Determining MTF with a Slant Edge Target

Douglas A. Kerr

Issue 2
October 13, 2010

ABSTRACT AND INTRODUCTION

The modulation transfer function (MTF) of a photographic lens tells us how effectively the lens transfers a luminance variation in the scene (by which detail is conveyed) onto the focal plane, and in particular how that varies with spatial frequency (which we can think of as the "fineness" of the detail). This function indicates, objectively, the "resolving potential" of the lens.

We often read of the MTF being determined using a *slant edge target* test. In this article we review the concept of the MTF and the principles of this testing technique.

THE MODULATION TRANSFER FUNCTION

We will examine the concept of the *modulation transfer function* by looking in sequence at the three words that make up its description.

Modulation

Modulation in this case refers to the variation in the luminance of a scene from point to point, and the corresponding variation in the illuminance from point to point in the image deposited by the lens on the focal plane (on the film or digital sensor). Detail is conveyed by such variation; if there is no variation in luminance, the scene is "uniform gray" and hardly worth photographing.[1]

Modulation can be quantified in terms of *modulation depth*, a way of expressing the ratio between the maximum and minimum luminance (or illuminance) across a certain small part of the scene (or image).

Transfer

For our purposes here, the job of the lens is to transfer the luminance variation of the scene into an illuminance variation on the image. It does this incompletely, for various reasons. We can quantify the degree to which it accomplishes this job in terms of the *modulation transfer ratio*. This is the ratio of (a) the modulation depth of the

[1] For the sake of simplicity, we will assume only "monochrome" scenes and "gray scale" photography, so that luminance/illuminance is the only property of interest.

illuminance deposited on the image to (b) the modulation depth of the luminance of the scene (usually within a small region).

The modulation transfer ratio is quite parallel to the *gain* of an amplifier stage in an electronic system.

Function

In mathematics, when the value of one variable quantity depends on the value(s) of one or more other variable quantities, in some specific way, the first variable is said to be a *function* of the other variable(s).

Under this concept, we may say that "*x* is a function of *y* and *z*". That means that the value of variable *x* (called the *dependent variable*) is determined by the values of variables *y* and *z* (called the *independent variables*).

A specific name identifying a function can, in common practice, mean three distinct things:

- The variable x itself (after all, we said "*x* **is a function** . . . ").
- The value of *x* for a certain set of values of the independent variables (the "value of the function" for that situation).
- The overall relationship by which *x* depends on *y* and *z* (the function proper).

This diverse use of the function name can be confusing if we have not been forewarned about it.

<u>Graphic representation of a function</u>

If a variable is a function of **one** other variable ("*x* is a function of *y*"), we can show the relationship graphically in the familiar way—a plot of "*x* against *y*".

If one variable is a function of **two** other variables ("*x* is a function of *y* and *z*"), we cannot show the relationship graphically in the familiar way.

Often what we will do then is to take one of the independent variables and (arbitrarily) consider it to be a *parameter* (it is still an independent variable; we just handle it a little differently). Suppose we decide to treat *z* as the parameter.

We adopt some specific value of *z* and, holding that constant, plot the variation of *x* with *y* (labeling the curve with the value of *z*). Then we take another specific value of z and, holding that constant, again plot the variation of *x* with *y* (labeling that curve with the new value of *z*).

The result is what we often describe as a "family of curves", one curve for each of our chosen values of the parameter, z.

But we can equally legitimately decide to treat *y* as a parameter. Then we choose a certain value of y and, holding that constant, plot the variation of *x* with ***z,*** and so forth.

Which of those we do will depend on the context in which we wish to visualize the variation of *x*.

The modulation transfer function

For a given lens with a given aperture (and focal length setting, if relevant), the modulation transfer ratio varies with several factors, most prominently:

- The spatial frequency[2] of the modulation (which we can think of as the "fineness" of the detail the modulation conveys). Typically the modulation transfer ratio decreases as the spatial frequency increases.

- The location in the image of the area of interest (notably its distance from the optical axis of the lens. Typically the modulation transfer ratio decreases as we move from the optical axis.

