Ph 77 - Advanced Physics Laboratory

Department of Physics, California Institute of Technology

- Optics Track -

Laboratory Optics & Interferometry

Overarching Objectives

The *Optics Track* in Ph77 is an eight-week set of laboratory and prelab exercises focusing on optics and optical instrumentation. Although we cannot include every hardware specialty in this short time, we try to cover some especially common types of optical equipment together with a foundational description of measurement techniques and mathematical principles. One of our primary goals is for you to gain familiarity with the kinds of optical hardware you might soon encounter in a university research lab in graduate school. However, given how ubiquitous optics has become in modern science and technology, you may well find this knowledge quite useful if your career path involves using or designing a large variety of laboratory experiments, lasers, spacecraft, telescope instrumentation, optical telecom technology, commercial products, or any other equipment that incorporates optics.

The supplementary documents *General Intro.pdf* and *eNotebook Example.pdf* (available on Canvas) contain important information that is common to all the Ph77 Tracks. Please read both these short documents before proceeding. If you are serious about learning optics at a deeper level, we recommend that you purchase the textbook *Optics*, by Eugene Hecht, and look through it independently while working on this Track.

Part 1. Laboratory Optics & Imaging

Because Ph77 is a laboratory course, our primary emphasis is on practical optical devices that are in widespread use. To this end, we begin with a primer on camera lenses and imaging sensors that you will use throughout this Track. You carry additional examples around with you every day, including your eyes and smartphone, and the basic tenets of optical physics readily extend to telescope and microscope optics, and to radiation outside the visible range. If you want to learn about optics in science and technology, basic laboratory optics is a good place to start.

Lenses

A basic lens is made from a single piece of glass, known as a **singlet lens**, and Figure 1 shows several different types of singlets. Often two or three singlet lenses are glued together to form **doublets** or **triplets**, and often many lens elements are assembled to form complex lenses like the example shown in Figure 2.

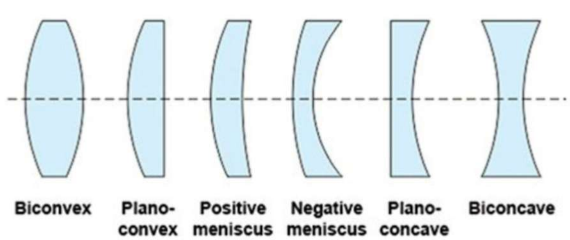

Figure 1. A "singlet" lens is a single piece of glass fabricated into a variety of shapes, such as those shown at left. These optical elements usually have spherical surfaces (as these are easy to manufacture), but aspherical surfaces are often used in high-performance applications.

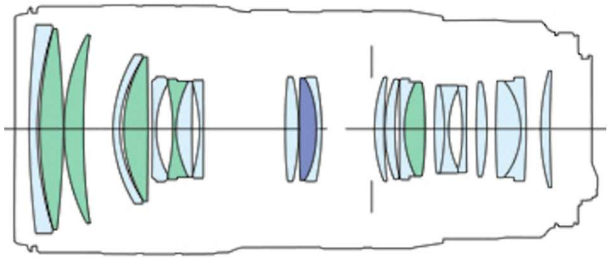

Figure 2. Complex camera lenses like this one are often used for imaging applications. The various lens elements move along the optical axis to change the focus and magnification, and an adjustable iris changes the lens aperture. These lenses incorporate many different types of glasses to reduce spherical and chromatic aberration, and anti-reflection (AR) coatings on the surfaces minimize reflections and light loss.

Both simple and complex lenses are characterized by several parameters:

Focal length. For a thin lens (see below), this parameter is the distance between the lens and the focal point of a set of on-axis parallel rays entering the lens. For a complex lens, the effective focal length can be defined from the magnification of far-away objects.

Aperture. For a thin lens, this is essentially the diameter of the lens (or, more accurately, the diameter of the central part of the lens that light passes through). With complex lenses, the effective aperture can often be changed using an adjustable *iris diaphragm* inside the lens body.

f-number. For a thin lens, this is equal to f/D, where D is the effective aperture and f is the effective focal length, and this ratio is nearly always greater than unity. On commercial camera lenses, the minimum f-number is usually printed on the side of the lens. The f-number is essentially an indicator of the angle of the unobstructed cone of light reaching the sensor. Note that the notation can be confusing, as the f-number is not the same as the focal length f. The f-number is sometimes abbreviated f#, and an f-number of three is often written "f/3", even though this does not mean anything is divided by three.

With a simple lens, these parameters are all fixed. With complex lenses, they are often adjustable, either manually using adjustment rings on the lens body or electronically using actuators inside the lens.

Thin-Lens Approximation

It is often useful to work in the *thin-lens approximation,* as shown by the ray diagram in Figure 3. In this diagram, simple geometry is sufficient to determine the size, position, and orientation of an image formed by a lens in the thin-lens approximation. You can look up a variety of thin-lens formulas

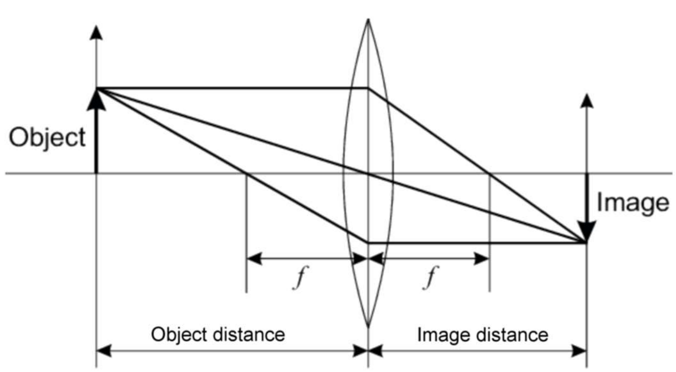

*Figure 3. In the **thin-lens approximation**, any two of the three light rays shown here are sufficient to determine the size and position of the image formed by a known object. Note that parts of the top and bottom rays run parallel to the optical axis (as shown), while the middle ray passes through the center of the lens with no deflection.*

online, but they can all be readily derived from the ray diagram. There are basically three rules for making a thin-lens sketch like that in Figure 3:

1) Starting from the tip of the object, draw one ray parallel to the optical axis. This ray goes through the focal point on the other side of the lens. In the thin-lens approximation, the focal points on the two sides of the lens are equidistant from the center of the lens.
2) Draw another ray (also starting at the tip of the object) that goes through the focal point on the near side of the lens. This ray comes out parallel to the optical axis on the other side of the lens.
3) Draw another ray through the center of the lens. In the thin-lens approximation, this ray goes straight through the lens.
4) These three rays all converge on the tip of the image, as seen in Figure 3. Any two of these rays are sufficient to define the position of the tip of the arrow. The geometry of the final ray diagram can be used to derive formulas for the position and height of the image arrow.

The thin-lens approximation is often quite accurate for singlet/doublet/triplet lenses provided the focal length is long compared with the lens thickness and aperture. Complex lenses like the one shown in Figure 2 are not well represented by the thin-lens approximation.

