actually have different temperatures, as indicated in the caption of Figure 7a. Using the technique just described, emissivity correction proceeds by finding a reference spot where a temperature of 68.7°C is indicated and calculating the emissivity at that location. By knowing the emissivity of the reference spot, the emissivity of the target spots can be calculated. The corrected temperatures are shown in Figure 7b.

As illustrated in these figures, this technique can be used with a camera’s area selection function (“AR” in the figures) and using the average temperature for that area. The reason for using the average temperature in the reference area is that there is usually a spread of temperatures within the area, especially for materials with low emissivity. In that case, using a spotmeter or an area

maximum value would give a less stable result. The isotherm function is not recommended either, as it is not possible to get the averaging effect with it.

It may also be possible to use a contact sensor to find the temperature of an area of unknown emissivity, but such measurements pose other problems that may not be easy to overcome. Furthermore, it is never possible to measure the emissivity of an object whose temperature is the same as the reflected ambient temperature from its surroundings.

Generally, a user can also input other variables that are needed to correct for ambient conditions. These include factors for ambient temperatures and atmospheric attenuation around the target object.

Using Camera Specifications

When considering IR camera performance, most users are interested in how small an object or area can be detected and accurately measured at a given distance. Knowing a camera’s field of view (FOV) specifications helps determine this.

Field of View (FOV). This parameter depends on the camera lens and focal plane dimensions, and is expressed in degrees, such as 35.5° × 28.7° or 18.2 × 14.6°. For a given viewing distance, this determines the dimensions of the total surface area “seen” by the instrument (Figure 8). For example, a FLIR ThermoVision SC6000 camera with a 25mm lens has an FOV of 0.64 × 0.51 meters at a distance of one meter, and 6.4 × 5.1 meters at a distance of ten meters.




Figure 8. A camera’s field of view (FOV) varies with viewing distance.

Instantaneous Field of View (IFOV). This is a measure of the spatial resolution of a camera’s focal plane array (FPA) detector. The configuration of the FPA in the FLIR ThermoVision SC6000 is 640 × 512 detectors, which makes a total of 327,680 individual picture elements (pixels). Suppose you are looking at an object at a distance of one meter with this camera. In determining the

smallest detectable object, it is important to know the area's IFOV covered by an individual pixel in the array. The total FOV is 0.64 × 0.51 meters at a distance of one meter. If we divide these FOV dimensions by the number of pixels in a line and row, respectively, we find that a pixel's IFOV is an area approximately 1.0 × 1.0mm at that distance. Figure 9 illustrates this concept.

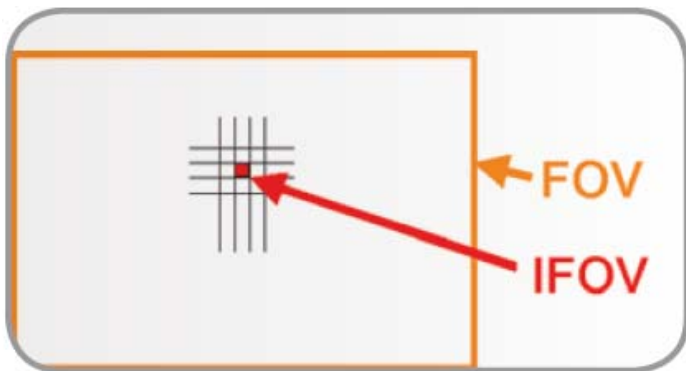


Figure 9. A camera's geometric (spatial) resolution (IFOV) is determined by its lens and FPA configuration.

To use this information consider, the pixel IFOV relative to the target object size (Figure 10). In the left view of this figure, the area of the object to be measured covers the IFOV completely. Therefore, the pixel will receive radiation only from the object, and its temperature can be measured correctly.




Figure 10. IFOV (red squares) relative to object size.

In the right view of Figure 10, the pixel covers more than the target object area and will pick up radiation from extraneous objects. If the object is hotter than the objects beside or behind it, the temperature reading will be too low, and vice versa. Therefore it is important to estimate the size of the target object compared to the IFOV in each measurement situation.

Spot Size Ratio (SSR). At the start of a measurement session, the distance between the camera and the target object should be considered explicitly. For cameras that do not have a calibrated spot size, the spot size ratio method can be used to optimize measurement results. SSR is a number that tells how far the camera can be from a target object of a given size in order to get a good temperature measurement. A typical figure might be 1,000:1 (also written 1,000/1, or simply abbreviated as

1,000). This can be interpreted as follows: at 1000 mm distance from a target, the camera will measure a temperature averaged over a 1mm square.

Note that SSR is not just for targets far away. It can be just as important for close-up work. However, the camera's minimum focal distance must also be considered. For shorter target distances, some manufacturers offer close-up lenses.

For any application and camera/lens combination, the following equation applies:

$$\frac{D}{S} = \frac{SSR}{1}, \quad \text{where } R$$

D is the distance from the camera to the target,

S is smallest target dimension of interest, and

SSR is the spot size ratio.

The units of D and S must be the same.

When selecting a camera, keep in mind that IFOV is a good figure of merit to use. The smaller the IFOV, the better the camera for a given total field of view.

Other Tools for Camera Users

As mentioned earlier, IR cameras are calibrated at the factory, and field calibration is not practical. However, some cameras have a built-in blackbody to allow a quick calibration check. These checks should be done periodically to assure valid measurements.

Bundled and optional data acquisition software available for IR cameras allows easy data capture, viewing, analysis, and storage. Software functions may include real-time radiometric output of radiance, radiant intensity, temperature, target length/area, etc. Optional software modules are also available for spatial and spectral radiometric calibration. Functions provided by these modules might include:

- Instrument calibration in terms of radiance, irradiance, and temperature
- Radiometric data needed to set instrument sensitivity and spectral range
- Use of different transmission and/or emissivity curves or constants for calibration data points
- Adjustments for atmospheric effects

In addition, IR camera software and firmware provide other user inputs that refine the accuracy of temperature measurements. One of the most important functions is non-uniformity correction (NUC) of the detector FPA. This

type of correction is needed due to the fact that each individual detector in the camera's FPA has a slightly different gain and zero offset. To create a useful thermographic image, the different gains and offsets must be corrected to a normalized value.

This multi-step NUC process is performed by camera software. However, some software allows the user to specify the manner in which NUC is performed by selecting from a list of menu options. For example, a user may be able to specify either a one-point or a two-point correction. A one-point correction only deals with pixel offset. Two-point corrections perform both gain and offset normalization of pixel-to-pixel non-uniformity.

With regard to NUC, another important consideration is how this function deals with the imperfections that most FPAs have as a result of semiconductor wafer processing. Some of these imperfections are manifested as bad pixels that produce no output signals or as outputs far outside of a correctable range. Ideally, the NUC process identifies bad pixels and replaces them using a nearest neighbor replacement algorithm. Bad pixels are identified based on a response and/or offset level outside user-defined points from the mean response and absolute offset level.

Other NUC functions may be included with this type of software, which are too numerous to mention. The same is true of many other off-the-shelf software modules that can be purchased to facilitate thermographic image display, analysis, data file storage, manipulation, and editing. Availability of compatible software is an important consideration when selecting an IR camera for a user's application or work environment.

Conclusions

Recent advances in IR cameras have made them much easier to use. Camera firmware has made setup and operation as simple as using a conventional video camera. Onboard and PC-based software provides powerful measurement and analysis tools. Nevertheless, for accurate results, the user should have an understanding of IR camera optical principals and calibration methods. At the very least, the emissivity of a target object should be entered into the camera's database, if not already available as a table entry.

