



How to Read MTF Curves

by

H. H. Nasse

Preface

„ The rules of optics are complex and nasty! “

This is a nice, honest sentence from an internet discussion about ‚How to read MTF curves’, telling us how difficult to understand this world of numbers might be for photographers.

Nevertheless I am going to show you on the following pages, that things are not that bad, and that you can understand the basic relationships without an excursion into higher mathematics of Fourier-optics.

After reading the paper you will be able to conclude about the character of a lens from MTF data published by makers or testing institutes. You will however also learn about the limitations of MTF, so that you can read lens reviews critically.

And those who see too many numbers and curves may be assured that these are not really necessary for good photography, since photography is mainly based on experience. But it is great fun to understand your tools in a better way, and this what I wish you to have during reading this first part. In a second part we will show you a number of illustrating images.

Point Spread Function

When photographers want to take a very natural-looking picture of a subject, they would like to have an ideal lens on their camera, one which allows all light rays emanating from one point on the object to meet again at exactly one point of the image. We now know that with real lenses we can go only part way to achieving this ideal. Image points in the geometric sense of the word do not exist in reality.

Aberrations of the lens systems, production tolerances, and ultimately the wave-like nature of light as well, are the reasons why the light originating from one point of the object is always distributed over an area around the ideal image point.

To a certain extent, this area is the 'smallest possible circle of confusion'; however, the light therein is not evenly distributed, the intensity usually decreases from the inside to the outside and the shape is rarely circular. This effect is known as "**point spread function**". Its shape and size characterize the image quality of a lens.

If it is possible to compare photography with painting, the point spread function is the handwriting, the brushstroke of a lens. Just as there are wide, flat, pointed or even bristly brushes, lenses also have very different styles of handwriting.




But why then is it still not used to quantitatively describe image quality?

There are **three reasons** for this:

First of all, the shape is sometimes very complicated and therefore defies a simple numerical description. This is illustrated in the following pictures taken with a microscope. The first six point spreads in the images 1 to 6 on the next page are examples of useable, but moderate image quality and are typical

of high-speed lenses at full aperture, wide-angle lenses on the edge, or slight defocusing.

A small white square has been pasted into each image for a size comparison; it represents an $8.5\mu\text{m}$ pixel area like those of a 12 MP, 35mm format full-frame camera. All these point spreads are thus considerably larger than this (relatively large) pixel area.



The point spread **Nr. 7** in the second row above is an example of outstanding imaging performance. A digital sensor generally does not see such small point spreads, however. The image **Nr. 8** on the far right shows the same point spread behind a low pass filter that is usually positioned in front of the sensor and is intended to suppress the Moiré effect. The image quality is therefore artificially deteriorated in the low pass filter by increasing the point spread considerably by means of several birefringent discs.

The **second reason** is that you almost never see such single, isolated point spreads. For example: Only if you take pictures of stars on a dark night do you achieve the same effect as that here in the lab. Most images are generated in the camera in a complicated way of combining the parts of a large number of single point spreads.

This is because a small area of the object consists of many densely packed points and these correspond to many densely packed ideal image points in the image behind the lens. Since the real point spreads cannot be infinitely small, this means that the individual point spreads overlap: The intensity at one point of the image (you could even say **in one pixel**) is generated by a two-dimensional integration (summation) of many point spreads. There is thus a not-so-easily manageable mathematical connection between the 'brushstroke' and the image that we see.

The **third reason** is that the entire imaging chain from the lens to the eye can be much more elegantly described with the method that I would now like to explain.


Modulation transfer

Since we are primarily interested in how extended objects are imaged, objects which, unlike stars, comprise an infinite number of points, we must find another way to quantitatively describe the image quality. We use a **sinusoidal brightness distribution** to examine how an object that looks as simple as possible is imaged. The sinusoidal brightness distribution is a pattern of bright and dark stripes in which the transition between bright and dark occurs gradually and continuously, i.e. sinusoidally, just as the electric power in our sockets varies with time. We use the sinusoidal stripe pattern because the result in the image is once again a sinusoidal pattern,

regardless of how complicated the shape of the point spread may be.

Several of its properties also remain stable or at least have nothing to do with imaging quality: The **direction** of the stripes does not change and the **frequency** – the number of stripes per unit length – only changes according to the imaging scale.

What is no longer identical to the original is the difference in brightness between the dark and bright stripes. This is because the extended point spreads ensure that part of the light falls on a position that would actually be completely dark, instead of falling on a bright location.



This graph shows a sinusoidal stripe pattern (black curve) as an intensity profile (a cross-section perpendicular to the stripes). It has 20 periods per millimeter, so one period is 50 μm long. The red and blue curves are cross-sections of the brightness distribution in a point spread. The brightness that would exist with ideal imaging at the point of the sine pattern marked in blue is distributed to the surrounding area according to the blue curve. You can therefore see that some of this light falls into the dark “valleys” at 25 μm next to the blue point.

Light also falls there from the red dot on the flank of the sine pattern. Although the sine pattern on the flank is darker, a larger fraction reaches the point at -25 μm , since the red dot is closer to the dark valley. Thus, the intensity in the dark areas of the pattern is the sum of many such contributions from the neighboring areas. The result is then the weaker modulated curve labeled “image”. The brightness of the dark stripes in the image is raised by the aberrations, while the bright stripes get darker.

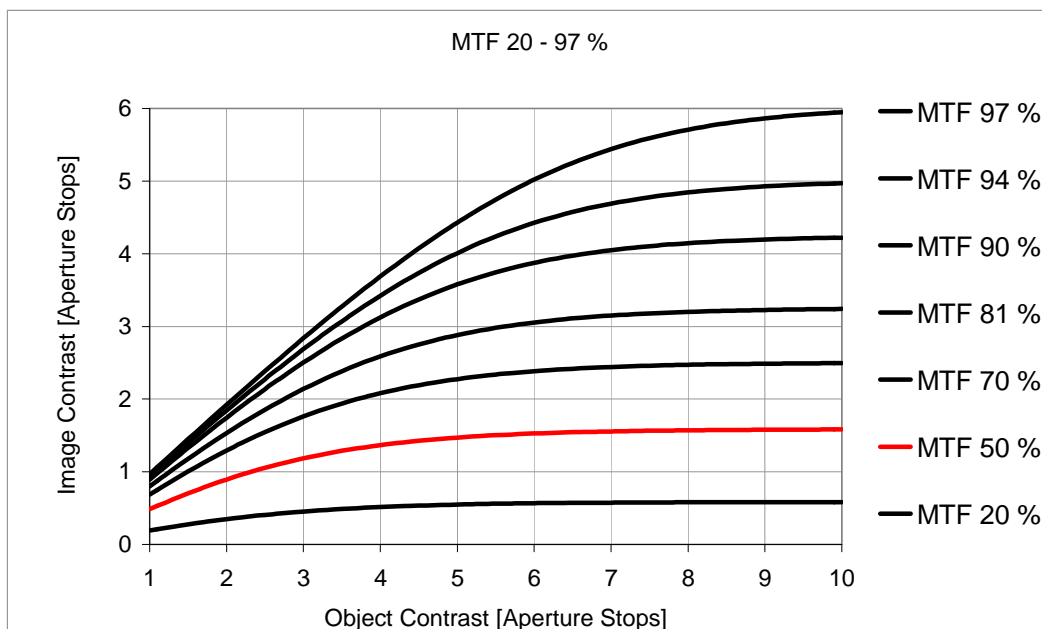
In optics, the difference between bright and dark is referred to as “**contrast**”. Seen from a more general point of view, the difference between maximum and minimum for all sinusoidal, periodically changing quantities is called “**modulation**.” If we compare the modulation of the image with the modulation of the object by simply dividing these two figures by each other, we get a simple figure that provides a statement about the imaging properties of the lens: The **modulation transfer**. Thus, we have already understood the first two letters of the term “MTF”. It is a number between 0 and 1 or between 0 and 100%.