In the thin-lens approximation, the image magnification is defined simply as M = (image size)/(object size), and this quantity is easily calculated from the ray diagram in Figure 3 (and formulas can be found online). Note that M in this context is not an intrinsic property of the lens, as it depends on the position of the object.

Note also that the magnification factor M is not the same as the zoom factor X you often see in common usage, as in "my smartphone has a 3X optical zoom lens". When a manufacturer says 1X, that usually means that a photo will look something like what you see with your naked eye, and a 3X lens gives a photo that is zoomed in a factor of three relative to 1X. In this sense, the X factor mostly tells you about the angular field-of-view of the final image, not the properties of the lens. In common usage, both the X factor and the word "magnification" are often ill-defined to the point of being meaningless or misleading.

Exercise 1. If the image distance equals the object distance in the thin-lens approximation, what is that distance? Sketch the thin-lens diagram for this case. This configuration is often called 1:1 imaging, because the image size is equal to the object size.

The Diffraction Limit

When you pass visible light through a lens, the finite size of the lens aperture will cause the light to diffract. Of course, you can find the fundamental physics of diffraction described online (or in optics textbooks) in as much detail as you want to know. In a nutshell, sending a plane wave through an aperture with diameter D will give a light beam that expands with an angle

$$\delta\theta_{diff} \approx \frac{\lambda}{D} \tag{1}$$

where λ is the wavelength of light. For an object placed far from the lens (at a distance $L \gg f$), the smallest feature on that object that can be resolved with this lens is given by

$$\delta x_{far-away-object} \approx \delta\theta_{diff} \cdot L \tag{2}$$

This *diffraction limit* is a fundamental barrier to optical imaging resolution, typically expressed as an angle in radians. In a microscope, the object is usually placed at the focal length, and then the diffraction limit tells us that the smallest resolvable features have

$$\begin{aligned} \delta x_{microscope} &\approx \delta\theta_{diff} \cdot f \\ &\approx \lambda \cdot f\# \end{aligned} \tag{3}$$

Because it is (essentially) impossible to make a lens with an f# less than unity, it is therefore (essentially) impossible to resolve objects smaller than λ with an optical microscope. Nobel prizes have been awarded for developing clever tricks that beat the diffraction limit, and this remains an area of active research.

Exercise 2. The aperture size for a smartphone camera is limited by what people are willing to carry in their pockets, so consider an aperture of D = 5 mm. If you take a picture of the Moon using such a camera, what is $N = \delta\theta_{Moon}/\delta\theta_{diff}$, assuming a wavelength of 550 nm? (It is customary to use 550 nm when talking about visible light, as this is in the green part of the spectrum.) This number gives you how many "real" pixels (where the pixel size is equal to the resolution limit) you get across the diameter of a full Moon in your photo.

Note that this fundamental resolution limit is independent of other properties of the imaging system, including the sensor properties, the lens focal length, etc. [Software post-processing can improve the effective resolution by maybe a factor of two, at least in general appearance, mostly using sharpening algorithms. And you can do slightly better still by taking multiple photos and carefully using information from all of them. These kinds of computationally intensive image-enhancing techniques are another area of active research, although software tricks cannot take you much beyond the nominal diffraction limit.]

If you wanted to capture a high-resolution photo (N = 1000) of someone's head using this camera, about how close would they have to stand?

Lens Aberrations

Another important bit of practical information to remember is that all lenses suffer from *aberrations* that limit their performance. The presence of aberrations often means that a lens will not focus parallel rays to a precise point (even in the absence of diffraction) and this fact reduces the image resolution that a lens can achieve. Put another way, lens aberrations often prevent a lens from reaching the diffraction limit. The most common lens aberrations are:

Spherical aberration. This occurs because a lens made from spherical glass surfaces will not focus parallel light to a single point. Spherical aberration is generally small if $\sin(\theta) \approx \tan(\theta) \approx \theta$, where θ is the cone angle of the light hitting the sensor. In lens-speak, spherical aberration is smaller when the f-number is larger. As a general rule-of-thumb, spherical aberration is probably not a big issue if the f-number is greater than around 10. Aspherical surfaces can be used to reduce spherical aberrations to some degree, so these are often used in high-end applications.

Curved mirrors can be used for imaging (e.g., optical telescopes), and they also suffer from spherical aberration. A parabolic mirror can focus parallel on-axis rays to a point (ignoring diffraction), but a spherical mirror will not. On the other hand, parabolic mirrors have terrible off-axis imaging properties, so there are many trade-offs in optical telescope design.

Chromatic aberration. Because the index of refraction of glass depends on wavelength, a lens will have different focal lengths at different wavelengths. If you are using monochromatic light (typically a laser), then chromatic aberration is usually negligible. Chromatic aberration is also essentially negligible with telescope mirrors, as reflection angles do not change with wavelength.

If you are taking a photo with your smartphone, then you need not worry much about lens aberrations; the design engineers did their best to minimize these problems, and all you can do is point and shoot. But if you are designing a new physics experiment, astronomical instrument, or commercial product, then you may need to worry about aberrations as you optimize your design. There are sophisticated software tools that can help you with this task (see below) but designing an imaging system from scratch can be a challenging problem. For basic laboratory science, here are some general rules regarding lens aberrations:

1) On-axis aberrations are typically smaller than off-axis aberrations, sometimes much smaller. As a result, the quality of an image formed by a lens may be quite sharp in the center of the field of view but worse near the edges. All imaging systems have this problem, and it might be an important issue if you need a sharp focus over a wide field of view.
2) Aberrations are usually smaller when the f-number is larger, requiring some compromise from the user. A low f-number means a large aperture, which collects a lot of light and has a better diffraction limit; but a wide aperture means aberrations will be worse. On the other hand, a high f-number will give reduced aberrations, but the smaller aperture will admit less light, and the diffraction limit may limit resolution.
3) Complex lens designs can be used to reduce both spherical and chromatic aberrations, but such lenses are expensive. Also, scattered light from all the lens surfaces may introduce more problems. Fortunately, most commercial camera lenses designed for photography are exceptionally well designed and constructed, so these are often used in scientific applications as well.
4) In laser applications, one can often get away with basic singlet lenses. First, there is no chromatic aberration with monochromatic laser light. Second, a 2mm-diameter laser beam going into a 100mm focal length lens means an effective f-number of 50, so spherical aberrations are probably negligible. This is why you often see optical tables using simple singlet lenses to manipulate dozens of laser beams.

In imaging applications over a broad spectral range (e.g. photography and related scientific endeavors), singlets generally do *not* give good results. Commercial camera lenses are often a better choice, and high-quality microscope objectives are also available if you need to install one inside a custom physics experiment. We mostly use commercial camera lenses in Ph77, as you will see in the coming weeks.

Imaging Sensors & Cameras

Along with imaging optics, we must also consider imaging sensors. Camera systems can be categorized into a few general cases:

Point-and-shoot cameras. Here the lens and sensor are designed together for optimal overall performance. As the name suggests, the user's main job is to point and shoot, often adjusting some small number of available camera functions, for example zoom. Everything else is done for you with auto-focus, auto-shutter-speed, etc. Smartphones and simple commercial cameras are mostly point-and-shoot. These are terrific tools for general lab photography, but their lack of customization and control often limits their usefulness in scientific settings, including in Ph77.

