

Measuring the spatial resolution of an optical system in an undergraduate optics laboratory

Calvin Leung and T. D. Donnelly

Citation: *American Journal of Physics* **85**, 429 (2017); doi: 10.1119/1.4979539

View online: <http://dx.doi.org/10.1119/1.4979539>

View Table of Contents: <http://aapt.scitation.org/toc/ajp/85/6>

Published by the [American Association of Physics Teachers](http://www.aapt.org)



American Association of **Physics Teachers**

Explore the **AAPT Career Center** – access hundreds of physics education and other STEM teaching jobs at two-year and four-year colleges and universities.

<http://jobs.aapt.org>



Measuring the spatial resolution of an optical system in an undergraduate optics laboratory

Calvin Leung and T. D. Donnelly

Department of Physics, Harvey Mudd College, Claremont, 91711, USA

(Received 13 June 2016; accepted 12 February 2017)

Two methods of quantifying the spatial resolution of a camera are described, performed, and compared, with the objective of designing an imaging-system experiment for students in an undergraduate optics laboratory. With the goal of characterizing the resolution of a typical digital single-lens reflex (DSLR) camera, we motivate, introduce, and show agreement between traditional test-target contrast measurements and the technique of using Fourier analysis to obtain the modulation transfer function (MTF). The advantages and drawbacks of each method are compared. Finally, we explore the rich optical physics at work in the camera system by calculating the MTF as a function of wavelength and f -number. For example, we find that the Canon 40D demonstrates better spatial resolution at short wavelengths, in accordance with scalar diffraction theory, but is not diffraction-limited, being significantly affected by spherical aberration. The experiment and data analysis routines described here can be built and written in an undergraduate optics lab setting.

© 2017 American Association of Physics Teachers.

[<http://dx.doi.org/10.1119/1.4979539>]

I. INTRODUCTION

What is meant when we say a camera is high quality? We might mean that the camera can do justice to an unevenly lit scene, or has a high dynamic range. Maybe the camera's sensor has many pixels and thus round objects do not appear pixelated. Or perhaps we mean the camera responds consistently to different colors of light, or has very little chromatic aberration. Maybe we mean to say that the camera has good spatial resolution, that it does a good job of reproducing small objects, sharp edges, and fine detail.

Spatial resolution measurements are a standard way to test the quality of an optical system. Having a resolution-measurement testbed in an undergraduate optics lab is an interesting and relevant way to learn about optics and electro-optical systems. For the hobbyist, the curious student, or anybody with a smartphone camera, it is interesting and worthwhile to be able to quantitatively compare different imaging options on the market. In particular, off-the-shelf DSLR cameras, such as the Canon 40D, are compelling systems to study because they are ubiquitous, the technological bang-for-the-buck is very high, and there is a great deal of consumer interest in selecting appropriate cameras and lenses. In this paper, we explore two ways in which a camera's spatial resolution can be measured and investigate the dependence of the resolution on various parameters of the camera. A widely accepted method for quantifying the performance of imaging systems is through the use of a set of resolution test targets. Another more sophisticated method uses Fourier analysis and measures the camera's response to the different spatial-frequency components of a known signal. Using the modulation transfer function to quantify resolution illustrates the power of using the discrete Fourier transform as a robust tool to extract subtle patterns from spatial data.

In the remainder of this paper, we first motivate the study of spatial resolution by discussing scalar diffraction theory, which predicts fundamental physical limits on the spatial resolution of any imaging system. Then, we introduce the experimental setup we will use and present procedures and

the necessary mathematical formalism for making measurements via both techniques. A set of measurements is made via both of these methods. The data analysis procedures are described in detail and the results are compared. Finally, as an application of this work, we characterize the Canon 40D and perform two tests that confirm that it is not diffraction limited.

There is a vast literature on spatial resolution, Fourier transforms,¹ and methods of quantifying optical systems, both in the peer-reviewed literature and in textbooks.² The importance of helping students to master the most relevant pieces of this literature through an improved upper-division undergraduate optics laboratory curriculum was highlighted over two decades ago.³ However, with regard to treatments of these topics that are appropriate for an upper-division undergraduate laboratory, the previous work on the spatial resolution measurement usually focuses on measuring the modulation transfer function of a system. In most cases, the literature assumes a strong background in Fourier optics and signal processing.⁴ In other cases, the techniques described are highly specialized⁵ or no longer applicable due to advances in computing power.⁶ To our knowledge, this manuscript is the first self-contained treatment of spatial-resolution measurement methods developed with the modern-day undergraduate in mind. In particular, we are careful to discuss experimental difficulties that specialists take for granted. An undergraduate in an upper-division optics laboratory can become fluent in the basic theory and practice of various methods of spatial resolution measurement with the guidance of this manuscript.

II. WHY SPATIAL RESOLUTION IS IMPORTANT

Spatial resolution is fundamentally limited by diffraction in any optical system. This limit comes from the diffraction of electromagnetic waves propagating through a finite aperture. Traditional imaging cannot overcome the diffraction limit, but optical engineering *can* design a system that balances cost and performance to get as close as possible to that limit for a given camera setting. One way, therefore, to judge

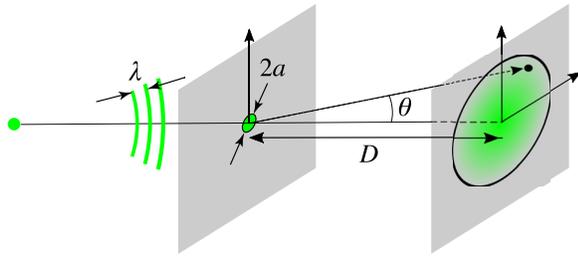


Fig. 1. An aperture (left) diffracts light to produce an image on the screen (right). If the screen is far away from the diffraction aperture, with $D \gg a$ and $D \gg \lambda$, a Fraunhofer diffraction pattern will be visible on the screen. The dark circles roughly indicate the relative sizes of the aperture and the diffraction pattern but are not drawn to scale.

the quality of a camera is to see how closely it performs to the diffraction limit.

The diffraction limit can be illustrated by studying the system depicted in Fig. 1, consisting of a point source of monochromatic light, a circular aperture of finite radius a , and an image plane some distance D away. A well-known result of diffraction theory is that small apertures act to blur sharp edges and smear point sources. This means that the image of the point source formed on the other side of the finite aperture fundamentally cannot be a point.

In fact, in the limit that the point source is far away, the intensity profile due to the diffraction of light can be analytically calculated. If the light has wavelength λ , its image formed on a faraway image plane takes the form

$$I(\theta) = 4I_0 \left[\frac{J_1(2\pi a \sin \theta / \lambda)}{2\pi a \sin \theta / \lambda} \right]^2, \quad (1)$$

where θ (the independent angular variable) and a (the aperture radius) are shown in Fig. 1, and J_1 is the first-order Bessel function of the first kind. This intensity distribution, plotted in Fig. 2, was first calculated by George Airy⁷ and is thus referred to as an “Airy disk.” From Eq. (1), it is evident that shortening the wavelength of the light or increasing the aperture size (increasing the ratio a/λ) both make the resulting diffraction pattern resemble a point source more closely, improving the resolution.