The photographer is used to expressing bright-dark differences in aperture stops, which is also very reasonable as the perception of our eyes follows such logarithmic scales. But, what, for example, does a modulation transfer of 50% mean if our pattern of stripes consists of a difference of 6 aperture stops between the brightest and darkest points, i.e. a brightness ratio of $1 : 2^6 = 1 : 64$? Is the difference in the image 3 aperture stops or 1:32, which would correspond to 5 aperture stops? Both would be wrong. In reality, we would then still have approximately 1.5 aperture stops in the above-mentioned case. This is because, in optics, the “contrast” parameter is defined as follows:

$$\text{Contrast} = \frac{\text{Maximum} - \text{Minimum}}{\text{Maximum} + \text{Minimum}}$$

Therefore, in our example, the contrast of the object is 63 divided by 65, or approx. 0.97. After imaging with a modulation transfer of 50%, the contrast in the image is only half as high, approx. 0.48. Minimum to maximum is then approx. 1:2.9. (1.9/3.9 = 0.48)

The following graph shows how object contrast and image contrast are related for different modulation transfers if they are measured in aperture stops:



We can then recognize three **important properties of MTF** here which we should remember when reading MTF curves:

1. Small differences in higher MTF values are particularly significant at high object contrast levels.
2. On the other hand, weak tonal value variations of less than one aperture stop do not require high MTF values. Differences above 70-80% are then hardly relevant.
3. With very low MTF values, it practically does not matter how high the object contrast is; the image contrast is always low.

Incidentally, this is why the datasheets of films always also gave the resolving power for the low contrast of 1:1.6. The resolution figures for the contrast of 1:1000 can only be measured using contact exposure. For the finest structures (i.e. very high spatial frequencies), no lens in the world is capable of producing a contrast of ten aperture stops. Estimating the amount of information of film images based on this higher resolution value is thus too optimistic.


Modulation transfer function, resolving power

It is obvious that one single stripe pattern is not sufficient to characterize the quality of a lens. A very coarse pattern with large separations between bright and dark stripes could, of course, also be imaged well by a lens with a relatively large point spread function. If we decrease the separation between the stripes, however, so that the separation between bright and dark approaches the size of the point spread, then a lot of light from the bright zone is radiated into the darker zones of the pattern and the image contrast becomes noticeably lower.

If we want to use the comparison with painting again: Coarse structures can be painted well with a thick brush, the pointed fine brush is required for fine details, however.

We therefore need to investigate how the lens images stripe patterns of various degrees of fineness, i.e. we need to determine a modulation transfer for every one of these patterns. We thus obtain a whole sequence of numbers, and if we plot them as a function of a parameter which describes the fineness of the stripe pattern, these numbers represent a curve, the **modulation transfer function**.

The fineness of the stripe pattern can be measured by counting how many periods of the pattern are contained in a distance of one mm in the image. A period is the separation between two bright or two dark stripes, or the width of a line pair consisting of one dark and one bright stripe. The number of periods per millimeter in the image plane is the **spatial frequency**, given in the unit **line pairs per millimeter**, abbreviated to **lp/mm**.

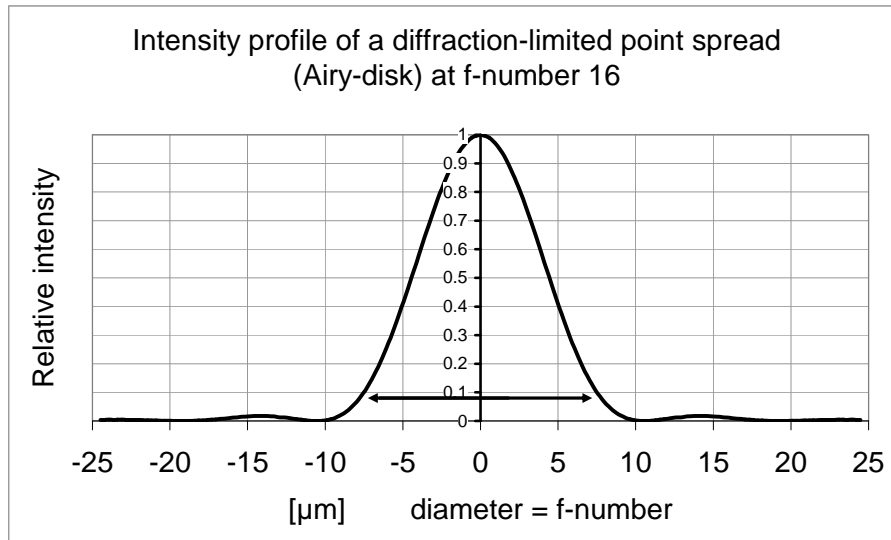


Modulation transfer function of a 50mm lens of 35mm format in the image center, measured at f/ 2 and f/ 5.6, for purposes of comparison the diffraction-limited transfer functions for f/5.6 and f/16 are also shown (solid line without circular dots). The diffraction-limited image is the best possible one. On the horizontal axis we have the spatial frequency in line pairs per mm.

A **diffraction-limited** image has an almost perfectly straight MTF curve which decreases in proportion to the spatial frequency. The zero MTF value is reached at the so-called **limit frequency**, which is determined by the f-number and the wavelength of the light.

A rough estimate for medium wavelengths of visible light is:

The width of the point spread in μm corresponds to the f-number, and the limit frequency is about 1500 divided by the f-number.



For **real lenses** with residual aberrations the MTF values initially decrease quickly and then very slowly approach the zero line. The curves therefore exhibit a pronounced sag. In the above example this can clearly be seen for the curve of aperture 2; for aperture 5.6 the lens is not very far removed from the optimum which is physically possible.


Neither should it be confused with the resolving power which is achieved in conjunction with an **image sensor**. And this leads to the above-mentioned third reason why we describe image quality with the modulation function:

The spatial frequency at which the MTF value reaches zero or falls below a low threshold (e.g. 10%) is the **resolving power** of the lens **in air**. Periodic stripe patterns can become this fine before their image changes to an unstructured gray.

We never observe the image of the lens directly with our eye, but require further links in the imaging chain: We always require an image sensor, analog or digital, or maybe a scanner, a printer or projection optics.