Thus, the modulation transfer ratio is a function of spatial frequency and distance off axis.

This function is called the *modulation transfer function* (MTF) of the lens. And of course, as we discussed earlier, the term MTF is also applied to the modulation transfer ratio (which we then never hear of under its own name), or to its value in a particular situation.

Two presentations

As we mentioned above, when a variable is a function of two other variables, there are two ways to present the relationship graphically, choosing either of the two independent variables to play the role of a parameter.

For scientific or optical engineering work with the MTF, we normally select distance off the axis as the parameter, and plot the modulation transfer ratio against spatial frequency (preferably in cycles/mm). But,

---

[2] Spatial frequency has dimensions of *cycles per unit distance*. In scientific work, the unit is typically *cycles per millimeter*. Often this unit is spoken of in optical work as "line pairs per mm", but sometimes as "lines per mm", a source of considerable confusion. There are historical justifications for both these conflicting practices; these are beyond the scope of this article.

given the dual use of the term MTF, we are almost forced to say, "we select distance off the axis as the parameter, and plot MTF against spatial frequency." In other words, this form of the MTF is a plot of MTF against spatial frequency.

However, when MTF data is presented by lens manufacturers, they customarily select spatial frequency as the parameter, and plot the "MTF" (meaning the modulation transfer ratio) against distance off axis. Usually, there are only two curves, for a "low" and a "not so low" spatial frequency.[3]

DETERMINING THE MODULATION TRANSFER RATIO

The classical concept

The classical concept of determining the MTF of a lens involves presenting it with patterns having repetitive variations in luminance (of a known modulation depth) at different spatial frequencies. Then, the pattern deposited on the focal plane is examined (perhaps with a special instrument, or perhaps by capturing it with precisely calibrated film) and noting the modulation depth for each test pattern. We make this determination both at the center of the image and then at locations at successively greater distances from the axis. The two modulation depths, for each combination of spatial frequency and distance off axis, are compared to get the modulation transfer ratio.

This is then plotted against the appropriate non-parameter independent variable for the desired form of presentation.

Although in the form of the MTF curves presented by lens manufacturers often only two spatial frequencies are treated, for scientific work it is important that we have the MTF at numerous spatial frequencies. Doing so requires test exposures done with numerous test targets, each having patterns of lines at various spacings.

A more modern method

The availability of computers to easily perform sophisticated manipulation of data, and the fact that a digital camera inherently has an instrument for measuring illuminance the focal plane (its sensor), have led to the adoption of a quite different technique for determining the MTF of a lens, the *slant edge target* technique. This technique is the actual subject of this article.

---

[3] Actually, there are often eight curves, accommodating two values of the parameters *aperture* and *modulation axis*.

THE SLANT EDGE TARGET TECHNIQUE—CONCEPT

An "analog" in electrical engineering

The underlying concept of the technique can perhaps be most clearly seen by considering an electrical engineering example.

The MTF (in the sense of a plot of modulation transfer ratio against spatial frequency) is quite parallel to the matter of the "frequency response" of an electronic amplifier, where we plot the *gain* of the amplifier (the ratio of the output voltage to the input voltage) as a function of frequency (in this case "temporal" frequency, in hertz).

Not surprisingly, the classical technique for determining the frequency response (we can call it the "gain function") involves presenting the amplifier with signals of known voltage at different frequencies, and in each case, measuring the output power. The plot of the gain (ratio of output voltage to input voltage) against frequency is the *voltage gain function*.

But there is a way to determine this with a "one shot" test (and the term is very apt). We submit to the amplifier what is called an *impulse*, a single pulse which (ideally) has zero duration (zero width) but still contains energy.

When we do this, a certain waveform comes out of the amplifier. It is called the *impulse response* of the amplifier. If we capture that (just one test is needed), we can from it determine the entire voltage gain function (gain as a function of frequency).

How can this be? Well, the impulse contains energy at all frequencies (in theory, up to infinity), with a uniform distribution. If we take the Fourier transform[4] of the output waveform, the result is a description of the frequency content of that waveform. And, given that the input signal contains "all frequencies", uniformly, that description will be the *voltage gain function* (or "voltage frequency response").