Clearly, diffraction poses a problem to scientists who study increasingly small systems or increasingly distant systems with microscopes and telescopes. Lord Rayleigh addressed the difficulties posed by diffraction by quantifying the resolution limit of optical systems due to diffraction.⁸

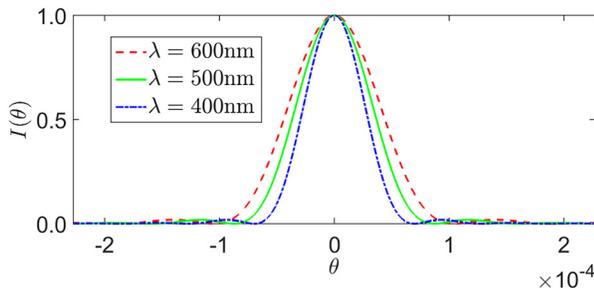


Fig. 2. The Airy disk is plotted for three point sources of different visible wavelengths. An aperture diameter of 3.5 mm, which is typical of a commercial DSLR camera and is achievable on the Canon 40D, is assumed. Note that the intensity function does not monotonically diminish outwards from the central maximum but rather oscillates as it vanishes.

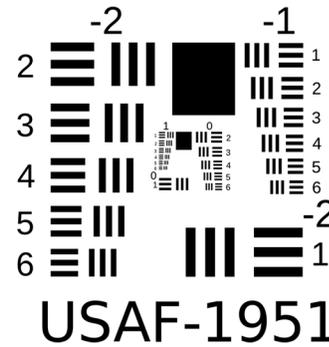


Fig. 3. The 1951 US Air Force Resolution Test Chart (Ref. 17). The horizontal/vertical line pairs are arranged in groups of six targets each. Each group has a group number vertically above or below the target. The six targets within a group are numbered, with the element number being horizontally adjacent to each target. This numbering scheme makes it possible to quickly quantify the resolution limit of a camera out in the field without having to worry about spatial or angular frequencies.

The Rayleigh Criterion says that two point sources are “just resolvable” if the Airy disk of one has a maximum at the first minimum of the other. If the Airy disks are separated any further than this, they are resolvable as two distinct point sources. If the Airy disks are any closer together, they appear as a single blur and are deemed not resolvable under the Rayleigh Criterion.

III. TEST TARGETS (CONTRAST TRANSFER FUNCTION)

Another experimentally straightforward way of quantifying resolution involves imaging test targets, which are typically objects upon which are printed lines of well-defined and varying sizes and separations, as shown in Fig. 3. An optical system’s resolution can be measured by imaging the alternating light and dark lines at successively finer spatial scales, as displayed in Fig. 4. The spatial scale at which the line pairs become indistinguishable defines a resolution cutoff for a particular camera. The resolution cutoff can be reported as a quantitative basis of comparison between different cameras.

The most direct way to report the resolution cutoff is by measuring the line spacing of the test targets with a pair of calipers, in line pairs per millimeter (lp/mm). However, instead of directly reporting a spatial frequency, it is often more convenient to report angular spatial frequencies,² such that the specified cutoff is defined independently of the target



Fig. 4. An image of the test chart, taken at $f/8.0$. The data plotted in Fig. 5 is taken by sampling a horizontally oriented row of pixels across vertically oriented targets within this image.

distance used during testing. While one may be tempted to measure and report angular spatial frequencies in inverse radians, it is more convenient to measure angular spatial frequencies as a fraction of the total angle subtended by the camera's field of view. Hence, the units of angular spatial frequency are "lines per picture width" (LPPW), where a "picture width" is not a unit of distance but rather the *angle* subtended by the camera's full field of view. Calculating the angular spatial frequency ζ of a set of lines in LPPW is straightforward for digital images, which can be quantitatively manipulated as large matrices; we simply divide the total number of pixels subtending the field of view by the imaged width of the line (also measured in pixels).

For example, the camera sensor in the Canon 40D subtends some angular field of view that spans 3888 pixels measured from left to right. The intensity profile plotted in Fig. 5, taken from a row of pixels in Fig. 4, cuts across three dark and two bright lines in the image. We can see that for this particular set of bright/dark lines, a "line pair" spans 110 pixels. Hence, the angular spatial frequency for this particular set of lines in the test target is

$$\zeta = \frac{3,888 \text{ pixels/picture width}}{110 \text{ pixels/2 lines}} \approx 71 \text{ LPPW.} \quad (2)$$

In addition, with digital imaging and image-processing software, we no longer have to rely on a resolution cutoff beyond which we deem the lines blurry. Instead, we can take a set of pixels such as those plotted in Fig. 5 and calculate a measure of resolution called the contrast

$$C = \frac{I_{\max} - I_{\min}}{I_{\max} + I_{\min}}, \quad (3)$$

as a function of the spacing of the line pairs or the angular spatial frequency. As the lines get more closely spaced and our optical system has trouble resolving individual lines, I_{\max} and I_{\min} tend towards each other and the contrast tends towards zero. A contrast closer to 1 corresponds to well-resolved line pairs, while a contrast of 0 implies that $I_{\min} = I_{\max}$ (or that there is absolutely no spatial variation in the image intensity).

Measuring the contrast as a function of the spatial or angular frequency gives us an elegant and quantitative way to represent the resolution of a system. In fact, the contrast as a function of ζ defines the "contrast transfer function," or CTF,² as plotted in Fig. 6.

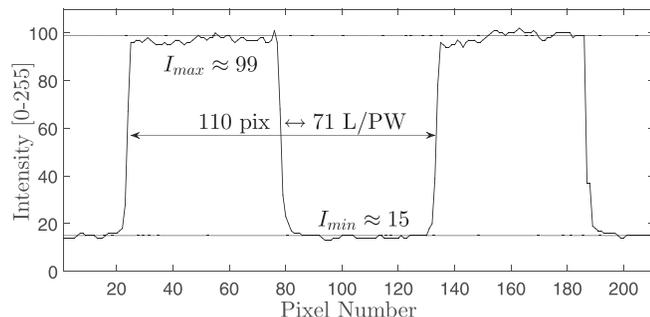


Fig. 5. A horizontal row of pixels is sampled from Fig. 4. The pixel intensity values are plotted as a function of the pixel index. The contrast is measured by estimating the maximum and minimum intensity levels as indicated by the horizontal lines. In this example, $C = 0.74$.

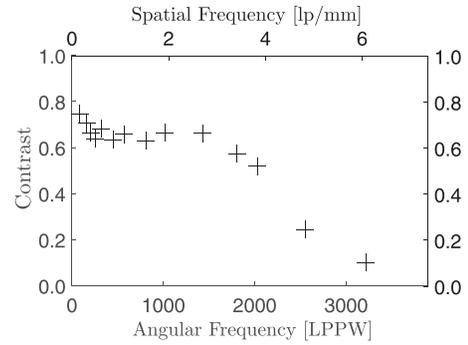


Fig. 6. The contrast is plotted as a function of the angular spatial frequency (bottom axis), which is linearly proportional to the spatial frequency of the targets on the screen, and which can be measured easily with the use of calipers. As one would expect, the contrast decreases as the angular frequency grows large, in correspondence with the observation that real cameras blur fine details.