Interchangeable-lens cameras. This is the next step up for photographers, both amateur and professional, and for science applications. The camera *body* includes the sensor and camera controls, and different types of lenses can be attached to the body. These include mid-range lenses for photographing people and scenery, *telephoto* lenses for long-distance photography, and *macro* lenses for close-up photography. These types of cameras are popular in the lab because they provide a lot of imaging flexibility at relatively low cost. With some cameras, you can operate the camera using your computer, allowing automated imaging data collection.

One common feature of commercial cameras is that their imaging sensors produce color images. This is done by covering the individual broadband camera pixels with an array of microfabricated red/green/blue filters, giving three interlaced R/G/B images that must be combined in software to produce a final color image. This task is done automatically by the camera, of course, and this is fine for many laboratory applications.

Research grade cameras. If you are building an expensive and highly specialized science experiment, you probably want a bare, high-performance sensor made of identical pixels, not covered with a complex array of RGB filters. In these circumstances, you can purchase research-grade cameras that are made especially for interfacing with a computer, allowing operation using existing laboratory software like LabView. These cameras may be ten times the cost of basic interchangeable-lens cameras with one-tenth the overall capabilities, and you may need to develop your own user interface software. But they will do exactly what you tell them to, so you go this route for high-end applications.

Laboratory Exercises

When working in any optics lab, it is important that you *do not touch any optical surfaces*. The oils in your fingerprints will degrade the optical coatings over time, reducing the overall quality of the optics. Also, do not try to clean any of the optics yourself. Cleaning any glass surface with a dry cloth will produce microscopic scratches that further reduce optical quality. If you see a smudge, let someone know and we will clean the optics with proper materials and solvents. The same goes ten times over for camera sensors, which are fragile, expensive, and certainly not designed to be touched by anything.

The Sony A6400 camera and 18-55 mm lens

For the next several weeks in this Track, your primary method of data collection will be using a Sony A6400 interchangeable-lens camera. Your first task, therefore, is to learn how to use this camera as a laboratory instrument. To begin, locate the camera on the small optical table (against the west wall in Room 210) and remove the lens. If you are unfamiliar with interchangeable-lens cameras, you will need to press the lens-release button (see Appendix I below) to twist off whatever lens is attached.

Next locate the Sony 18-55mm F3.5-5.6 lens and attach it to the camera. This lens is mostly silver in appearance and the lens body is labeled “18-55” (referring to the adjustable focal length in millimeters) and “3.5-5.6” (referring to the minimum f-numbers associated with the different focal-length settings). You will need to line up the white dot on the lens with the white dot on the camera and twist on the 18-55mm lens. (This is called a *bayonet mount* in camera world.) If you have any difficulties, ask for help. Attaching and detaching lenses from camera bodies remains a cumbersome process that manufacturers have not much improved over the years. You get used to it.

The camera should already be connected to power and the TV monitor, so turn the camera on and set its Mode Dial to Auto. Make sure the lens cap is off, and after ~30 seconds you should see a live

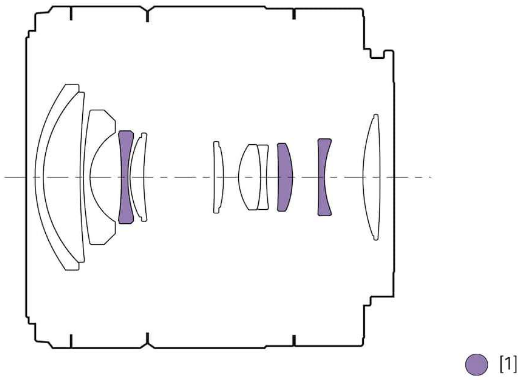

Figure 4. The internal layout of the Sony 18-55 lens. The various lenses and lens groups will move to change the focal length and focus position. Not shown is an internal iris diaphragm that changes the effective aperture of the lens. The purple lenses in this sketch have aspherical surfaces; the rest have spherical surfaces.

image on the monitor. The monitor is probably already powered on, and it can stay on indefinitely. The camera has no auto-off feature, however, so please turn it off when you leave.

The large lens ring sets the camera focal length from 18 mm to 55 mm, so turn this ring and see what happens. When you set the lens barrel to 18, you will see that the lens is short at this setting. If you set the lens to 55, the lens is longer. This makes sense; the lens needs to focus a far-away object on the camera sensor, so a longer focal-length lens needs to be farther from the sensor. Changing the focal length also changes the image zoom, as you can see on the monitor. This is what you would expect from the thin-lens picture in Figure 3.

The construction of this lens is quite sophisticated. Figure 4 shows the internal optics, consisting of various lens groups that move when you rotate the focal-length ring. Even better, the focus and aperture can be changed electronically using internal actuators via electrical contacts on the lens mount. Designing lenses like this involves quite a lot of complex engineering, driven by a large photography market.

Next set up the camera with the 18-55mm lens to view a scene like that in Figure 5. Use ambient room light only for now, set the camera to Auto mode, and set the lens to a 55mm focal length. The primary object in this photo is a *resolution target*, which can be used to measure the image resolution. Move the camera and/or resolution target so the target is 1 meter from the center of the lens. If the resolution target is mounted in the bulky translation stage, move it to the "USAF" mount (labeled on the base), as the latter is more convenient for what follows.

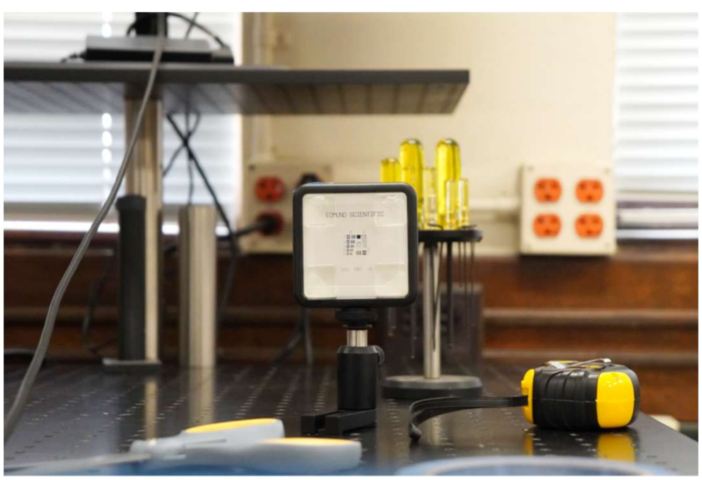

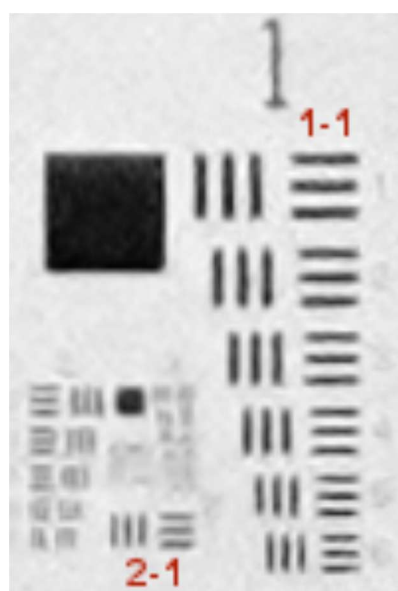

Figure 5. (Left) A photo of the USAF target at one meter from the camera, using the 18-55mm lens set to 55mm. (Right) The same image with extreme cropping.