Well, clever as this sounds on paper, there are some practical problems with actually doing it. One is that our impulse, if it is truly to have a zero duration (zero width in time) but nevertheless contain some energy (and of course, if it didn't there would be no output from it), it must (theoretically) have infinite amplitude (voltage). Let's be thankful we can't actually do this; if we could, our amplifier would blow up during the test.

---

[4] A mathematical "process" that takes a description of a waveform and from it develops a description of its frequency content.

And making a pulse have zero width isn't possible either.

So we resort to a variation of the theme. Here, instead of using an impulse as our input we use a *step function*. This is a waveform that, for example, starts out at +1.0 volt and then, at a certain point in time, **instantaneously** changes to −1.0 volt. Again this is not possible to actually achieve, but it is a lot easier to approximate than an impulse with zero time width and infinite voltage.

After applying this (just once) to our amplifier, and capturing the output waveform, we then take the Fourier transform of that. The result, as before, will be the frequency response (gain function) of the amplifier (although in this case, it is in terms of power gain rather than voltage gain).

Now, back to optics

If we present a "zero-width" bright line to a lens, it is the optical equivalent of the impulse in the electrical situation. Unless the lens has "infinite resolution", the image of that line on the focal plane will be a pattern having non-zero width, across which the illuminance varies in some way. This is called the *line spread function* (LSF) of the lens.

If we take its Fourier transform, we get what turns out to be the square root of the modulation transfer ratio as a function of spatial frequency: the modulation transfer function (MTF).

But of course, just as for the electrical impulse, this "zero width" line is impractical to make, and for it to have enough photometric energy that we can see the pattern of illuminance on the focal plane, it would have to have "essentially infinite" luminance.

So we follow the same ploy used in the electrical situation. We use a "test scene" that is black up to a straight line boundary and white beyond it—the optical equivalent of the electrical step function.

For any "real" lens, the image of that test scene will not have a zero width boundary between dark and light regions, but rather a boundary of some finite width, across which the illuminance varies in some way. The plot of illuminance across that boundary is called the *edge spread function* (ESF) of the lens.

If we measure this illuminance pattern take its Fourier transform, we get the modulation transfer ratio as a function of spatial frequency: the modulation transfer function (MTF).

Wow! Is this neat or what!

THE REALITIES

The need

In order to do this, for MTFs of the kind we fortunately encounter with modern lenses, we have to be able to measure the illuminance pattern—the edge spread function—with very high resolution.

Of course, a practical advantage of this technique is that we can use the camera sensor itself to measure the illuminance pattern. But the theoretical resolution of the sensor array is not sufficient to discern the illuminance pattern with sufficient resolution. We see this illustrated in Figure 1.