A. Image acquisition and data analysis

To measure the contrast transfer function, we image a printed version of the test target (shown in Fig. 3) using the Canon 40D connected to an EF 28–135 mm $f/3.5$ – 5.6 lens.

The primary experimental difficulty with test targets is ensuring that all areas of the test target are evenly illuminated. A higher background illumination increases I_{\min} and decreases the contrast artificially due to glare. Since the line pairs are arranged on the test target in a spiral fashion, spatial variations in the illumination of the test target can make it seem like the CTF repeatedly increases and decreases. Thus, to eliminate spatial variation in the lighting of the test target due to glare, the entire setup is illuminated with ambient light rather than the camera's built-in flash.

In addition, mechanical vibrations can degrade the measured resolution of the camera, especially for long exposures. Mounting the camera on a tripod or breadboard and using a remote trigger reduces any resolution degradation due to mechanical vibrations as to be negligible compared to the contributions from optical aberrations, diffraction, etc., of which we are interested. In our setup, the camera is mounted on a breadboard at a fixed distance of about 0.5 m from the test targets. Great care should be taken to ensure that at least the smallest line pairs are in the center of the camera's field of view so that we can compare the test target method to other methods of resolution measurement in the high-frequency limit. There exist open-source software packages that facilitate computerized control of the Canon 40D over a USB transfer cable. Having a live stream from the camera sensor displayed on the computer facilitates alignment and ideal focusing of the optics.⁹

Once the setup is constructed and the camera aligned and focused, it is important to choose the camera settings well. For example, high ISO increases shot noise on the sensor, while long exposure times coupled with mechanical vibrations can degrade the image resolution.² Since these two parameters both affect the exposure of the image in a well-understood way, it is possible to experiment with a good combination of ISO and shutter speed that balance electronic noise and mechanical noise. For the best signal-to-noise ratio, we adjust the ISO and shutter speed to maximize the intensity throughout the image without saturating the sensor.

The camera's aperture size, often described in terms of its f -number, can affect a camera's resolution limit in a number

of ways. By decreasing the f -number, one opens the aperture wider. This lets in light from a wider range of angles than only near the center of the lens, where lenses typically perform best. In addition, light incident on the detector at an oblique angle increases the probability of pixel cross-talk, which worsens the resolution. On the other hand, the diffraction limit increases as the aperture radius a is increased, as seen earlier in Fig. 2. We fix the aperture diameter to $2a = 3.5$ mm when using test targets.

When the shutter button is pressed, light incident on the camera's sensor is captured and stored in a raw format. Then, a slew of (potentially proprietary) algorithms are applied to the raw sensor output to denoise and enhance the image, which then is converted to a JPG image. While it can be interesting to analyze the sensor output in the minimally processed raw format, it is equally meaningful to simply study the final JPG such as that in Fig. 4, taking into account all of the processing done by the camera.

MATLAB is used to process the image as a large matrix of pixel intensities, and the resolution of the camera-lens combination is measured by making CTF measurements. First, the pixel intensities over a contiguous row of pixels are plotted as a function of the pixel indices. Here, I_{\max} and I_{\min} are taken as shown in Fig. 5, allowing the contrast to be calculated according to Eq. (3). For wider lines, the frequency ξ , can be computed as discussed in Eq. (2), but precise determination of the frequency is more difficult as the lines become only a few pixels wide. However, since the frequency of any target is $2^{1/6} \approx 1.12$ times the previous one, it is possible to precisely measure a wider target as in Fig. 5 and extrapolate the frequencies of smaller ones.

Repeating this procedure over many frequencies allows us to plot the CTF, as in Fig. 6. As a general rule of thumb, a system's resolution is considered poor if $C \leq 0.2$. Knowing the CTF can inform us as to what spatial frequency corresponds to poor resolution.

IV. ILLUMINATED SLIT (MODULATION TRANSFER FUNCTION)

While test targets offer a straightforward method for measuring the spatial resolution of an optical system, the measurements can be time-consuming and tedious. This begs the question: Is there a more efficient way to measure the spatial resolution of an optical system?

Fortunately, the answer to this question is yes. Another method, using the formalism of the "modulation transfer function," trades off simplicity and straightforward calculation for greater precision, speed, and lower sensitivity to noise. As in the test target method, a quantity called the modulation is defined. Like the contrast, the modulation is another way to quantify the resolution at a certain spatial frequency. Furthermore, in the same way that contrast measurements define the CTF, measuring the modulation as a function of frequency defines the modulation transfer function, or the MTF. To understand what the modulation is, consider a test target that contains not a set of solid black and white bars, but a sinusoidal intensity profile that sweeps from the left to the right of the field of view as a function of the azimuthal angle θ , as shown in Fig. 7. Suppose this sinusoidal intensity profile projects onto the camera's field of view a sinusoid of a certain angular frequency ξ with some positive amplitude A_{in} and background offset B_{in} , as described by

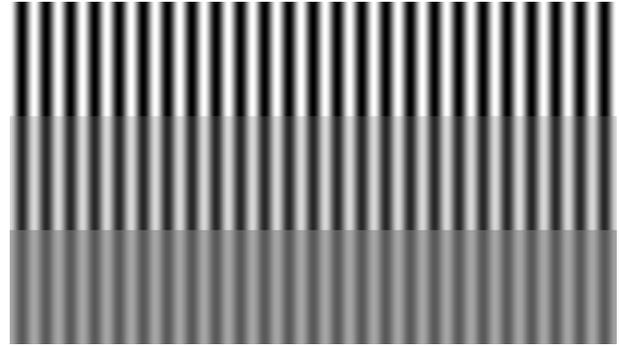


Fig. 7. Simulated responses to a sinusoidal target filling a camera's field of view are shown here. The 25 bright and 25 dark lines spanning the picture width means our frequency is $\xi = 50$ LPPW. The top third of the image is the response of an ideal camera, with $\text{MTF}(\xi = 50 \text{ LPPW}) = 1$. The middle third is the simulated response of a camera with $\text{MTF}(\xi = 50 \text{ LPPW}) = 0.70$, and the bottom third is a camera with $\text{MTF}(\xi = 50 \text{ LPPW}) = 0.30$. Among the many advantages of using the MTF is the fact that it is straightforward to simulate the response of cameras to arbitrary input signals.

$$\text{Input}(\theta) = A_{\text{in}} \sin(2\pi\xi\theta + \phi_{\text{in}}) + B_{\text{in}}. \quad (4)$$

To a very good approximation, the output intensity profile captured by the camera will also look like a sinusoid. However, the output sinusoid will likely have a different amplitude, offset, and phase (A_{out} , B_{out} , and ϕ_{out}), as given by

$$\text{Output}(\theta) = A_{\text{out}} \sin(2\pi\xi\theta + \phi_{\text{out}}) + B_{\text{out}}. \quad (5)$$

The modulation at that frequency ξ can be defined as $A_{\text{out}}(\xi)/A_{\text{in}}(\xi)$. Note that the modulation is always positive and does not specify any information about the phase; thus, even though the modulation turns out to be a very useful quantity, in this sense it is an incomplete description of the camera. If our camera is perfect, then the imaged profile will be exactly the same as the input, and we would have $A_{\text{out}}/A_{\text{in}} = 1$ for all frequencies ξ . If our camera is imperfect, oscillations will not be resolved fully, especially as ξ increases. Mathematically stated, we will always have $A_{\text{out}} < A_{\text{in}}$, and as ξ increases we will have $A_{\text{out}}/A_{\text{in}} \rightarrow 0$. In a manner very similar to the CTF, the MTF (modulation transfer function) is defined by making measurements at various angular spatial frequencies ξ , with the important restriction that $\text{MTF}(\xi = 0) \equiv 1$.