As you rotate the smaller lens ring to adjust the focus, the camera is set up to give you a (digitally) magnified image to help you set the correct focus. In Auto mode, you do not need to deal with many settings – the camera behaves much like a point-and-shoot.

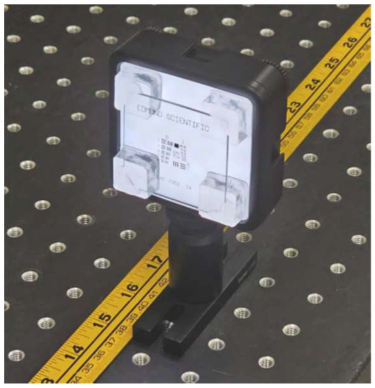

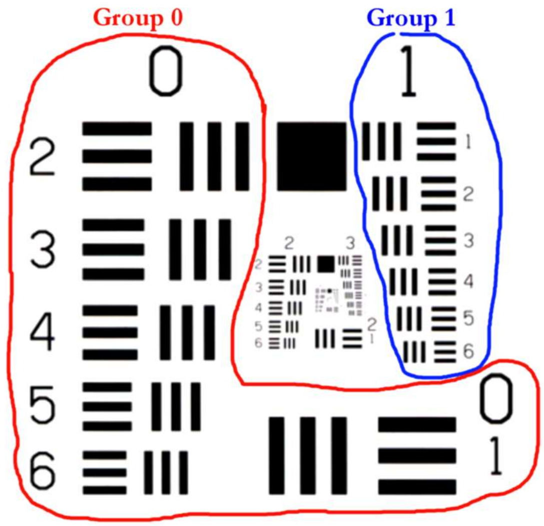

	1	2	3	4	5	6
0	1.00	1.12	1.26	1.41	1.59	1.78
1	2.00	2.24	2.52	2.83	3.17	3.56
2	4.00	4.49	5.04	5.66	6.35	7.13
3	8.00	8.98	10.08	11.31	12.70	14.25
4	16.00	17.96	20.16	22.63	25.40	28.51
5	32.00	35.92	40.32	45.25	50.80	57.02
6	64.00	71.84	80.63	90.51	101.59	114.04
7	128.00	143.68	161.27	181.02	203.19	228.07

	1	2	3	4	5	6
0	1000.00	890.90	793.70	707.11	629.96	561.23
1	500.00	445.45	396.85	353.55	314.98	280.62
2	250.00	222.72	198.43	176.78	157.49	140.31
3	125.00	111.36	99.21	88.39	78.75	70.15
4	62.50	55.68	49.61	44.19	39.37	35.08
5	31.25	27.84	24.80	22.10	19.69	17.54
6	15.63	13.92	12.40	11.05	9.84	8.77
7	7.81	6.96	6.20	5.52	4.92	4.38

Figure 6. The USAF 1951 resolution target (above left). Note that the six elements of Group 1 are placed together, while the six elements of Group 0 are placed in an L configuration (above right). Smaller groups on the target are similarly arranged.

The tables at left show the bar spacing in lines/mm (top table) and microns (bottom table). The notation is Ngroup-Nelement, (as in 5-2), where ***the leftmost column shows the group numbers****.* ***Do not confuse 3-1 with 1-3 (for example)… these are different.***

The USAF 1951 resolution target

Figure 6 describes the layout of the USAF 1951 resolution target. As the name suggests, this standard was created in 1951 for the US Air Force, and the target consists of a 2x2 inch glass slide with a series of metal-film bars fabricated onto the surface using photolithography. (Note that these bars are not held on with great adhesion, so please do not touch the glass surface, and do not try to clean it.) The bar triplets are arranged into *groups* and *elements*, and the spacing of the bars is given by

$$\frac{\text{lines}}{\text{mm}} = 2^{\left(\text{Group\#}+\frac{\text{Element\#}-1}{6}\right)} \qquad (4)$$

(To be clear, this expression is 2 raised to the number in parentheses.) For example, the 5-1 element has 32 lines/mm, so the spacing between lines is 31.25 μm. The tables in Figure 6 list all the spacings for this resolution target.

Exercise 3. Use the camera and 18-55mm lens to capture a photo like what you see in Figure 5, using the settings described above (ambient light only, with the lens set to 55mm). Note that autofocus has been disabled on this camera (please leave it that way), because autofocus does not work well in a laboratory setting. Autofocus is great for people and scenery, but the algorithms are not well suited for resolution targets and scientific imaging. The best way to focus is press the C1 button to cycle through the available zoom settings, and then stop at the maximum magnification. You can use the *control wheel* (see Appendix I for all the various buttons on the camera) to move the magnification window (try it). With this maximum-zoom image on the TV, turn the focus ring back and forth, back and forth, back and forth … until you feel you have achieved the best focus you can. If you cannot see that the 2-1 bars are pretty clearly resolved (see Figure 5), then something is wrong. And remember this method for obtaining the best image focus; you will need to focus carefully throughout this lab.

Once you are satisfied with the focus, use the remote shutter release to take a picture without shaking the camera. It works best to point the shutter-release at the IR sensor on the front of the camera body (shown in Appendix I below). Ask if you cannot find the remote shutter release, or press the shutter button on the camera with as little shake as you can. If the shutter time is less than $1/100^{\text{th}}$ second, then the remote shutter release is probably not necessary.

Remove the SD card from the camera (inconveniently located on the bottom of the camera, in the battery compartment) and open your saved photo using your own computer (an SD card reader is available in the lab). If you do not have your computer handy, you can use the iMac in the lab. Make sure your image is 6000x4000 pixels in size, then make a copy and crop to zoom in on the resolution target, yielding an image something like the second one shown in Figure 5. Crop way in, so you can easily see the barely resolved sets of bars. Add both the uncropped and cropped versions of your photo to your e-notebook. You should be able to see the resolution limit clearly in your .pdf file, even if it were simply printed onto 8.5x11-inch paper.

From your image, what is the smallest set of bars (labeled by group# & element#) in which the individual bars in a triplet are barely resolved? From this, record the resolution limit in μm on the target. *Be careful not to reverse the group and element numbers in the table*. For example, if you can barely resolve the 3-1 bars, then record the resolution as 125 μm (as tabulated in Figure 5). Just for comparison later, record the resolution **also in μrad** (microradians).

Note that this definition of resolution is somewhat qualitative, as it depends on exactly what you mean by "barely resolved." One can define the resolution with better mathematical precision, but doing so would involve a lot of numerical image analysis, and going down that road is not especially interesting at this point.