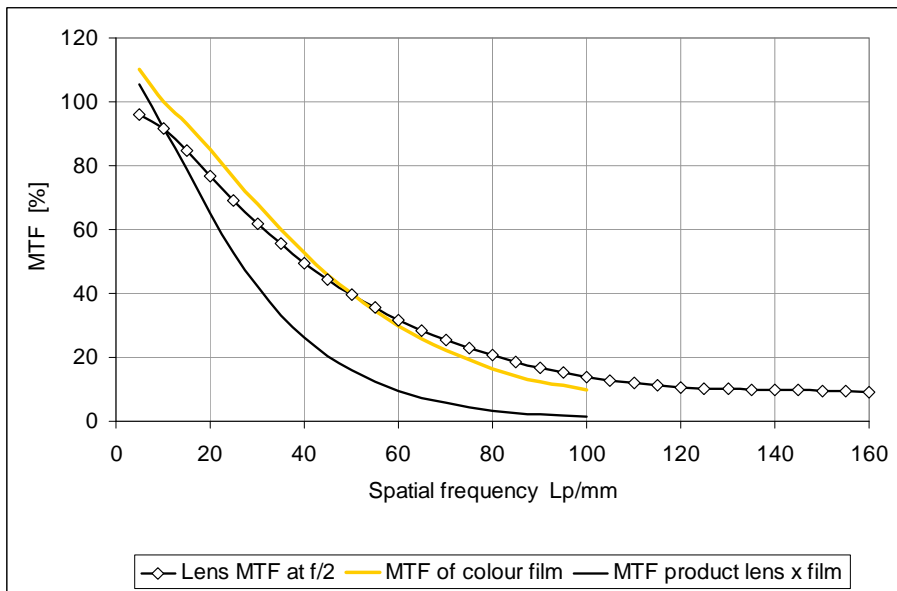
The curve above for aperture 2, in particular, shows that this resolution limit is difficult to measure; this is the case because the very flat slope at high spatial frequency values means the result depends very sensitively on the minimum contrast required. The measurement is therefore very imprecise. For this reason alone the resolving power in air is therefore not a suitable quality criterion for lenses!

All these components, even the human eye, have their own imaging properties, each of which can also be described by a transfer function. And the nice property of MTF is that the MTF of the entire imaging chain is (approximately) the **product of all individual MTFs**. Let us consider a few typical examples:



Product of two modulation transfer functions: Very good 35mm format lens and color negative film. The product is always smaller than the smallest factor in the imaging chain.

In this case, the total modulation is essentially limited by the film. If one specifies a minimum of 10% modulation transfer, one must expect a resolving power of 80-100lp/mm. If further elements such as projection optics or the eye are taken into account, the product is even slightly smaller.



Product of two modulation transfer functions: 35mm format lens with moderate performance and color negative film. The product is now determined to an almost equal extent by the lens and by the film.

If one considers the curve of the product of only two transfer functions and takes into account the fact that, in reality, even more transfer functions are involved, which can only make the product smaller, then it becomes clear that it is not necessary to consider the complete range of the very high spatial frequencies.


Digital sensors with 24 megapixels in the 35mm full-frame format or 15 MP in the APS-C format have Nyquist frequencies of about 90lp/mm. Their theoretical maximum resolutions are thus roughly comparable to the color negative film. It is therefore usually sufficient to consider the spatial frequencies up to 40lp/mm for these formats, although with larger numbers of pixels the 40lp/mm are slightly more important than usual.

Another consideration also suggests that this is a reasonable limit: If one looks at an A4-size print from a distance of 25cm, and thus sees the picture width at an angle of 60°,

the human eye can resolve 1600lp/picture height at most, because it resolves a maximum of 8lp/mm at this distance, which is called the 'least distance of distinct vision'. For the 35mm film format with 24mm picture height this corresponds to 66lp/mm. The spatial frequencies important for the human eye are therefore also in the range up to 40lp/mm.

If one considerably enlarges the image viewed, however, and nevertheless views it from a relatively short distance, the eye can use the highest spatial frequencies of the system and it suddenly sees weaknesses which would not be noticed during normal viewing of the image. This is, incidentally, what happens when one views digital pictures on a large monitor in 100% representation. In this case, the image of a 12MP camera is more than one meter wide.


Incidentally, a sensor which can also use the lens performance at higher frequencies is the low sensitivity black-and-white film:




*Good lens combined with high resolution B&W film
(this data is from T-Max 100)*

The following images illustrate graphically that 40lp/mm is already a spatial frequency which is quite high, at least for the 35mm format.

They show a picture of the well-known Siemens stars, which are used by many to test cameras. The complete image of a 12 MP camera in 35mm full-frame contains nine stars:



A greatly enlarged section shows the center of a Siemens star and how close to the center the spatial frequency 40lp/mm is:



Edge definition, image contrast

Maybe we should briefly recapitulate at this point: We know now why the modulation of sinusoidal stripe patterns decreases with increasing spatial frequency in optical imaging and also in the further stages of the image generation up to the perception. But what do these numbers tell us about the quality of real pictures? What is the relationship between terms such as **definition**, **brilliance**, **resolution of detail** etc. and these numbers?

Our subjects do not contain sinusoidal patterns, of course. They can only be generated as an approximation with a lot of effort, even in the laboratory, and instead one uses other test objects from which the sinusoidal modulation is mathematically deduced.

Stripe patterns with a rectangular intensity profile, with a sudden change between black and white,


are used on the special test charts for evaluating lenses and cameras and to determine the effective resolving power.

The modulation transfer for rectangular patterns is, incidentally, usually a little bit better than for sinusoidal patterns of the same spatial frequency. These precise rectangular shapes are also rarely found in real photographic subjects, however.

Fine periodic patterns represent only a small fraction of the subject properties which our eye uses to recognize imaging quality. Most important are really the **edges**, the borders between two areas with different brightness or color. We would therefore now like to understand what the relationship is between MTF and the reproduction of edges, and this brings us back again to our starting point, the point spread function.


*The following images show **from left to right**:*

- **Intensity profile of the point spread function**, in a logarithmic scale down to 1/1000 of the maximum intensity in the center. The width of the point spread function is given in μm , $1\mu\text{m}$ is 1/1000mm.
- **Intensity profile of two edge images** with large and small step in brightness. The vertical scale is the logarithmic aperture scale familiar to the photographer: Every graduation signifies a halving of the intensity. The horizontal scale is again a measure of the distance in the image in μm . The bright and the dark side of the edge are to the left and right respectively.
- The corresponding **modulation transfer** for five spatial frequencies 5, 10, 20, 40 and 80lp/mm is shown as a bar chart.

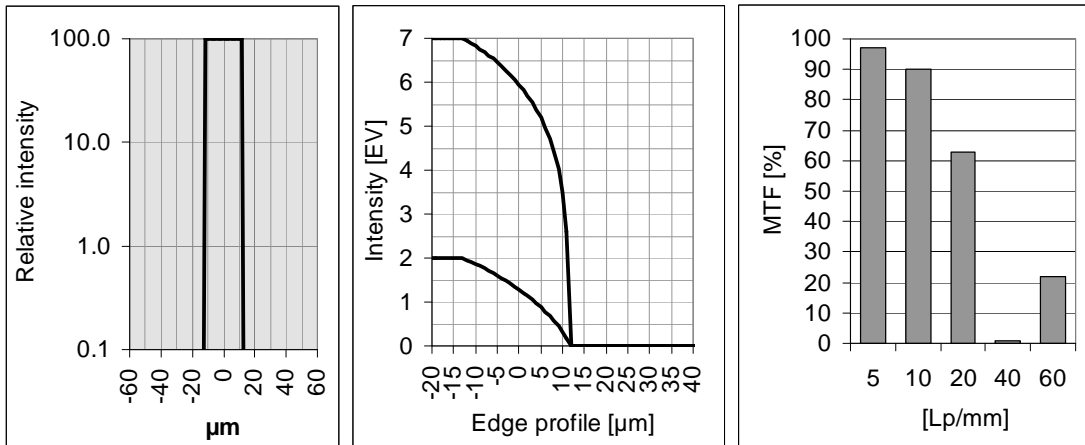


This is an example of a very good imaging performance in 35mm format; the point spread is narrow, the transition at an edge from white to black is no wider than about 10 μm , i.e. very steep. The photographer then says: The image of the edge is sharp. In the language of modulation transfer, this characteristic is recognized by the fact that all values at the important spatial frequencies are very high and do not decrease so strongly towards the higher frequencies.