The choice of using a sinusoidal pattern to define the modulation may seem arbitrary. Why is the modulation defined in terms of sinusoidal inputs and outputs? After all, smoothly varying sinusoidal test targets are more difficult to manufacture than alternating black and white bars. The answer lies in Fourier synthesis—by considering a linear combination of sinusoids of different spatial frequencies, we can construct intensity profiles that are experimentally convenient to realize and whose frequency content is known. Taking a picture of such a target contains spatially superimposed information at each spatial frequency in the linear combination. This information can be separated by taking a discrete Fourier transform.

The advantages of this method are manifold. First, instead of isolating smaller and smaller line pairs for contrast measurements at only one spatial frequency, using the Fourier transform gives us data over many frequency components in a single measurement. Second, because a Fourier transform

takes into account the information distributed over the entire spatial domain, our measurement is robust against spatially localized defects such as dead pixels, dust on the detector, etc. This is very different from using test targets, where the information about the camera is localized in spatial frequency and in space to a small region of the test target that is easily susceptible to experimental imperfections (such as uneven illumination). Finally, knowing the MTF of a camera allows users to simulate output images for arbitrary input signals.¹⁰

A. Selecting an intensity profile

To measure the (one-dimensional) modulation transfer function, it remains to select an intensity profile to image and find out what its frequency components are as a function of spatial frequency. In principle, most signals will have spatial frequency content over a wide range of frequencies, but in the presence of noise, we want to maximize the signal-to-noise ratio at high frequencies. That is, we want $A_{\text{in}}(\xi)$ to drop off as slowly as possible as ξ increases. The ideal intensity profile would have spatial frequency content distributed evenly over all frequencies. Such an intensity profile exists: it is the well-known Dirac delta function, which can be modeled as

$$\text{Input}(\theta) = I_0 \frac{L}{D} \delta(\theta) = \begin{cases} \infty & \theta = 0 \\ 0 & \theta \neq 0, \end{cases} \quad (6)$$

where $I_0 L/D$ is a scaling constant describing the “strength” of the input signal. This profile is infinitely localized in space and contains information at all spatial frequencies. Thus, by imaging an intensity profile that is a Dirac delta function, or a close approximation to it, our output signal will also contain a wide range of spatial frequencies, and thereby allow us to sample over a wide frequency spectrum with a single image.

This idea can be represented more quantitatively by taking a continuous Fourier transform of the Dirac delta function to see how the total input intensity is distributed as a function of the frequency ξ . This calculation is analogous to measuring the input amplitude A_{in} as a function of ξ . Denoting this function $X_{\delta}(\xi)$, we have

$$X_{\delta}(\xi) = \int_{-\infty}^{\infty} \frac{I_0 L}{D} \delta(\theta) e^{-i\xi\theta} d\theta = \frac{I_0 L}{D} e^{-i\xi(0)} = I_0 \frac{L}{D}. \quad (7)$$

We can make sense of the fact that $X_{\delta}(\xi)$ is constant over all angular frequencies in the sense that the signal present in the Dirac delta function is uniformly distributed over all angular frequencies. Then I_0 can be interpreted as the power per 2π radians per unit angular frequency emitted by the source, and $I_0 L/D$ is the power coming through an arc that subtends an angle of L/D radians.

In practice, arbitrarily thin and bright intensity profiles are impossible to realize. To circumvent this problem, we will consider a thin slit of finite width, which is a rectangular profile of small, finite width L illuminated by a bright light of finite power I_0 per unit angular frequency per 2π radians. This profile can be modeled by

$$\text{Input}(\theta) = I_0 \left[\mathcal{H}\left(\theta - \frac{L}{2D}\right) - \mathcal{H}\left(\theta + \frac{L}{2D}\right) \right], \quad (8)$$

where \mathcal{H} is the unit step function. The spatial frequency content theoretically present in the input signal is given by taking the Fourier transform¹¹ of $\text{Input}(\theta)$. We compute

$$\begin{aligned} X_L(\xi) &= \mathcal{F}\{\text{Input}(\theta)\} = \int_{-L/2D}^{L/2D} I_0 e^{-i\xi\theta} d\theta \\ &= \frac{I_0 L}{D} \text{sinc}\left(\frac{\xi L}{2D}\right), \end{aligned} \quad (9)$$

and observe that the input signal vanishes at high frequencies ($\xi \rightarrow \infty$), a direct consequence of $\text{Input}(\theta)$ having finite width.

$\text{Output}(\theta)$ is then measured by taking a picture of $\text{Input}(\theta)$. Taking the discrete Fourier transform of $\text{Output}(\theta)$ gives $Y_L(\xi)$, the spatial content present in the output signal; this is analogous to finding the output amplitude A_{out} at all frequencies simultaneously. A more sophisticated definition of the modulation transfer function, therefore, is the Fourier transform of the output signal divided by its corresponding input signal, with absolute value bars inserted to keep the modulation positive

$$\text{MTF}(\xi) = \left| \frac{Y(\xi)}{X(\xi)} \right|. \quad (10)$$

B. Diffraction theory and the MTF

We can revisit some of the results from diffraction theory using the MTF formalism. Since diffraction through apertures smears point sources (Dirac delta functions), we can think of apertures as filters that remove some of the high spatial frequency information from the system. This property is exploited, for example, when a laser’s spatial profile is smoothed by spatial filtering.¹ In particular, we can derive a limit on the MTF from the Rayleigh criterion. We estimate the spatial frequency and the modulation corresponding to the just resolvable condition by arranging many identical Airy disks next to each other in a periodic fashion such that they are just resolvable, as shown in Fig. 8. The separation distance defines the corresponding spatial frequency. The modulation corresponding to this cutoff frequency can be estimated by dividing the peak-to-peak amplitude of the

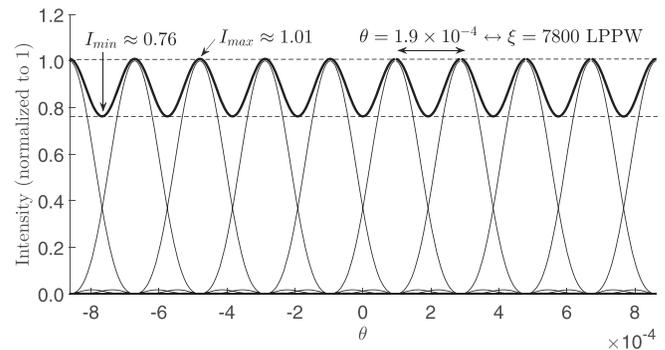


Fig. 8. A sequence of point sources are placed such that they are just resolvable as defined by the Rayleigh Criterion. The angular spacing of the point sources and the camera’s angular field of view define a particular angular spatial frequency in LPPW. The summed intensities, plotted in bold, can be used to estimate the maximum contrast/modulation allowable by diffraction theory at Rayleigh’s cutoff wavelength, which turns out to be ~ 7800 LPPW.

summed intensities by the height of each individual Airy disk.