If you want to do the Optics Track right, then you should find a suitable image-processing program for your personal computer. Just about every photo you take in the lab will need to be cropped and processed, often with some brightness adjustments, etc. All of this (and much more) is easy to do with an image-processing program, and difficult by any other means. Some examples include Adobe Photoshop (an industry standard, for both Windows and Mac, but expensive and cumbersome), Photoshop Elements (the "light" version of Photoshop, much cheaper), PaintShopPro (like Elements, but even cheaper, and a bit better … although Windows only), Inkscape (free, for Windows, Mac, and even Linux), and others.

Beyond Ph77, you will find that learning an image-processing program (any variety, even with little mastery) will be highly useful if you ever intend to give a talk, write a report, or display visual information in just about any circumstance. If you are planning a career that might involve communicating with other humans using images (and that includes the majority of STEM careers), then you would do well to find an image-processing program that works for you.

Next turn on the light behind resolution target (always use the lowest brightness level; it saves battery, and you do not need more light than this). The LED array is powered by a rechargeable battery, so please turn this back-light off when not in use. If the battery indicator (on the back) looks low (4 lights = fully charged), please plug it in while you are using the light (there is a USB-C cable for this purpose on the wall near the optical table) and leave it plugged in if it is not fully charged when you leave.

The resolution target is meant to be used with back illumination, so using this light source will yield better photos. If you try to reproduce Figure 5, however, you will find that you cannot take an acceptable photo using the camera's Auto mode. This is a common experience in scientific photography – Auto mode is designed to take nice photos of people and scenery in ambient light conditions, but the Auto algorithms often fail miserably in the lab. The solution to this problem is simply to move beyond Auto mode, so beyond a simple point-and-shoot experience.

To see how this goes, set the Mode Dial to M, which puts you into full *Manual Mode*. Unlike Auto, now you must consider a host of camera settings to get a good picture. Try these:

1) Press the top side of the Control Wheel (see Appendix I) to cycle through the various display settings. Stop when you have the simple display that includes a brightness histogram. Here you see the shutter time, f-number, metering mode (M.M.), and ISO number. If M.M. is negative, the image might be too dark. If M.M. is positive, then the image might be oversaturated. Zero is about right, although setting M.M. a bit negative is often better for science photos. A bit too dark is easily corrected in software later, while an oversaturated image is difficult to correct.
2) Rotate the Control Wheel to change the shutter time. You will see the image get darker for smaller shutter times, as you would expect.
3) To get out of this adjustment mode, press the shutter button halfway down (not enough to take a photo). This works to get out of most adjustment modes. Pressing the center button of the Control Wheel often works also.
4) Press the right side of the control wheel to change the ISO setting, then rotate the control wheel. As described in Appendix I, the ISO setting refers to the amplifier gain in the sensor. High ISO is good for low-light situations, but a low ISO setting produces less electronic noise and yields a less grainy image.
5) Rotate the Control Dial to change the lens aperture, as indicated by the f-number on the display. Lenses have fewer aberrations at high f-numbers, but low f-numbers let more light into the camera.

In practice, full manual mode is quite inconvenient to use, requiring many careful settings just to get a reasonable image brightness level. A better choice in this lab is to set the Mode Dial to A, which means *Aperture Priority*. In this mode, the f-number stays wherever you set it, while the shutter time automatically adjusts to set M.M. to zero automatically.

Remember, however, that Aperture Priority is not full auto… you still need to manually set the ISO. If the shutter time is longer than $1/100^{th}$ second (note that 1" means a one second exposure time), then set the ISO manually to a higher value. Also if the shutter time is $1/4000^{th}$ of a second (the fastest setting), then choose a lower ISO.

Using Aperture Priority mode, again cycle the display until you see the histogram of image brightness levels. Often you will see that some pixels are perilously close to saturation, and this problem can be especially common in scientific imaging. For better results, press the bottom side of the Control Wheel and set the *Exposure Compensation* to -1.0 (called -1 "stops" by photographers). Now the histogram will show that all pixels are comfortably below saturation. These settings should give you good results over a broad range of circumstances.

As a general rule in this Track, staying with Aperture Priority is probably best. Just remember to:

1) Set the mode dial set to A.

2) Always keep an eye on the shutter time. If it is more than 1/100th of a second, then image shake might be a problem. If it is at 1/4000th second (the camera minimum), then the image will likely be over exposed and saturated. In most circumstances, therefore, the shutter speed should be between 1/100th and 1/4000th second. To achieve this, manually adjust the camera ISO.
3) At the same time, try to avoid really high ISO settings, as electronic noise can yield images that are quite grainy when the ISO is high.
4) Also keep an eye on the aperture (the f-number on the display). A lens usually works best in the middle of its aperture range (by design), although you might want something different.
5) Watch the histogram on the display. If you see values near maximum, the image may be overexposed. If the values are all quite low, then the image will be underexposed. Manually adjust the Exposure Compensation to get the histogram you want.
6) When focusing, cycle through the zoom settings using the C1 button. Move the focus window around to find a good place to focus on.

Practice using Aperture Priority using the 18-55mm lens with the resolution target, both with the back-illumination light on and off. You will quickly see that different ISO values must be used in these two circumstances. Aperture Priority is not as convenient as auto mode, but it allows easy access to a greater range of camera settings. [A quick aside relating to camera lingo: one "stop" is basically one octave – a factor of 2. The camera's ISO, aperture, and exposure compensation numbers can each be adjusted in steps of 1/3 stop. This is just logarithmic scaling in camera-speak.]

Note that this camera has an enormous *dynamic range* regarding light intensity. In low-light conditions you might set the shutter time to as high as 30 seconds (to capture as much light as possible) and use an ISO setting as high as 102400 (a high amplifier gain on the optical sensor). In bright-light conditions, the shutter time can be set as low as 1/4000th second and the ISO can be set as low as 100. If you do the math, together the shutter and ISO settings give you a factor of over *100 million* in dynamic range. Adjusting the lens aperture gives you even more.

Exercise 4. Take another picture of the resolution target for your e-notebook, much like you did in the previous exercise, but now using back illumination. Set the aperture to f/14, near the center of its range. Focus carefully once again at maximum magnification (if you are not turning the focus knob back and forth at least a dozen times, then you are probably not getting the optimal focus), and again record the resolution limit in both μm and μrad. This new image gives you a reference for the 18-55mm lens going forward, and *from now on you should always use back-illumination when taking more photos of the resolution target*. If you are careful with the focus, you should be able to resolve the 2-1 bars, and you can see that the resolution is clearly limited by *pixelation* (meaning that there are just a few camera pixels making up a set of bars, as illustrated at right). Check your image resolution now, before you leave the lab. If you are not resolving the 2-1 bars, ask for help.