For a lens with such imaging performance, the image quality achieved is usually limited by the sensor or by other factors such as focusing accuracy, camera movements etc.




Here, the diameter of the point spread function is significantly larger; the image of the edge from white to black is by no means as sharp, the edge profile is flat, because the transition from maximum brightness to black takes between 30 and 50 μm depending on the size of the brightness step. A deep black is nevertheless achieved after this distance, the contrast between the ends of the above scale is therefore high. The MTF values reveal these characteristics by quickly falling off towards the higher spatial frequencies, while differing only slightly at the lowest frequency when compared to the previous example.



A wide, box-shaped line image naturally gives rise to poorer edge definition. The MTF values for the lower and medium spatial frequencies up to 20lp/mm are normal, even at 60lp/mm there is still an acceptable modulation transfer. If one were to look only at these frequencies, one would have the impression of a quite respectable imaging performance.

But: There is no contrast at 40lp/mm! The curve of the modulation transfer can drop to zero and then increase again. This is then called “**spurious resolution**”, which is a somewhat unfortunate expression because the structure with 60lp/mm is reproduced with a clear resolution. One usually does not notice that black and white are interchanged (except with the Siemens stars) and the next zero point would come again at 80lp/mm and then a resolution again where even black and white are in the correct position again. The term ‘spurious resolution’ is intended to express that the isolated measurement of a high resolution at a single, coincidentally favorable spatial frequency can simulate an image quality which is not even present. You will not find this type of imaging in published MTF curves, but it is of practical importance with respect to focus errors and motion-induced blurring.



Here the point spread function is about as narrow as in the first example, but is surrounded by a weak halo. The edge definition is high in parts, but at the same time a broad, bright fringe stretches into the dark zone. The photographer would say that the lens exhibits flare. The contrast near the edge is low.

The MTF values of this 4th type are characterized by the fact that they decrease only slowly with increasing spatial frequency, as in the first example. But the values of the low spatial frequencies of 5 and 10lp/mm are conspicuously low.

The imaging properties of a lens with this character can be somewhat inconsistent and are judged differently depending on the image content.

Edges with low to medium contrast are reproduced with the same degree of sharpness, particularly if the exposure is rather brief. Fine structures with a lot of contrast appear a little bit flat, however, edges with a lot of contrast and lights even show flare or appear broadened at generous exposure.


Most of the high-speed standard lenses from the sixties were corrected like this at large apertures. At 10lp/mm they had only 60-70% MTF, whereas nowadays 80-90% is usual here.

At the time, one said the lens was 'optimized for resolution', which is not exactly correct, because they simply had a good edge definition, but the resolving power for fine periodic structures was no better than in lenses with a different design.

When black-and-white photography was still dominant, one could compensate the low contrast reproduction of these lenses by enlarging on paper with hard gradation. Colour photography with its less flexible laboratory processing later demanded a change in the correction to better contrast rendition.

Lenses with this imaging character are however favourable tools for particular subjects. One should always be careful with judgements about lenses.

For instance the famous soft portray lens 'IMAGON' has a modulation transfer function like this:




Incidentally, it is not the case that, when designing a lens, a decision has to be made between high resolving power and good contrast rendition; both are possible for lenses with good correction.

But what does 'contrast rendition' actually mean? We must not forget that when we talk about 'contrast' we always mean **micro contrast**, i.e. structures, which we can just about see or just cannot see with the naked eye, for example on a slide. But if we photograph a chessboard, for example, so that it fills the format, the contrast between the black and the white squares has nothing to do with this.

MTF measurements say nothing about this **macro contrast**. They gauge only the correction of the lens, i.e. the small deviations of the light beams, while the macro contrast depends on the veiling glare of the lens, i.e. on the large deviations.

These result from undesirable reflections between the optical surfaces and from light scattering at the interior barrel components, so that they usually reach the image plane a long way from the original target. All these characteristics are often mixed up with each other in the term 'brilliance of the image'. Good MTF values at low spatial frequencies are necessary, but they are no guarantee for brilliant images.



Enlargements of format-filling images of a chess board, left with perfect imaging quality, in the middle with low micro contrast, right with high degree of veiling glare.



The characteristics of the images above are also illustrated by their histograms: In the picture of the lens with poor micro contrast (middle) the right peak in particular is broadened towards the left, because the white at the edges with glare lights up the areas which are in reality black. The separation of the two peaks on the gray scale is the same as with the good picture on the left, however.

In the picture on the right with the high level of veiling glare, the lower peak of the histogram is shifted upwards because the black is lit up by the veiling glare in the whole area.

The four basic types of point spread function shown above and the corresponding MTF curves can be found in all lens data, not always in the exemplary forms shown here, of course, but usually as mixtures and combinations thereof.

From these examples we also learn that one must always consider the MTF at **several** spatial frequencies. The meaning of a value of 75% at 10lp/mm would only be completely unambiguous for the image of a sinusoidal pattern. In real images, it always also depends on what the values at 20 and 40lp/mm are.

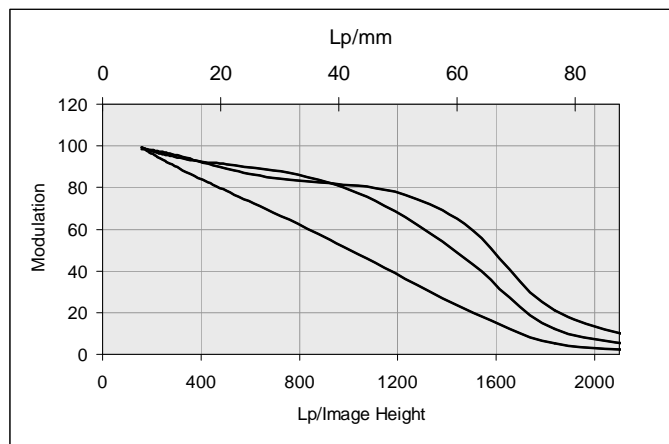
If they are very high, the lens exhibits glare at edges rich in contrast and with lights, as in our fourth example. And if they are lower than normal, the lens is simply less sharp, maybe a little out of focus, but it is free of glare.

Test procedures which measure only one point of the modulation transfer function, for example the resolution or the spatial frequency where 50% MTF is achieved, are of little value! This applies to optics just as it does to a HIFI system: If I know at which frequency the loudspeakers will have their maximum transmission or how loud 440Hz are, I still don't know how music sounds.