To estimate the Airy disk size, we write down some typical parameters used in our setup. We fix the aperture size of the Canon 40D to $2a = 3.5$ mm. We assume the light is green on average, with $\lambda = 550$ nm, and take the object to be at a distance $D = 70$ cm away. The angular separation implied by the Rayleigh criterion corresponds to a spatial frequency of ~ 7800 LPPW. This can be estimated using the angle subtended by the field of view, which is roughly 0.76 rad from left to right at a focal length $f = 28$ mm, according to online sources including the Canon 40D manual.¹² As shown in Fig. 8, the contrast at that spatial frequency can be estimated to be $C = 0.14$. As we will discuss later, the contrast is a good approximation to the modulation in the high-frequency limit, so we write $\text{MTF}(7800 \text{ LPPW}) \sim 0.14$.

C. Image acquisition

To measure the modulation transfer function, it is necessary to precisely realize the illuminated slit of width L as described in Eq. (8), and then to acquire an image. To precisely realize a slit of width L , a pair of razor blades placed on an adjustable micrometer drive and a black cardboard screen together obscure all but a thin slit. The razor blades are aligned by eye to be parallel and oriented such that the sharp edges face each other. Assuming the lens is symmetric about the lens axis, rotating the slit about the lens axis simply rotates the image formed on the sensor, providing no new information about the optics. However, choosing to orient the slit along one of the axes of the pixel array maximizes the Nyquist frequency attained without resorting to sophisticated supersampling methods.¹³ An incandescent lamp is placed about 2 m behind the razor blades to provide collimated, spatially uniform light at the slit. The camera is placed about half a meter away from the razor blades, and mounted to the breadboard to ensure repeatability. A schematic of the entire setup can be seen in Fig. 9, and a closeup of our slit is shown in Fig. 10.

To compare the results from the illuminated slit with the results from the test targets, camera parameters that significantly affect the camera's resolution must be held constant from before. The intensity profile of interest, in this case the illuminated slit, should be in the center of the field of view, just as the test targets were. Measurements should be made at similar distances so the camera's focus is similar. The

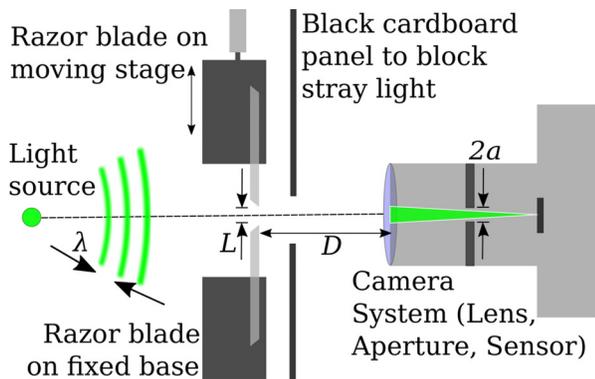


Fig. 9. A schematic of the setup used to realize and image a uniformly illuminated slit is shown. An idealization of the camera optics and aperture are shown, with the camera's CMOS sensor represented by the dark bar at the far right.

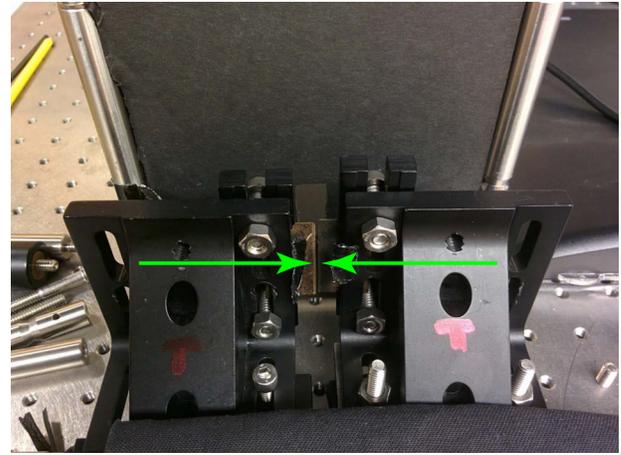


Fig. 10. The slit setup as seen from the source side (i.e., the left side of Fig. 9). The arrows (green online) indicate the razor blades that define the slit. The blade to the left is mounted on a micrometer drive while the right blade remains fixed. The black cardboard and the pieces of tape on the razor blade block stray light from the source.

camera's aperture size should be held constant. We use $2a = 3.5$ mm, which corresponds to $f/8.0$, to begin with. The ISO and shutter speed are selected as discussed previously. After the camera is aligned, focused, and set to the desired settings, an image of the illuminated slit is taken. As described earlier, a JPG image consisting of 2592×3888 pixels is produced, which can be treated as a $2592 \times 3888 \times 3$ matrix, where the three layers represent the intensity in each of the red, green, and blue color bands.

Because modulation measurements are faster to make than contrast measurements, it is straightforward to quickly measure the MTF at different aperture diameters to characterize the camera. In addition to measuring the MTF at $2a = 3.5$ mm, we open up the aperture to $2a = 8.6$ mm corresponding to $f/3.5$. Measuring the MTF at two different aperture sizes allows us to observe the opposing effects of lens nonideality and the diffraction limit on the camera's resolution. As the lens is opened up, the diffraction limit increases. At the same time, spherical aberration is increased because the curvature of the lens is no longer negligible as the aperture is opened up. In addition, by separating the image into its constituent red/green/blue channels, we may investigate the MTF as a function of color band.

D. Data analysis

This section describes how the JPG image of a slit of width L is used to determine the MTF. The input and output spectra, $X_L(\xi)$ and $Y_L(\xi)$, can both be fully determined from the image.

Given a JPG image, a small rectangular region of interest (ROI) consisting of $(M \times N) \times 3$ pixels is selected. Cropping the image only affects the frequency resolution, leaving the upper limit of the angular frequency (or the Nyquist frequency), unchanged. Increasing the number of points in the ROI is a tradeoff between increasing the frequency resolution and increasing the noise, since the signal is localized near the bright slit and pixel readings sufficiently far away from the bright slit contain less information about the camera's response. We experimented with different ROI sizes and found that $M \times N = 79 \times 779$ gave good frequency resolution and low noise for a variety of different slit widths, as



Fig. 11. The region of interest shown is cropped from a larger JPG image of the slit. The dimensions of the region of interest are chosen to strike a balance between frequency resolution and increased noise. Discrete Fourier transforms are taken along the horizontal axis to obtain a number of different output spectrum estimates. This particular image was taken at a slit width of 1.06 mm.

shown in Fig. 11. The color bands can either be analyzed separately or aggregated with a weighted average

$$I_{\text{grayscale}} = c_r I_{\text{red}} + c_g I_{\text{green}} + c_b I_{\text{blue}}, \quad (11)$$

where (c_r, c_g, c_b) are weights whose exact values depend on different conversions to grayscale taking into account the different energies of photons in the different bands, the response of the human eye, the efficiency of the detectors, etc. By default, MATLAB has a built in function that implements this conversion, taking $(c_r, c_g, c_b) = (0.2989, 0.5870, 0.1140)$.¹⁴ Regardless of whether color is separated or combined into a single intensity value according to Eq. (11), we can now discuss how to obtain $X_L(\xi)$ and $Y_L(\xi)$ from a matrix of pixel intensities with M rows and N columns in which the same spatial information is encoded in each of the M rows, corresponding to Fig. 11.