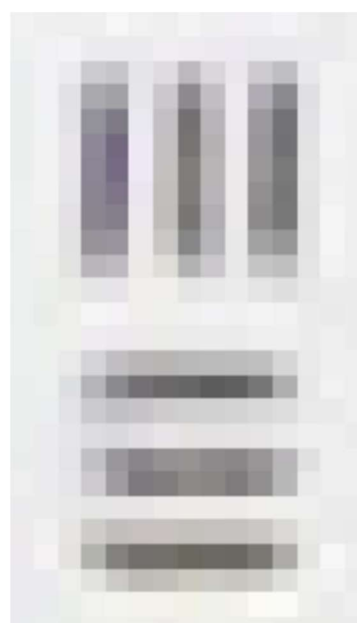

Crop your image way down and make an image like the one at right for your smallest "barely resolved" bars. (Seriously, make a *huge* image of just these bars; see that pixelation.) Note that having ~3 pixels spanning two adjacent bars is typically enough to "barely resolve" the bars.

The Sony 55-210 mm lens

Now that you have had some practice using this camera, locate and install the Sony 55-210mm lens located on the optical table. This lens is similar in operation to the 18-55 mm lens above, but it covers a range of longer focal lengths. Not surprisingly, as you can see, the physical length of the 55-210mm lens is longer than the 18-55mm lens.

Exercise 5. With this lens focal length set to 210mm, measure the resolution limit (in microradians) as a function of the f# of the lens, using the resolution target placed one meter from the camera. Start with f/32 and view the image on the monitor at the highest magnification (by pressing C1 to cycle the magnification setting). If you are careful, you can determine the smallest "barely resolved" bars directly on the TV monitor. Use the blue lab jack to (gently!) support the end of the lens, as this greatly reduces image shake as you focus. You have to focus especially carefully in this exercise, so take your time.

When you are satisfied with f/32, repeat the process for every second f# setting (so f/32, 25, 20, 16, 13, 10, 8, 6.3). You will find that the resolution does not change a huge amount as you change the f#, but press on. If you are careful to always get the best possible focusing, the data quality will be better than you think. If you are unhappy with your first pass through the f#s, call this a "practice run" and try again for your final data. Getting good results in experimental physics often takes several attempts!

Next plot your resolution data (in microradians) along with a theoretical model of this lens, including contributions from pixelation and the diffraction limit. Figure 7 shows roughly what your model should look like (not including the data points you will have in your plot). The discussion above should give you all the information you need for the theory curve and the measured resolution. If your baseline resolution is higher than expected, this might be an indication that you are not focusing the lens well enough. It is especially tricky to focus at low f-numbers. If you make your measurements carefully, the data points should fall on the theory curve reasonably well.

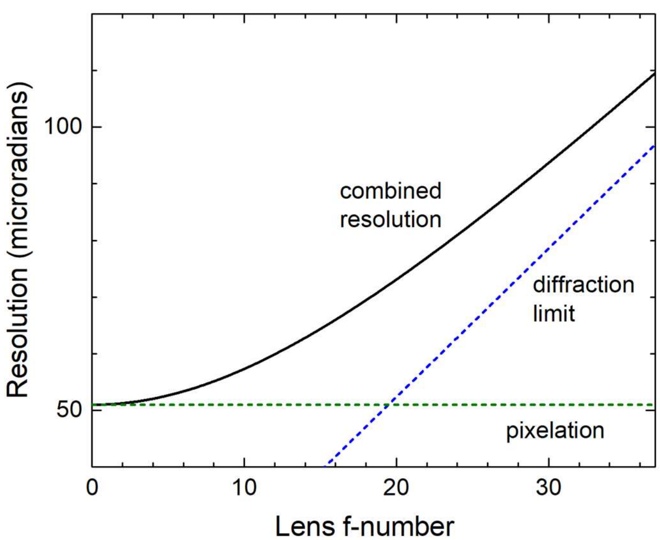

Figure 7. A theoretical model of the resolution of the 55-210 lens when set to a focal length of 210mm. The blue dashed curve shows the diffraction limit, which changes with the size of the aperture. The green dashed curve is from pixelation, assuming a spatial resolution of δx=2.7 pixels (roughly), converted to angular resolution. The combined curve is then the two components added in quadrature.

Optimization of camera lenses

One important take-away message from these exercises is that image resolution is usually limited by some combination of: 1) pixelation, 2) lens aberrations, and 3) the diffraction limit. Commercial lenses are often designed so that all these factors are important, because this is how engineering designs tend to evolve. Imagine you are designing a camera/lens system. If pixelation was a negligible factor, then you would want to use fewer pixels to save money. If lens aberrations were negligible, you would use fewer lens elements and thus increase the aberrations slightly, to save money. And if the diffraction limit was not a factor, you might just build a smaller lens, again to save money. This continual optimization explains why, in the commercial marketplace, lens performance is often limited by all these factors.

Given these considerations, what limits the image resolution often depends on the lens aperture. When the aperture is set to the middle of its range, the resolution is likely limited by all the above factors. With the aperture wide open, the diffraction limit is less important while lens aberrations are worse, so the latter may limit the resolution. And finally, when the aperture is at its smallest setting, lens aberrations are reduced, and the image resolution is then likely limited by diffraction.

Learning from a bad example

The Sony lenses you have been using so far are well-crafted multi-element lenses, designed to reduce spherical and chromatic aberrations as much as possible. Now let us look at what can happen if you try using a simple one-element lens. First locate the *150mm singlet plano-convex lens* in the lab, which is mounted in a black tube about 25cm in length. As the name suggests, this lens consists of a single plano-convex element with a focal length of 150mm mounted in a hollow tube. The only adjustment you can make is the lens aperture, using the iris diaphragm in front of the lens (be gentle with the iris; it will break if you twist it hard enough). Mount the tube to the camera as usual and support the end of the lens tube using the lab jack located on the optical table. There is no focus adjustment on the lens tube, so move the resolution target to focus.

Exercise 6. Manually adjust the iris in front of the plano-convex lens so the aperture is about 1cm and you should see a reasonable-looking image. But set the iris wide open and you will see a badly washed-out (low contrast) image. Take photos at both these settings, crop down to the center of the resolution target, and measure the image resolution in both μm (at the object) and μrad for your e-notebook. Take a third photo with the aperture set to 3mm and you should notice that the resolution is degraded by the diffraction limit. Verify that diffraction is indeed the limiting factor when the aperture is this small.

When you see a low-contrast image like the one at full aperture, your first thought might be scattered light. Stray light striking a camera sensor will add a roughly uniform background brightness to the image you are trying to obtain. With this singlet lens wide open, however, the uniform background comes from lens aberrations. Light from the edge of the lens is so badly distorted that it basically spatters all over the image plane, washing out the dark bars in the image of the resolution target. These aberrations are worse when the light rays enter or leave the glass at steep angles, and this happens most at the edges of the lens. The f-number of this lens is about f/3 when the iris is wide open, and that is low enough to produce a terrible image using a plano-convex singlet. Using a 1-cm aperture (f/15) gives a much better image.

There are a few take-away lessons from this exercise:

1) Aberrations can be large with singlet lenses, and plano-convex lenses are the worst. For this reason, they are rarely used for imaging applications.
2) Nevertheless, you can always reduce the aberrations by "stopping down" the lens (i.e., reducing the size of the aperture). As a general rule, even singlet lenses work reasonably well when the f-number is higher than f/10.
3) Image resolution is not the only parameter to consider with an optical system. Bad lens design can also yield low-contrast images, as can stray light. If you see images with poor contrast and washed-out dark levels, you probably have a poor optical system.