Edge definition in digital images

When image data is processed digitally, the transfer function of the camera can be strongly influenced. **Edge enhancement** involves making the bright side of an edge a little bit brighter, and the darker side a little bit darker. This improves the micro contrast and the edge steepness, and the subjective impression of sharpness is significantly improved, without significantly increasing the resolution of detail. This is convincing proof that definition and resolution are not the same thing.

In the transfer function this manipulation can be seen by the fact that the normal decrease with increasing spatial frequency is partially or completely cancelled, as happens with lenses with high edge definition. In digital image processing it is even possible to exaggerate this edge enhancement and to generate a transfer function which increases with increasing spatial frequency. In the language of transfer theory it then has a partial high-pass character – and such systems show marked **artifacts at edges**.



Modulation in the picture of a DSLR in the 35mm format, 24 MP, with various parameters of edge enhancement of the camera JPEG processing. The curves with a flat slope up to about 50lp/mm belong to pictures with very high edge definition.




Modulation in the picture of a 2/3" camera, with minimum, medium and maximum edge enhancement. Significant artifacts must be expected with the buckled curve, additional bright lines usually appear next to the edges of dark areas.

Tangential and sagittal

Up to now, we have been concerned with the question of the relationship between the modulation transfer and the point spread function. We have seen how the shape of the point spread function and the distribution of the light intensity within its complete area influence the modulation transfer at various spatial frequencies. For this, we have plotted the MTF as a function of the spatial frequency parameter.


Such a curve is valid only for one single spot in the image, however, and even for this spot we really require several curves,

because we have seen from our point spread examples that they are not necessarily circular. Some can be better compared with a flat brush with which one can draw fine lines in one direction only. If we rotate the orientation of the stripe pattern, we must expect different MTF curves depending on whether the shorter or longer elongation of the point spread function is perpendicular to the stripe pattern.



The main orientations, i.e. the shortest and longest elongations of the point spread functions, are always parallel or perpendicular to the radius of the picture circle, because lenses are rotationally symmetric. Stripe patterns where the longitudinal orientation of the stripes is toward the center are therefore called **radial** or **sagittal** (*sagitta* = Latin for 'arrow') in optics. This direction usually has the better modulation transfer.

Stripes perpendicular to this have, of course, the same orientation as a tangent to a circle around the image center. This orientation of the stripes is therefore called **tangential** or also **meridional**.




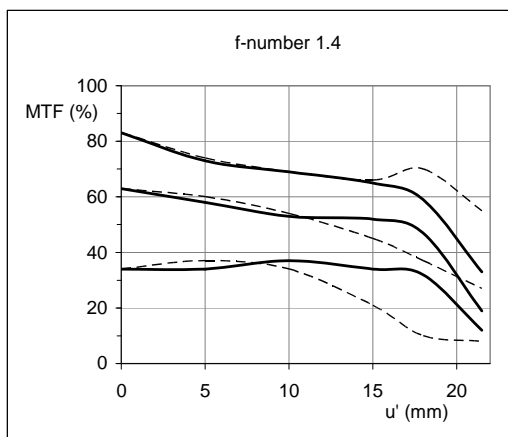
MTF curves for lenses

Since the imaging quality of lenses generally changes from the center to the edge and since precisely these differences are of special interest to us, we naturally need more curves than the two for tangential and sagittal orientation. Half a dozen or so measuring points between the center and the corner are required in order to be able to describe the spatial changes of the imaging properties with sufficient precision. This would be 12 curves in total – not very clear and legible.

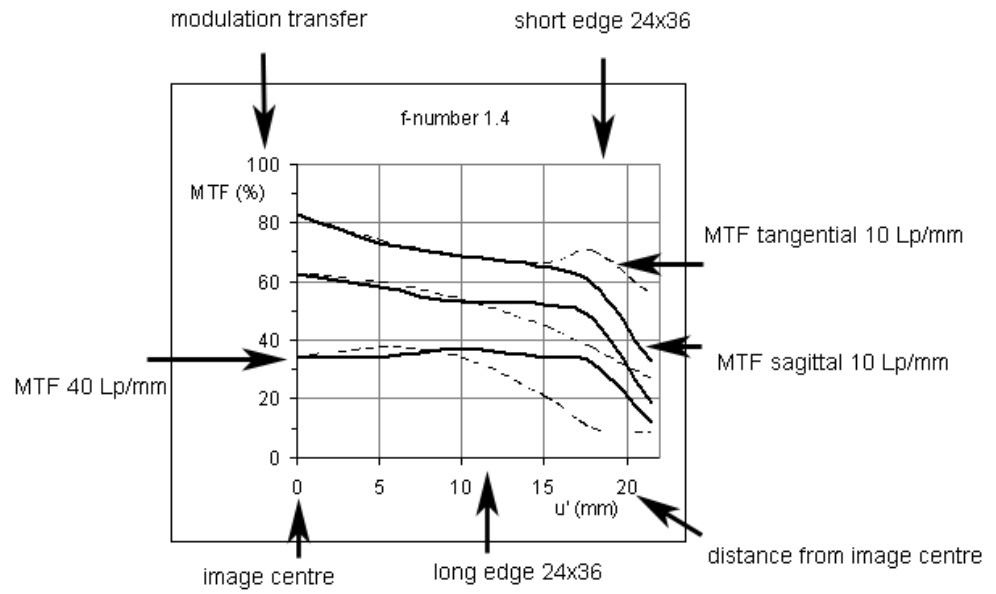
The MTF curves which we have got to know so far, where we have plotted the modulation transfer on the vertical axis and the spatial frequency on the horizontal axis, are really only suitable for sensors where there are no spatial variations. This representation is not so suitable for lenses.

Since the MTF curves as a function of the spatial frequency always slope to the right and fall off more or less rapidly, it is sufficient to read off only three numerical values from each curve, i.e. from three sufficiently different spatial frequencies, usually 10, 20 and 40lp/mm. If one shows how these MTF values change in the image area for three frequencies, one obtains a graphical representation which is much better suited to lenses.

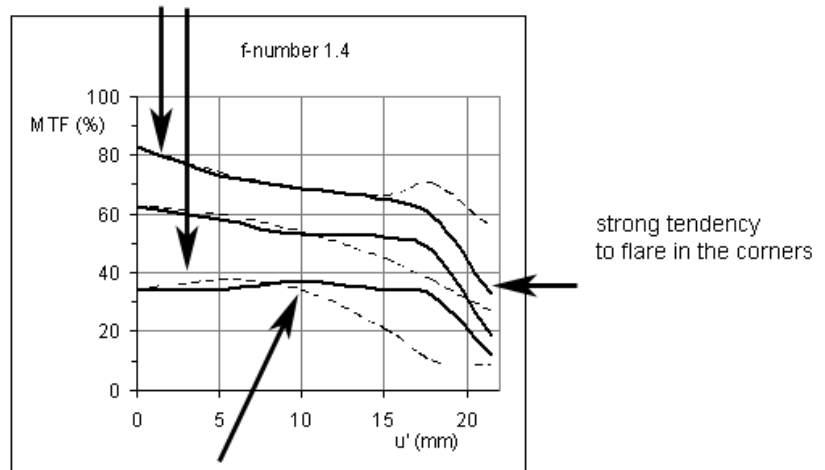
This is why you will find MTF curves in our datasheets, where the modulation transfer is plotted on the y-axis and the image height, the separation from the optical axis, on the x-axis. The diagram contains six curves, i.e. the tangential (broken lines) and sagittal values (solid lines) in each case for three spatial frequencies. The top curve of the six always relates to the lowest spatial frequency, of course, and the bottom one to the highest.