In the case of the input spectrum, the functional form of $X_L(\xi)$ is known from Eq. (9). Since the zeros of $X_L(\xi)$ will show up as dips in $Y_L(\xi)$, we can fit the first zero of the functional form of $X_L(\xi)$ experimentally, as shown in Fig. 12. The intensity scale of $X_L(\xi)$ is set by $X_L(0) = Y_L(0)$, since $\text{MTF}(0) = 1$ by definition. The remainder of this section will focus on how to obtain smooth estimates of $Y_L(\xi)$.

Taking the discrete Fourier transform of each of M rows in the ROI yields M estimates of the frequency content of the output signal $Y_L(\xi)$. These estimates, labeled $Y_{1,L}(\xi), \dots, Y_{M,L}(\xi)$, are in general complex-valued to reflect the phase offset of each frequency component. However, since the MTF does not encode phase information and is a real, positive quantity, we ignore the phase factor and average the magnitudes over the M estimates to provide a smoother estimate of the real, positive output spectrum $Y_L(\xi)$, or

$$Y_L(\xi) = \frac{1}{M} \sum_{i=1}^{M \text{ rows}} |Y_{i,L}(\xi)|. \quad (12)$$

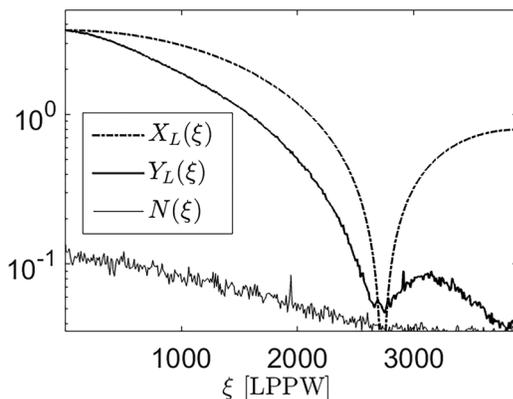


Fig. 12. Assuming an input spectrum of the form $X(\xi) = (I_0 L/D) |\text{sinc}(\xi L/2D)|$, we fit the first zero of the input spectrum by estimating the minimum of the experimental data.

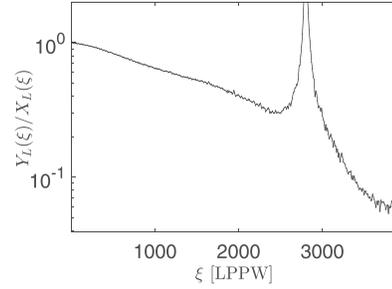


Fig. 13. Plotting $Y_L(\xi)/X_L(\xi)$ from the data shown in Fig. 12 reveals a spurious peak in the MTF estimate, a result of the fact that the analytically defined input spectrum $X_L(\xi)$ vanishes for certain values of ξ while the experimentally measured output spectrum $Y_L(\xi)$ contains noise.

The same procedure repeated with the lamp turned off provides an estimate of the (real and positive) noise floor due to electronic cross-talk, stray light, etc. To a very good approximation, this noise floor $N(\xi)$ is independent of the slit width, so we do not denote it with a subscript L . The three spectra $X_L(\xi), Y_L(\xi), N(\xi)$ are plotted in Fig. 12. The resulting MTF estimate $\text{MTF}_L(\xi) = Y_L(\xi)/X_L(\xi)$ is plotted in Fig. 13.

The obvious problem with Fig. 13 is that it features a spurious peak which is clearly unphysical in nature. This behavior is a result of our measurements having $Y_L(\xi) > 0$ everywhere due to the finite noise floor, whereas $X(\xi) = I_0 L/D |\text{sinc}(\xi L/2D)|$ has zeros at $\xi = 2n\pi D/L, n = 1, 2, \dots$, one of which is seen in Fig. 12. One solution is to use values of L that are so narrow that the finite pixels cannot distinguish the slit from an infinitely thin impulse. This approach runs into difficulties because as the slit gets narrower, there is not enough light incident on the detector to maintain a good signal-to-noise ratio, as discussed in Eq. (9). On the other hand, as the slit gets wider, the spurious peaks in the MTF estimates are shifted to lower and lower frequencies.

A robust solution is to take a weighted average of multiple MTF estimates over several slit widths chosen to be not too narrow and not too wide. In our experiment, six such widths ranging from 0.21 to 1.06 mm are used and are listed in Table I. The “weights” in the weighted average are taken to be the relative signal-to-noise ratio at each frequency. For a particular slit width L , we can define the signal-to-noise ratio as a function $\text{SNR}_L(\xi)$ of the angular frequency ξ as

$$\text{SNR}_L(\xi) = \frac{X_L(\xi)}{N(\xi)}, \quad (13)$$

Table I. Angular spatial frequencies ξ for different slit widths. The slit widths were measured to a precision of $1 \mu\text{m}$ with a micrometer drive on the translation stage. The first few frequency values were measured experimentally, and since the first zero of $X_L(\xi)$ scales as $1/L$, the higher values are extrapolated to some extent (the extrapolated values are indicated with an asterisk).

Slit width [mm]	ξ of first zero [LPPW]
1.060	1013
0.860	1273
0.660	1812
0.460	2670
0.260	4480*
0.210	5546*

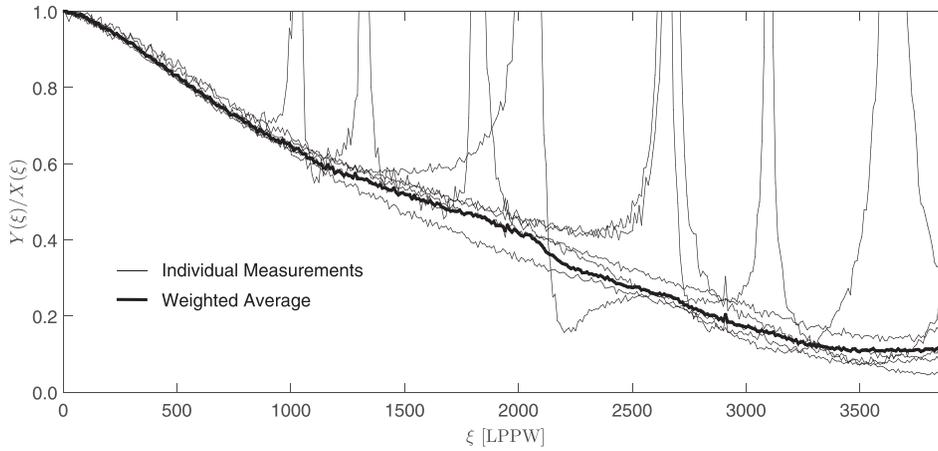


Fig. 14. A collection of MTF estimates made at six different slit widths as listed in Table I. The weighted average of the six estimates does not feature any of the spurious features present in each individual measurement.

with the relative signal-to-noise ratio defined as

$$\text{RSNR}_{L_i}(\xi) = \frac{\text{SNR}_{L_i}(\xi)}{\sum_{\text{all } L} \text{SNR}_L(\xi)}. \quad (14)$$

As we would like the weight $\text{RSNR}_L(\xi)$ on a particular measurement $Y_L(\xi)$ vanishes at the frequencies where the input signal $X_L(\xi)$ vanishes, the sum of the RSNR weights is 1. Hence, the spurious peaks, such as the one seen in Fig. 13, do not end up affecting the MTF estimate significantly.