The human eye
As another singlet example, the lenses in your eyes are positive-meniscus singlets with an f-number that varies from about f/8 in bright daylight to about f/2 at night. Most of the focusing happens at the front surface of the eye (which has the largest jump in refractive index), which is kept smooth by surface tension in a "tear film" (pronounced like the tears you cry) that you constantly replenish when you blink. With 20/20 vision in daylight, your eye can resolve down to about 300 μrad, but this gets worse at night because of singlet aberrations at f/2 (along with sensor noise). The aberrations are worse off-axis, and the spacing of your retinal cells is lower off-axis as well. The diffraction limit is not important in this optical example, mainly because the focal length of the eye lens is so short.

Macro lenses
To continue our series of camera lenses, locate and install the Sony E 30mm F/3.5 Macro Lens, the "macro" signifying that this lens is designed for close-up imaging. It has a remarkably broad focusing range, with a *working distance* (the distance from the object to the front of the lens) from infinity down to about 1cm. At the closest working distance, the field of view with the Sony camera is only about 2cm.

Exercise 7. Set the working distance to 1.5 cm (close to the minimum value but leaving some wiggle room to focus), capture a photo of the resolution target at f/5, then measure the resolution limit, expressing your answer in both μm (at the object) and μrad. The thin-lens approximation is not so accurate here, so assume a lens-object distance of 30 mm. Calculate the diffraction limit for this lens, and you will see that your measured resolution is limited mainly by pixelation. (This should also be fairly obvious if you zoom into your image.)

If you think about all this, you can draw several conclusions about macro lenses like this one:
1) One big advantage of using a macro lens is the short working distance. This means a fairly ordinary resolution limit in μrad translates into an impressively small resolution limit in μm at the object. In general, macro lenses are useful when you want to photograph small objects at resolutions down to around 10 μm.
2) The best way to improve the pixelation problem is to either: 1) produce a sensor with smaller pixels (difficult), for 2) increase the distance from the lens to the camera sensor… while maintaining a small working distance. And doing the latter basically gives you an optical microscope….

Microscopy
Let us next turn our attention to optical imaging at the smallest scales, as this is something that comes up a lot in experimental physics. (Physicists also put a lot of effort into imaging at the largest physical scales, but we will not be delving into telescope optics in this Track. Nor will we venture into electron microscopy or other advanced imaging techniques at sub-micron scales. Experimental physics is quite a broad subject.) The main thing that distinguishes microscopy from macro imaging is that the image resolution in microscopy is usually *limited entirely by the diffraction limit*. Of course, inexpensive microscopes may be limited by optical aberrations as well, but well-crafted microscope optics are nearly always limited by diffraction.

Commercial microscopes are a lot like point-and-shoot cameras, in that you can productively use them without having any real understanding of their inner workings. Laboratory microscopes come in a large variety of different types, including:

Stereo microscopes. These have two separate optical paths, one for each eye, and are meant to be used for quick examination of relatively large objects. Stereo microscopes are real workhorses in the lab, especially useful for electronics work (with small surface-mount components) or assembling small parts. Novices squint when looking at small things; pros use appropriate magnifying optics to get a much better view.

Stereo microscopes are generally low power (i.e. low magnification), often with adjustable magnification. Some have imaging sensors, either using a third optical path (a trinocular microscope) or using a flip mirror to send one of the dual paths to the camera. The brain is much more comfortable working with two eyes than one, plus the converging optical paths give some degree of stereoscopic depth perception.

Another quick word about magnification. A 10X stereo microscope means that the object appears (to your eyes) 10X larger than it is. The 10X only makes sense when you are using your eyes. Once you switch to an electronic imaging sensor, the 10X moniker becomes ill defined. In that realm, the things that matter are the field-of-view and image resolution.

High-power microscopes. These are usually made for biology & medical imaging because that is where the market is. But they find their way into physics experiments as well. The cost is higher than stereo microscopes, often much higher. These instruments usually have a single optical path (stereoscopic viewing doesn't work well at high magnification) with interchangeable fixed-magnification *microscope objectives* mounted on some kind of rotating turret. High-power microscopes have a small field of view and often a wicked-small depth of field, making them generally difficult to work with.

Once you get into the realm of high-power optical microscopy, you soon encounter a plethora of advanced techniques like fluorescence microscopy (imaging fluorescent molecules within the sample), confocal microscopy (using scanning laser illumination), phase contrast microscopy (using interferometry to examine optical phase shifts in the sample), and more. Many of these techniques were developed to increase the image contrast in biological samples (which can be mostly semi-transparent with indices of refraction close to that of water – tough to image). These techniques are all quite specialized (and expensive), so beyond the scope of this Track. But you may encounter them in physics labs at some point in your career.

Working with microscope objectives

Commercial microscopes are nearly always your best option if you can bring your sample to the microscope. However, if you are working on a cutting-edge physics experiments done inside a vacuum chamber, or inside a cryostat, then you may have to bring the microscope to the sample. And that is often accomplished by buying a microscope objective and building it into your experiment, perhaps projecting the image through a window to an imaging sensor in the outside world. This kind of DIY microscopy is remarkably useful in the lab, and it is easy to demonstrate using the set-up you have just been using. So, let us take a quick look at the world of microscope objectives.

A microscope objective is just another complex lens, so the underlying physics is essentially the same as with camera lenses. But the language is a bit different. A microscope objective is usually characterized by several parameters, including the *Numerical Aperture* (NA), the *Resolving Power* (RP), the *Working Distance* (WD), and the *Depth of Field* (DoF).

The Numerical Aperture is related to the f-number by

$$NA = \frac{1}{2 \cdot (f\#)} \tag{5}$$

and the Resolving Power is what we have been calling the image resolution. Same stuff, just different names. The Working Distance is the distance between the object and the front of the objective (quite an important parameter if you are trying to peer into a vacuum chamber or cryostat), and the Depth of Field is essentially the distance between the nearest and the furthest objects that are in acceptably sharp focus in a single image. Because the resolving power is set by the diffraction limit, we have

$$RP = \lambda \cdot (f\#) = \frac{\lambda}{2 \cdot NA} \tag{6}$$

The depth of field is given by

$$DoF = \frac{\lambda}{NA^2} = \frac{4}{\lambda}(RP)^2 \tag{7}$$

Note that you want a small Resolving Power to image small objects, and this immediately means a very small depth of field. The math giving these relationships is quite simple once you are solidly in the diffraction-limited regime.

Mitutoyo Plan APO Objective	Numerical Aperture	Working Distance (mm)	Depth of Field (µm)	Resolving Power (µm)
2X	0.055	34	180	5
10X	0.28	34	7	1.0

Table 1. The two microscope objectives we use in the Optics Track.