MTF curves for the Planar 1.4/50 ZF lens, at 10, 20 and 40lp/mm, white light and distance to object on infinity.



Good contrast and very acceptable sharpness in the centre



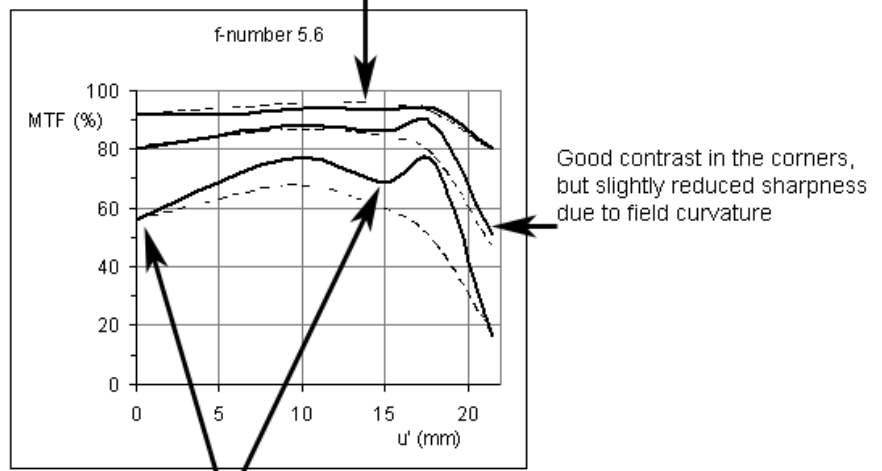
Good edge sharpness, moderate micro contrast, high contrast edges with flare in the outer parts of the frame

In the center, this lens achieves more than 80% MTF at 10lp/mm even at full aperture and decreases to just below 40% at 40lp/mm. This means good contrast rendition and medium definition, which has a slightly soft effect only after the picture has been greatly enlarged.

Away from the center the MTF decreases at 10lp/mm to 70%, the tendency to produce flare at edges rich in contrast therefore increases.

In the corner of the image the sagittal curves, in particular, are very close together at low level, we must therefore expect significant glare with open light sources at the corners.

MTF at 10Lp/mm in nearly the whole frame above 90%, above 50% at 40 Lp/mm, that means images in rich contrast, well defined, excellent image quality





curve for 40Lp/mm varies a little, while 10Lp/mm remains flat, this points to small focus shift and residual field curvature

If the lens is stopped down, all MTF values increase greatly; the curves are very close to each other at a high value. The MTF values therefore decrease only relatively slowly with increasing spatial frequency. This means excellent edge definition and very good micro contrast right down to the finest structures which can be reproduced by the sensor or the film.



due to the field curvature which suddenly appears.

In the corner of the picture, all the curves tail off somewhat, those for 10lp/mm a little, those for the higher frequencies more. This indicates that the very good flatness of the visual field stretches to about 18mm picture height and that there is a defocusing in the corner of the picture


One should not take the small variations of the curves for 40lp/mm too seriously, they are visible only at extreme enlargements of the picture and when taking photographs of flat objects, in most pictures they are therefore invisible. They are caused by field curvature and shifting of the focus. I will explain why these lead to such variations in the section on the three-dimensional characteristics of MTF.



A comparison of short Tele focal length for 35mm format, stopped down to $f/5.6$. On the left a high quality prime lens (Planar 1.4/85 ZF), on the right an inexpensive 5x zoom lens. Image quality with the prime is in the whole frame practically limited by the sensor and allows highest magnifications. The zoom lens is quite good in the centre, but drops continuously to the edges. Except from the corners a good overall contrast may be expected, but the lens lacks a ‚biting‘ sharpness, since the MTF values drop rapidly at higher spatial frequencies. The lens could only be recommended for moderate enlargements.




A comparison of two super-wide lenses, which are much more difficult to make, stopped down to $f/5.6$. On the left data measured with a Distagon 2.8/21 ZF, on the right a lens, where the lateral chromatic aberration is corrected less good. Its sagittal MTF shows some focus shift, but is otherwise not too bad. But the tangential MTF is very low towards the edges of the frame; and what this means can be seen in the following two thumbnails (200x200 pixels from a 12MP-image, image height about 12 mm):



Three-dimensional characteristics

It is, of course, a truism that image definition also depends on whether the lens is focused correctly. It should thus also be possible to describe this with MTF curves; and therefore you are now being introduced to a third type of MTF curve which is not so commonly known.

Here the MTF values are not plotted as functions of the spatial frequency or the picture height, but of a focus parameter. For this we measure how the MTF changes in the longitudinal direction on the image side of the lens and thus obtain the following curves:



The MTF for 10, 20 and 40lp/mm is again plotted on the vertical axis. The zero point on the horizontal axis corresponds to the best focus: The MTF value for the medium frequency 20lp/mm is a maximum there, so this is where the sensor or the film should be positioned, as symbolized by the yellow line. To the left we are closer to the lens, to the right we are behind the sensor.


Incidentally, you can see that the tolerance range for using the best MTF values at this aperture is only a few hundredths of a millimeter. The two black triangles show the depth of focus on the image side, calculated purely geometrically for a circle of confusion of 0.03mm diameter.

Incidentally, with this criterion for the depth of focus, the MTF values at the limit of the area which counts as focused are around 20% at 40lp/mm.

Incidentally, it is quite possible for the maxima for different spatial frequencies to be at different positions. And the curves are often skewed, which means that the type of blurring is different in front of and behind the focus.

So what can happen when one stops down the lens aperture? We decrease the aperture of this lens by three stops and repeat the measurement, but do not change our focusing scale, i.e. the zero still means:

MTF maximum in the center of the picture at 20lp/mm and f-stop 1.4.




The maximum values of the MTF curves increase significantly, of course, because stopping down greatly reduces the residual aberrations. At the same time, however, we see a shift of the curves to the right, i.e. further away from the lens. The lens is now not at all optimally focused onto our sensor position (yellow line), the MTF increase does not take full effect at that position. The geometrically calculated depth of focus on the image side is quite wrong, its length is still alright, but the position is wrong.

This phenomenon is called 'focus shift'; it is usually more marked with extremely high-speed lenses and is connected to the spherical aberration, because this means that beams of light which pass through the aperture area at different distances from the optical axis have a different focus.


The focus shift here is about 0.05mm. The black dots in the above graph show how this shift in the image space is connected to the distances in the object space in front of the camera (scale on the right hand side). If the lens was originally focused with aperture 1.4 at 3m distance, for example, the best focus has now moved to 3.25m if the lens settings are not altered.


Should this focus shift be corrected when taking a photograph? Not really, unless the best possible performance really has to be in the center of the picture. But 0.05mm is about 20% of the distance between the markings of the depth of focus scale for aperture 4, not so easy to control, therefore. And anyway, it is a completely different story in the other parts of the frame. We therefore again measure the MTF in the longitudinal direction, not in the center of the picture, but at a distance of 10mm.



Away from the center of the picture we still have to distinguish between tangential and sagittal orientation; in order to keep the graph clear, the curves for 10 and 20lp/mm have now been left out.