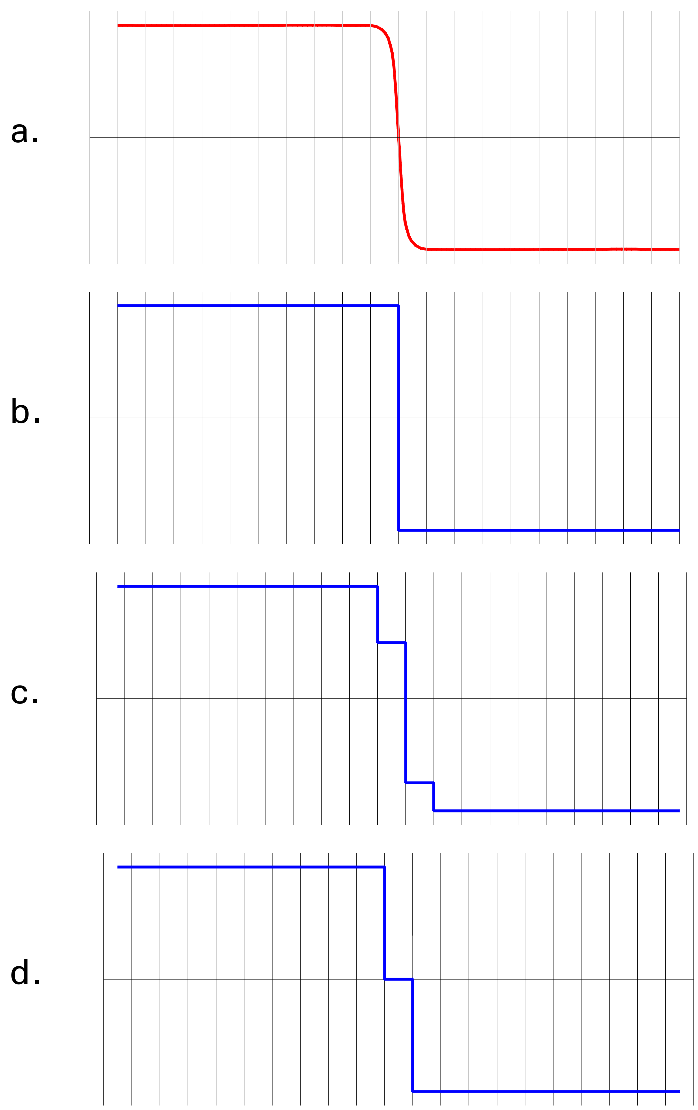

Figure 1. Resolving the edge spread function

In panel **a**, we see a hypothetical edge spread function (as would be observed downstream from the lens under test). The gray grid is at the pixel pitch of the camera sensor array, in order to give an idea of the scale.

In panel **b**, we see what would happen if the edge image was located in a certain way on the pixel grid. (We only consider pixels along a line perpendicular to the boundary). The plot line across the “band” for each pixel shows the pixel output (only a single value for any pixel, of course). Note that the overall sensor output for this row of pixels seems to be a perfect “step function” (in electrical terms).

In panel **c**, we see a slightly different location of the image. Now we see a different pixel output—still certainly not a precise representation of the illuminance pattern itself.

In panel **d**, we see yet another possibility—again not even close to a precise representation of the illuminance pattern.

So regardless of which one of these happens—and this is essentially beyond our control—the illuminance pattern suggested by the sensor output is useless for precise analysis.

So we must “fake” enhanced resolution of the sensor.

The slant edge target

Enter now the title character of this drama. As before, we present the lens with a target with a black portion and a white portion, with a sharp boundary between. But we intentionally orient it so that the boundary does not match the pixel axis of the sensor array, by a small angle. Now, a fascinating drama can play out; we can follow it on Figure 2.

We see the image of the target laid out on the sensor pixel detector grid. (The black portion is shown in gray to allow the entire grid to be seen.) Each square represents the “domain” of one pixel detector. But we will assume that each detector actually only responds to the illuminance at the center of its domain (where we will show a dot if we are interested in the output of that detector).

The variation in illuminance (the edge spread function) happens along the “ESF axis” direction, and of course it happens identically all across the edge. That is, the *illuminance* will be constant along any line parallel to the boundary (a certain distance from the edge); the *variation in illuminance* will be the same along any line parallel to the ESF axis (which is drawn in an arbitrary location).