We form a weighted average of MTF estimates over the values of L tabulated in Table I, which turns out to be equivalent to measuring the sum of the output signals divided by the sum of the input signals. The individual MTF estimates at six different slit widths and their weighted average are

$$\text{MTF}(\xi) = \sum_{\text{all } L} \text{RSNR}_L(\xi) \text{MTF}_L(\xi) = \frac{\sum_{\text{all } L} Y_L(\xi)}{\sum_{\text{all } L} X_L(\xi)}, \quad (15)$$

and are plotted (for grayscale) in Fig. 14. Meanwhile, Fig. 15 shows the results when the color bands are analyzed separately, for both $f/3.5$ and $f/8.0$.

V. RESULTS

In summary, we have measured the grayscale CTF of the Canon 40D at $f/8.0$, and the MTF of the camera at $f/3.5$ and $f/8.0$ in each color band. Because the measurements were performed under similar conditions, we can now answer two interesting questions. First, what is the relationship between the MTF and the CTF? Second, what are the limiting factors on our camera's resolution?

A. Comparison between MTF and CTF

The first of these questions can be answered mathematically and confirmed experimentally. Given a set of contrast measurements, there exists a series expansion

$$\text{MTF}(\xi) = \frac{\pi}{4} \left[\text{CTF}(\xi) + \frac{\text{CTF}(3\xi)}{3} - \dots \right], \quad (16)$$

which gives the MTF in terms of the corresponding CTF.¹⁵ By converting the CTF measurements made with the test targets into MTF values using the approximation in Eq. (16), we can compare the results from the test target method against the Fourier analysis method, keeping in mind that the approximation we make underestimates the MTF at low

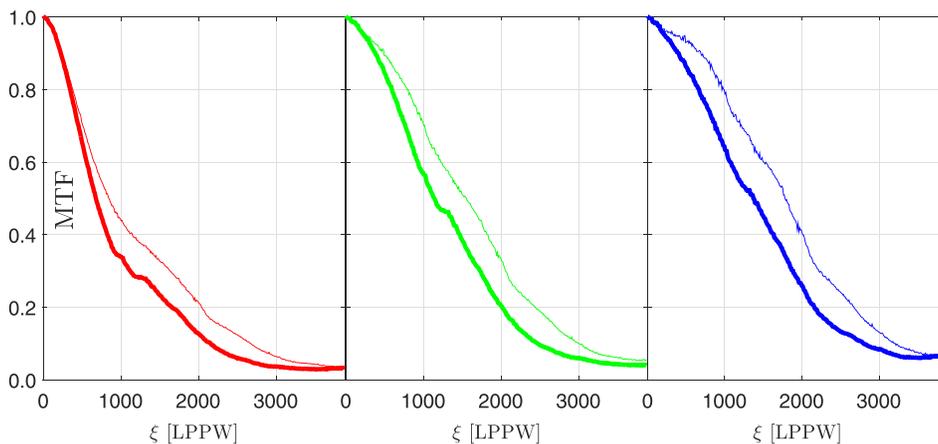


Fig. 15. The MTF of a Canon 40D is measured as a function of color band (from left to right: red, green, and blue band) and as a function of aperture size ($f/3.5$ in thick traces and $f/8.0$ in thin traces). Note that shortening the light's wavelength improves resolution, but increasing the aperture size does not. This shows that the camera is not diffraction-limited in this regime.

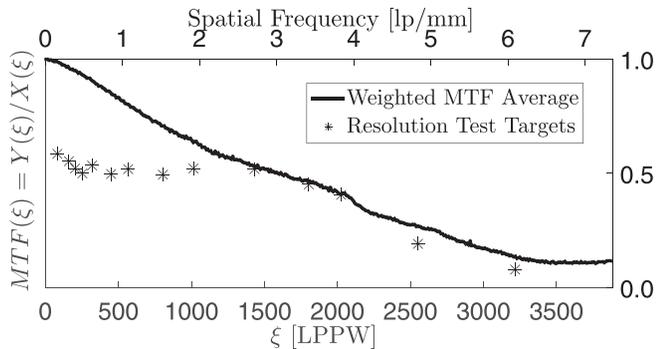


Fig. 16. Contrast measurements made from the test targets at $f/8.0$ shown in Fig. 4 are converted to modulation using Eq. (16). These values are compared to our MTF measurements made under similar conditions with the illuminated slit. As expected, the test-target measurement underestimates the modulation at low frequencies, and the two measurements converge as ξ increases.

angular frequencies. Interestingly, since both the MTF and CTF tend toward zero in the high-frequency limit, the higher-order terms in Eq. (16) vanish faster than the $\text{CTF}(\xi)$ term. That is, the CTF converges to the MTF up to a constant as ξ gets large, so that

$$\text{MTF}(\xi) \approx \frac{\pi}{4} \text{CTF}(\xi). \quad (17)$$

Experimentally, this agreement is seen at high frequencies as expected (see Fig. 16).

The tradeoffs of using contrast versus modulation to quantify a system's optical quality are clear from Fig. 16. On the one hand, measuring the contrast is much more intuitive and straightforward than measuring the modulation and does not require the same amount of averaging to obtain meaningful results. On the other hand, measuring the contrast at many points is much more tedious. Naturally, choosing between one quantification method or another depends on what information is relevant. If the falloff rate of the spatial resolution as a function of frequency is not of interest, it is more practical to specify a contrast threshold and measure the frequency corresponding to that threshold. In many situations, however, the MTF is more valuable to have. For example, astronomers routinely measure the MTF of telescopes and imaging cameras in order to estimate the size of astronomical objects from telescope images.¹⁶

B. Limits on resolution

We know that for even the most ideal camera, diffraction is the ultimate downfall of resolution. But how close to ideal is our Canon 40D? In this section, we discuss possible limiting factors on the resolution and conclude that spherical aberration, rather than pixelation, chromatic effects, or diffraction, is the most likely culprit.

Recall that we derived theoretically that $\text{MTF}(7800 \text{ LPPW}) \sim 0.14$ corresponds to an angular separation implied by the Rayleigh criterion. Since this spatial frequency corresponds to an angular size smaller than a pixel (whose angular size corresponds to 3888 LPPW), we determine that our consumer-grade imaging system is not operating near the Rayleigh limit.