We can now see that both curves are shifted less and even to the left. The position of the maximum thus moves when we move in the image area.





Here we have reached the edge of the image at a distance of 18mm from the center of the image and we can see that the sagittal maximum has now returned precisely to the zero of our focusing scale. The field curvature should therefore not be imagined as a uniform curvature of the image area, but there are reversal points.

This combination of residual field curvature and focus shift leads in any case to the fact that the MTF curves for the same lens can look completely different if we do not focus onto the local maximum for every picture height, but measure strictly in a fixed plane:



So these two graphs do not mean: *“The left lens is a little worse in the center than in the field, the right lens by contrast is very good in the center of the picture, but has a significantly deteriorating picture definition in the zone around 15mm picture height.”*

Both measurements are from the same lens, but were taken with a slightly different focus. This difference of 0.05mm is of the same order as conventional mechanical camera tolerances such as adjusting the AF and the focusing screen.

Limits to the significance of MTF curves

The relationships presented in the last section are a suitable transition to now move on to talk about the limits of this world of numbers. If the shape of the curves is so sensitive to small changes in the focusing, one cannot, of course, expect to recognize the curves in every picture if the subject is three-dimensional, i.e. when even the different distances cause some details to be focused well and others not so well.

The measurement conditions of MTF curves are comparable to reproduction photography, where one plane is strictly imaged onto another plane. The only other member of this category is the photography of subjects which are very distant with short focal lengths.

The reasons why the scale does not match our perception are partly that the MTF curves of lenses describe only the first link of the imaging chain, of course, and do not take into account those that follow. Sensor, scanner, projectors, the eye, in short everything afterwards also always has a transfer function which decreases towards high spatial frequencies. And thus they lead to a decrease in the variations of the lens at high spatial frequencies because all transfer functions are of course multiplied. Let us take slide projection as an example: The eye does not see what the 40lp/mm curve is doing when one is sitting behind the projector. A further cause is that our eye's logarithmic perception of brightness is not taken into account.

The scales of the measured quantity MTF do not do justice to our perception. Some experience is needed to be able to translate the graph of the curves into a prediction of the subjective perception of the picture. One must take the viewing conditions into account; there is a difference whether one looks at an A4 print or at a significantly larger 100% representation on a large monitor from the same distance.

The graphic appearance of the MTF curves usually leads to the significance of the 40lp curve for normal picture sizes being overrated and the significance of the 10lp curve being underrated. If one looks at a projected picture, for example, from about the projector distance for conventional projection focal length, then the normal human eye can resolve about 20lp/mm at best from a 35mm format.

Thus there have already been many investigations regarding the question of how to translate MTF measurement data into a scale which is related to our perception, including the **Heynacher numbers** used at Zeiss, and other values based on psycho-physical factors such as **SQF** (subjective quality factor), **MTFA** (modulation transfer area), **SQRI** (square root integral). Their common feature is that they all compute areas below the 'modulation over spatial frequency' curve.

A further common feature is therefore that they all attempt to describe the quality at one image point by one single number. As we have seen above, this is, of course, sometimes an undue simplification of the data.

Unfortunately, the scope of this article precludes me from going any deeper into this.

Phase transfer function

The desire for simplification is also the reason why, until now, I have withheld something from the reader: The MTF values are nowhere near the whole truth about the correction state of a lens. But no-one should really be surprised that a system as complex as a lens cannot be completely described by only these few numbers. The performance data of a lens, irrespective of whether it is calculated by computer or measured in the laboratory, fill a small file.

Simplification is necessary in order to make it digestible and clear, but one has then to put up with the fact that the precision of the description suffers.

So much to the introduction, and now to more concrete details again: It is possible that two lenses which have the same MTF data produce quite different images of one detail of the subject, not randomly, but systematically – here is an example:



Images of two high-speed wide angle lenses at full aperture, details near the edge.

The picture shows the roof of a house and a tree in front of a bright sky, i.e. a typical picture of a horizon which is rich in contrast. The MTF at low spatial frequencies is particularly important at the edges of the dark foreground objects, because it determines the amount of glare at these edges. In the picture on the left, the roof exhibits no glare, but the tree does, in the picture on the right it is the other way round. If there were no tree in this picture, one would judge the picture on the left to be the better one (in black-and-white at any rate). At these edge picture heights both lenses have the same MTF values for all spatial frequencies, however.

The really complete **optical transfer function OTF** also has a second part, the **phase transfer function PTF**, which is usually neglected. It has something to do with the symmetry of the point spread function.

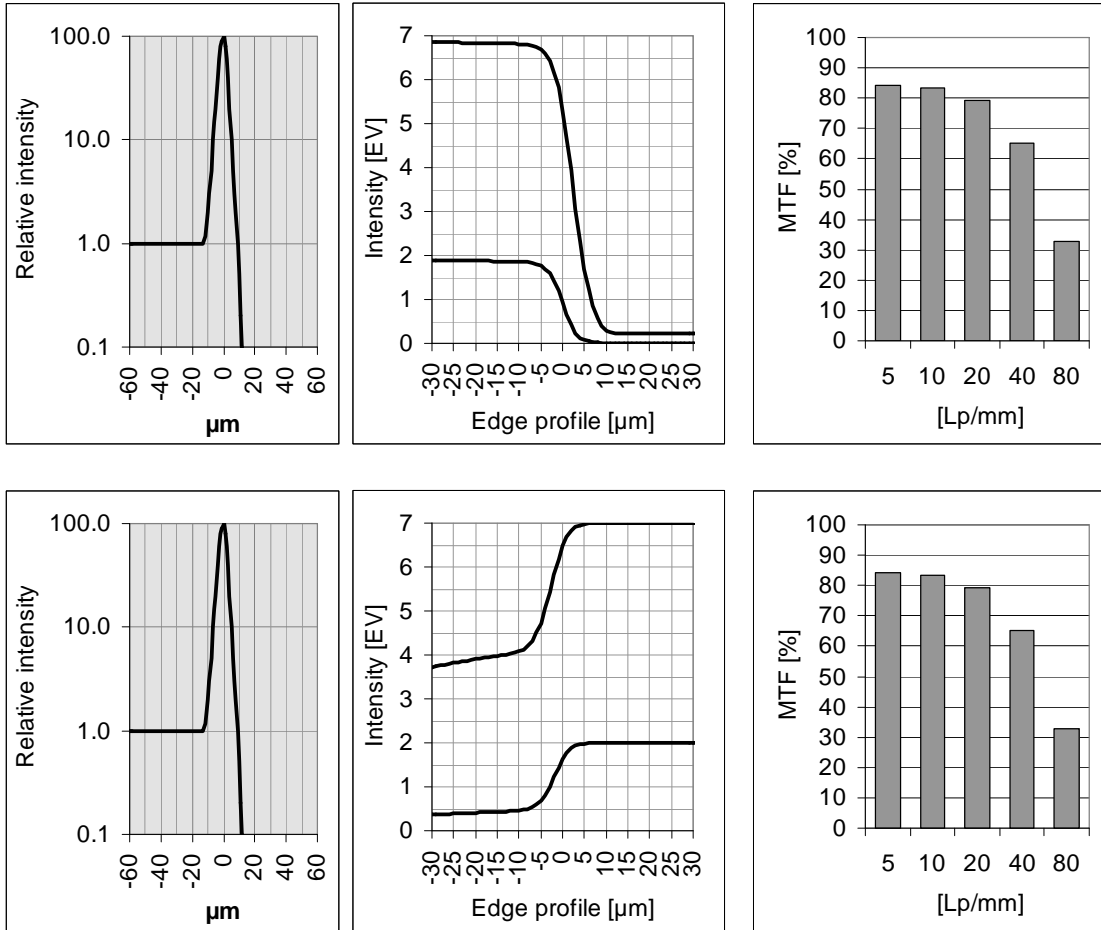
We did take into consideration the fact that the point spread functions can be elongated, that they therefore have different spreads in the tangential and the sagittal directions. We therefore measure two MTF curves for each image point.