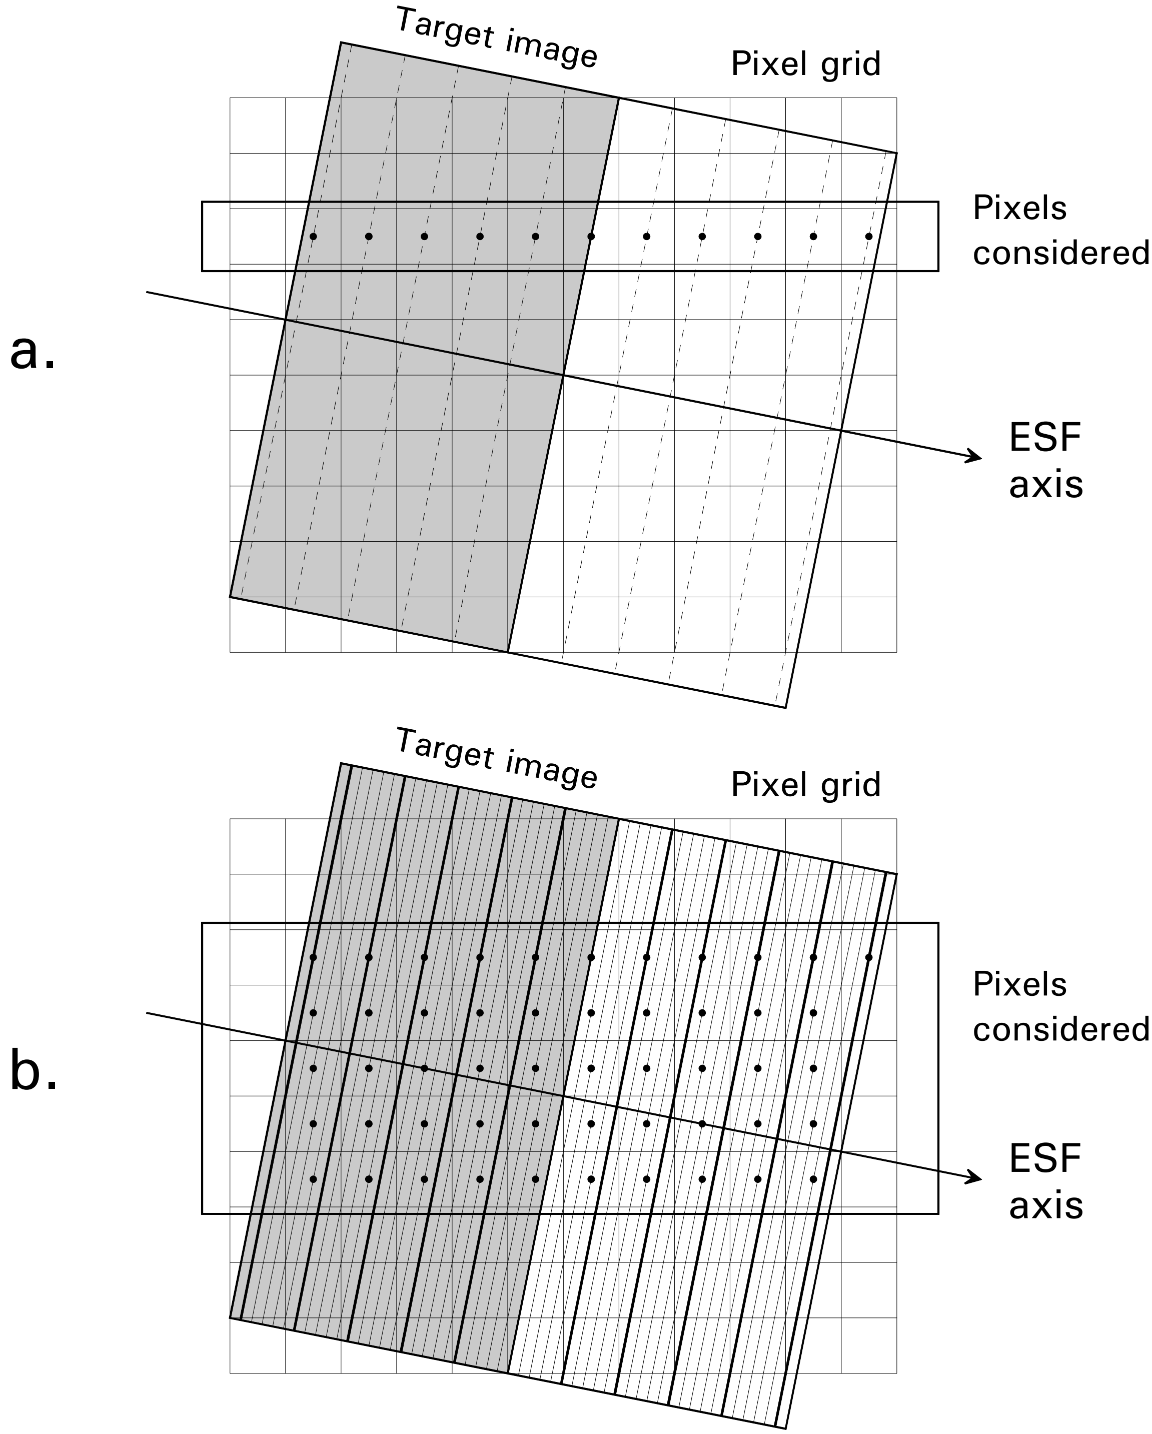

Figure 2. Operation of the slant edge target

We first consider the response of the line of pixel detectors (hereafter, just "pixels") highlighted in panel **a**. These pixels pick up the luminance of the "edge spread pattern" at various distances from the boundary, which are evenly spaced. That illuminance is the same all along the associated dotted line, drawn parallel to the boundary. Thus a measurement taken at any point along such a line represents the illuminance every place along it (including where the line crosses our arbitrarily-drawn "ESF axis", along which we are interested in the variation of illuminance).

The reason we have only concentrated on one row of pixels in this panel is not because they have any special role, but merely because if we started by considering all the pixels, the drawing would have been so busy that it might have been hard to grasp the principle from it.

But now that we know what we are looking for, in panel **b** we consider the response of all the pixels over a larger region. Recall that the output of any pixel represents the illuminance any place along a line parallel to the boundary. Thus we have again drawn the lines parallel to the boundary through each pixel point. The illuminance is the same along any of these lines. We've not drawn them dotted as that is just too "busy" for this already-too-busy drawing. But we have drawn slightly **bolder** the ones shown in panel **a**.

We see now that the suite of output data from all these pixels has told us the luminance along each of many lines **parallel** to the boundary, and very closely (and evenly) spaced. These values are in fact the luminance at points with that particular spacing along our arbitrarily-drawn ESF axis.

Accordingly, this suite of data gives us a "high-resolution" description of the variation of illuminance along the ESF axis; that is, a high-resolution description of the ESF itself, which we require to make a precise determination of the MTF.

The spacing of the "samples" of the ESF is in fact the pixel pitch multiplied by the sine of the angle of "rotation" of the target. In our illustration (where the rotation is about 11.3° [5]), this is a little less than one-fifth the pixel pitch. Thus, our clever approach gives us an effective resolution of about five times that which could be given by the sensor array in normal use.