Pixelation is another concern. The fact that pixelation causes the image of a point source to instead look like a

square the size of a pixel allows us to infer the functional form of the MTF contribution due to pixelation. If pixelation were the limiting factor on the resolution, the resolution should fall off like $\text{sinc}(\xi/3888 \text{ LPPW})$. However, as seen in Fig. 15, the resolution falls off far more rapidly and has a strong dependence on the aperture size, leading us to rule out this hypothesis.

From the same set of data, Fig. 15, we can rule out chromatic, or color-dependent, aberration as the dominant contribution to our resolution. The data show that the qualitative wavelength dependence is as predicted by scalar diffraction theory for both large and small aperture sizes—the resolution improves as the wavelength is decreased. We conclude that chromatic aberration is not the limiting factor on the Canon 40D's resolution.

Finally, we consider spherical aberration, an effect that occurs in imaging systems where the thin-lens approximation breaks down. This effect becomes more pronounced as the aperture is opened wider.¹ Our data show that when the aperture size is increased to $2a = 8.6 \text{ mm}$, the resolution decreases. We conclude that spherical aberration is likely the limiting factor on the resolution.

It is plausible that the camera's designers have carefully optimized the resolution limits of all of the components in the body of the camera—the pixel array, electronics, etc.—to minimize the cost of the camera. At the same time, photography connoisseurs have the option to improve the system's performance by replacing the stock lens with a specialized lens to suit their imaging needs. In this way, the Canon 40D can cater to a wide range of consumers at minimal base cost. This sensible design choice is a direct reflection of the rich science and engineering that goes into making a camera.

VI. CONCLUSION

We have demonstrated a number of ways to quantify the resolution of a camera with a relatively simple experimental setup, including using the Rayleigh Criterion, the contrast transfer function (CTF), and finally the modulation transfer function (MTF). Our measurements suggest that the resolution of a typical DSLR such as the Canon 40D is limited by its optical subsystem, namely, spherical aberration. This experiment provides a concrete introduction to cornerstone concepts of modern optics, including Fourier analysis, diffraction theory, and analysis techniques to resolve information in spatial data. The emphasis placed on overcoming practical problems dealt with in the MTF measurement (i.e., making measurements that are robust against noise) makes the experiment described well-suited for the undergraduate laboratory and for students going on in the optics/photonics industry.³

ACKNOWLEDGMENTS

The authors thank Professors Richard Haskell, Peter Saeta, and Patricia Sparks for helpful discussions. This investigation would not have been possible without their support.

APPENDIX A: SUGGESTED BILL OF MATERIALS

Included in Table II is a suggested bill of materials to use in building this experiment. We did not have an adjustable mechanical slit and built our own without much difficulty out of translation stages, right angle brackets, and two

Table II. A suggested bill of materials to be used in building this experiment. Standard hardware like an optical breadboard and bolts are not included.

Item	Cost
Digital camera, e.g., Canon EOS 40D (used)	\$300
Lens, e.g., Canon EF 28–135 mm <i>f</i> /3.5–5.6 IS USM	\$300
Camera remote trigger	\$10
Mini-USB transfer cable	\$5
Thorlabs adjustable mechanical slit VA100	\$248
1951 Air force resolution test targets	\$10
Table lamp	\$10
Black cardboard for stray light blockage	\$5
Total	\$890

commercially available razor blades. Though our setup was mounted to a breadboard, mechanical vibrations were not found to be the limiting factor in our ability to make robust measurements.

¹M. Born and E. Wolf, *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light* (Cambridge U. P., Liberty Plaza Floor 20, New York NY 10006, 2000), pp. 370–433.

²G. D. Boreman, *Modulation Transfer Function in Optical and Electro-Optical Systems* (SPIE Press, Bellingham, 2000), pp. 1–61.

³D. C. O’Shea, “Transition to reality: Preparing the student at Georgia Tech through an undergraduate sequence of optics laboratories,” *Educ. Opt.* **1603**, 317–324 (1992).

⁴A. P. Tzannes and J. M. Monney, “Measurement of the modulation transfer function of infrared cameras,” *Opt. Eng.* **34**, 1808–1817 (1995).

⁵H. Fujita, D.-Y. Tsai, T. Itoh, J. Morishita, K. Ueda, and A. Ohtsuka, “A simple method for determining the modulation transfer function in digital radiography,” *IEEE Trans. Med. Imaging* **11**, 34–39 (1992).

⁶D. C. O’Shea, “A modulation transfer function analyzer based on a micro-computer and dynamic ram chip camera,” *Am. J. Phys.* **54**, 821–854 (1986).

⁷G. B. Airy, “On the diffraction of an object-glass with circular aperture,” *Cambridge Philos. Soc. Trans.* **5**, 283 (1835).

⁸J. Strutt, “On the theory of optical images, with special reference to the microscope,” *Philos. Mag.* **42**, 167–195 (1896).

⁹In this study, we used <http://digiCamControl.com/digiCamControl>, which is open source, available at no charge, and is compatible with a wide range of cameras.

¹⁰Of course, to handle two dimensional images, one needs to generalize the definition of the MTF to two dimensions.

¹¹Readers familiar with scalar diffraction theory will notice the similarity between the functional form of $[X_L(\xi)]^2$ and the Fraunhofer diffraction pattern projected onto a screen by a coherently illuminated rectangular slit.¹ The connection between Fourier transforms and diffraction is profound but outside the scope of this paper.

¹²*EOS 40D Instruction Manual* (Canon Inc., Lake Success, NY 11042, USA, 2007), pp. 181–188.

¹³P. Burns, “Slanted-edge MTF for digital camera and scanner analysis,” in *Proceedings of the IS&T’s PICS Conference* (2000), pp. 135–138.

¹⁴L. Benedetti, M. Corsini, P. Cignoni, M. Callieri, and R. Scopigno, “Color to gray conversions in the context of stereo matching algorithms,” *Mach. Vision Appl.* **23**, 327–348 (2012).

¹⁵J. W. Coltman, “The specification of imaging properties by response to a sine wave input,” *J. Opt. Soc. Am.* **44**, 468–471 (1954).

¹⁶R. Lupton, J. E. Gunn, Z. Ivezic, G. R. Knapp, S. Kent, and N. Yasuda, “The SDSS imaging pipelines,” *ASP Conf. Ser.* **10**, 269–278 (2001).

¹⁷U.S. Department of the Air Force, *Military Standard Photograph Lenses* (U.S. Government Publishing Office, Washington DC., 1959), p. 21.



Fire Syringe

From the 1856 catalogue of Benjamin Pike, Jr. of New York City: “This instrument is used for producing instantaneous light, by means of the condensation of air. It consists of a stout condensing syringe, having a solid piston with a hole, for containing a small piece of tinder, at the end. Upon forcing down the piston quickly, the air within the tube is rapidly condensed, and the condensation occasions so great an evolution of heat as to light the tinder.” The deluxe model cost \$6.00. This example is in the collection of the National Museum of American History at the Smithsonian Institution in Washington, D.C. (Picture and Notes by Thomas B. Greenslade, Jr., Kenyon College)