The **MTF** does not tell us anything about this difference, because it does not yet completely describe the characteristics of the point spread function.

In the previous examples, however, we have tacitly assumed that the brightness distribution is symmetric in one cross-sectional direction of the point spread function. In reality, this is often not the case, however.

Point spread functions can be as skewed as in the following example. The most frequent causes are coma errors

which produce point spread functions with a tail in the radial direction.



The orientation of an edge is very important for such a skewed point spread function intensity profile, of course. This point spread function has a halo of 1% of the maximum intensity on the left, on the right it stops suddenly. If the bright side of the edge is to the right, it will produce glare to the left (bottom). If, however, the opposite is the case and the left side of the edge is bright (top), then the contrast of the edge image is high because the point spread function only extends a short distance to the right.

The MTF values do not take account of this dependence on orientation. This is contained in the phase transfer function, which differs depending on the orientation of the "tail" of the point spread function. The name stems from the fact that such a skewed point spread function shifts the phase, i.e. the position of its maxima and minima, of the sinusoidal pattern sideways.



Color correction

The fact that the optical characteristics of glass depend on the wavelength of the light can also be seen in our pictures: Lenses have color aberrations. While it is true that each lens has a sophisticated compensation system which uses a combination of different types of glass, so that this type of aberration is usually no longer critical, some residual aberration is still present.

There are lenses where the color aberrations are more critical; this is mainly at long focal lengths, where only very recently has it been possible to significantly improve the image quality thanks to the development of completely new types of glass.

Long telephoto lenses without these types of glass having extremely low dispersion or anomalous partial dispersion have only mediocre MTF values. It is nevertheless possible to achieve astonishingly good imaging results for many subjects.

This is because the MTF of these lenses is strongly dependent on the spectral composition of the light. If the measurement is done with green light instead of the conventional white light, where all wavelengths of the visible spectrum are present with a certain weighting, the MTF curves are dramatically different:




MTF curves of a 300mm telephoto lens, on the left measured with white light, on the right with green light, 100nm bandwidth.

This is why green filters were an important accessory in the time of black-and-white photography. This same effect can also be achieved in color photography if the subject is predominantly monochrome (nature photographs, red roofs). This is yet another reason why the imaging performance is not completely represented by MTF curves.

But it is not always the case that MTF curves gauge a lens too pessimistically. On the contrary, it is possible that a weakness in the color correction is not visible in the MTF data for white light. In other words: MTF says little about color fringes.

Comparing MTF in white and in coloured light helps to understand the reasons of colour fringes in images of high contrast edges and highlights.

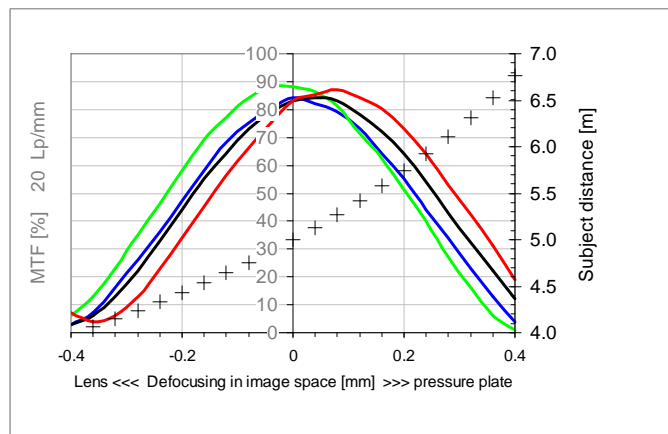
The following curves illustrate the longitudinal chromatic aberration of a high speed short telephoto lens by measuring MTF as a function of focus:



Focus MTF of the wide open Planar 1.4/85 ZA in white light (black curve) and in blue, green and red light. The cross symbols connect position on the image side (horizontal scale) to the subject distance (vertical scale on the right), the lens has been focused in white light to a distance of 5m.

MTF values in coloured light are higher than in white light, but at the same time the maximum is at different positions, they don't have a coincident focus. In the best focus for white light (position 0) the red light MTF is the lowest of all. From that follows, that the red line spread has the largest diameter; the image then shows a slight reddish fringe. This is getting even stronger, when the subject is at slightly shorter distance, where the green MTF is at its maximum. Thus this kind of fast lenses produces fringes at highlight details which are red or purple, if the detail is in front of the focal plane, and which is green,

if the detail is behind the focal plane. The saturation of these colours, called secondary spectrum, depends on the distance of the MTF peak positions and on the slope of the focus MTF curves. If lenses exhibit more monochromatic aberrations (like old lenses), the curves are more flat and the colours look pale. Just modern, highly corrected fast lenses tend to show more saturated colours. Since the distances between the peak positions can't be made infinitely small, the only chance to make the fringes disappear is stopping down, since then the depth of focus is large compared to the longitudinal colour aberration, and coloured MTF differences are getting small.



Focus MTF of the Planar 1.4/85 ZA at f/ 5.6

Bokeh

Curves where the tangential and sagittal values are nearly identical in the whole visual field are often called ideal MTF curves because, in these cases, the “bokeh”, i.e. the representation of the markedly defocused background, is particularly good.

Such statements should be regarded with caution. MTF only makes statements about the focal plane or its immediate surroundings. And in that case, a circular point spread function is indeed an advantage, because it reproduces small details in a way which is as faithful to the original as possible, with the best trueness of shape. This is important for the legibility of writing, for example.

It is not possible to use MTF data to draw conclusions about the brightness distribution of the strongly defocused point spreads, however. There are lenses with nicely parallel tangential and sagittal MTF curves but which are spherically strongly overcorrected. This correction state causes annular defocused point spread functions, which are visible as rim-lights and as pairs of lines and produce a restless-looking background.

This unpleasant characteristic cannot be deduced from the MTF data.

Comparability of MTF data

MTF data is published in many publications, by manufacturers and now also in many independent tests. Unfortunately, one has to be very cautious when comparing this data, because the measurement conditions can vary greatly.

The least of the problems would be to overlook the fact that the spatial frequencies are different. And, likewise, different spectral weightings of the visible light can also leave a comparison wanting. There are also manufacturers who are not afraid to publish data which are better than diffraction limitations allow, i.e. it is physically impossible.

This tells you that these values are merely computed and that this was done by taking only geometrical optics into account, without considering the wave nature of light. If the lenses have excellent correction, the values stick to the 100% line. But, please, don't believe that these numbers are realistic. Real lenses are always a little worse than the calculation of the optical design program.

MTF data published by Zeiss always originates from measured lenses.