Because the pixel detectors actually do not pick up the luminance at a point (as suggested by our example), but rather respond to an average of some sort over a region approaching the domain of the pixel, certain special steps have to be taken in the evaluation of the edge

---

[5] This is a greater angle than that usually used for such tests, adopted here for clarity of the illustration. One widely used test target uses an angle of about 5.7°, specifically a slope (tangent of the angle) of 1:10. The tidy repetitive pattern of "sample" distances we see in the example requires an angle whose tangent is a ratio of integers, preferably "1/n".

spread function from the set of collected pixel detector values. This is a well-known matter in digital signal processing.

Note that the axis along which the edge spread function is considered (by definition, perpendicular to the "edge") is not either axis of the pixel array. This is not really of any consequence to us; the edge spread function exists in two dimensional space regardless of the orientation of the target.[6]

Target orientation

Any given scheme for determining the MTF with the slant edge technique will have an "intended" rotation of the target edge. However, we cannot always assure that this angle is exactly achieved.

MTF analysis software for use with the slanted edge target technique typically contains provisions for first deducing the exact rotation of the target edge from the data (you can visualize from Figure 2 how this generally could work) and then using the result in the actual analysis.

SUMMARY

The slant edge target approach allows a convenient "one-shot" determination of the MTF (in the sense of the modulation transfer ratio as a function of spatial frequency) by exploiting two clever ploys:

- The use of the Fourier transfer to get the MTF from the edge spread function.
- The use of the "slanted" target to get an effective resolution of the sensor array much greater than would be dictated by its pixel pitch so that the edge spread function can be adequately measured by the sensor array itself.

#

[6] Actually, when we get into one of the esoteric subtleties of the MTF (the matter of "axis of modulation"), the direction of the ESF axis is of concern. We can deal with that by thoughtful choice of at what points in the image (at different distances from the center) do we run tests.