Table 1 shows these parameters for the two microscope objectives we have in the lab. These come straight from the manufacturer's spec sheet, assuming λ ≈ 550 nm for visible light. As you can see from the above equations, the RP and NA are essentially redundant parameters, but they are often both listed in the specs.

Exercise 8. Locate the microscope objective tube in the lab, along with the 2X and 10X objectives. (These objectives cost >$1000 each, so please handle them carefully.) Use both objectives to image the resolution target and measure the resolving power (same as optical resolution) as you did before. And include cropped images of both photos in your e-notebook. For best results, mount the target in the *micrometer stage* provided, as the depth of field decreases rapidly with higher resolving power (see Table 1). [Translation: obtaining a good focus means moving the resolution target by very small distances, which requires a micrometer.] Be sure to adjust the ISO and shutter speed to produce a good image. Compare your measured resolutions with the manufacturer's specifications. If you cannot *clearly* resolve the 7-6 bars using the 10X objective, then something is wrong.

A technical fine point (optional). You may have noticed that the microscope objective tube includes an additional lens placed directly behind the objective. This lens is there because these objectives are "infinity corrected", which means that they were designed to have an image plane at infinity (behind the objective). The second lens (relatively weak with a 300 mm focal length) simply moves the image plane from infinity to the position of the camera sensor. If this lens is absent, then the objective works about the same at the center of the field of view, requiring just a slightly different working distance to move the image plane from infinity to the camera sensor. However, the off-axis focus is substantially worse.

In a nutshell, these objectives were designed to work with the image projected to infinity. With the second lens in place, the focus is quite sharp across the entire field of view. With the second lens removed, you are using the objective in a way that was not intended, and that has consequences. In particular, the focus will not be sharp across the entire field of view unless that second lens is included.

This is a general rule when using commercial optics: if you buy a high-quality commercial lens or microscope objective, use it as it was intended. Otherwise the image quality will be degraded, and that could mean that you bought a really expensive lens just to get low-quality images.

Lens design

For small-scale science applications, your best option is usually to stick with commercially available optics like camera lenses or microscope objectives. If you work on big projects, however, like a next-generation optical telescope, a spacecraft imaging system, or a microscopy system built into a cutting-edge condensed-matter experiment, you may have no choice but to design (often together with contractors) a set of custom optical components optimized for the task at hand. Needless to say, that level of design expertise is beyond the scope of this Track! But it is instructive to peek behind the curtain to see how the process works.

Ray tracing in optical design

The first thing to note in lens design is that dealing with pixelation is straightforward – you need only consider how many pixels are needed to accomplish your goals. Similarly, the diffraction limit is a straightforward design issue because it is easy to calculate. In most optical design projects, therefore, dealing with aberrations is the challenging part.

When it comes to the theory of lens aberrations, useful analytical results are a bit scarce. Thus the primary tool in lens design (as well as mirror design in large telescopes) is numerical ray tracing. As

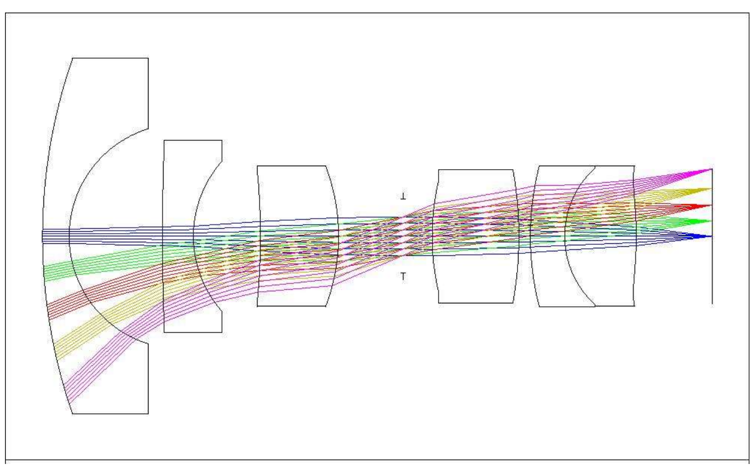

Figure 8. This example shows a series of light rays passing through a complex lens. The colors indicate different bundles of rays (nothing to do with wavelength), each bundle containing a set of parallel input rays. As the rays pass through the lens, each bundle focuses on a different point in the image plane.

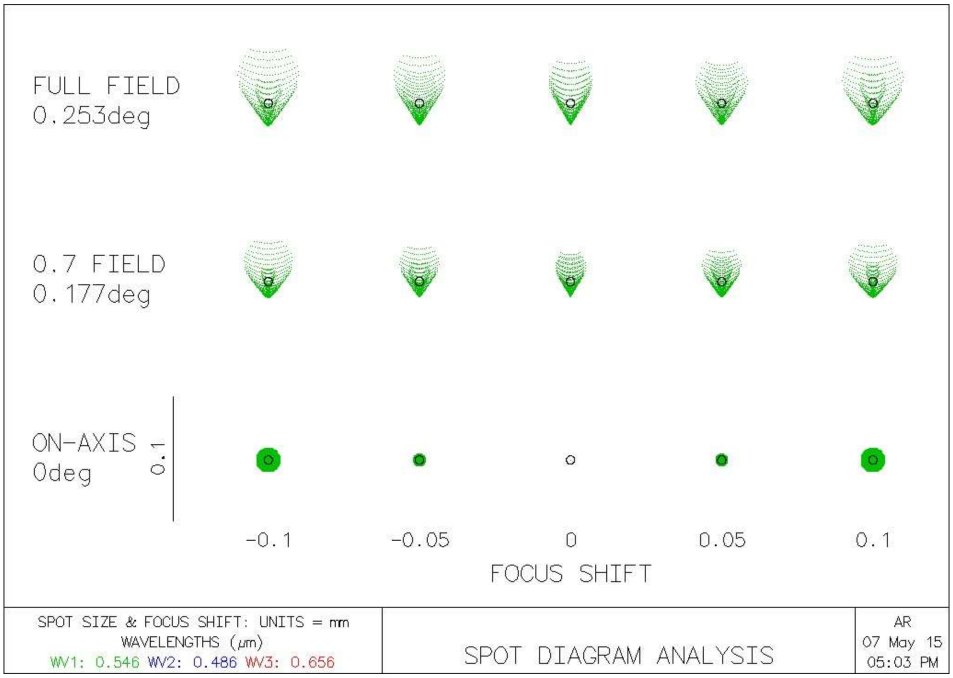

Figure 9. An example showing how several ray bundles strike the image plane, as a function of the position of the focal plane along the optic axis (focus shift) and the angle of the incoming ray bundle relative to the optic axis.

the name suggests, you enter your entire lens design into a sophisticated software package, including all the properties of all the lens elements (mostly surface shape and refractive index, the latter being wavelength dependent), and then let the computer spit out the results. Ignoring diffraction is a reasonable approximation in many cases, so the software simply calculates the refraction of narrow light "rays" at each lens surface, giving results that look something like what you see in Figure 8. In this sketch, each bundle of rays could represent light from different stars in the sky. Because the stars are at infinity, light from each star enters the lens as a parallel set of rays, and these rays come to a point on the image plane. Note that this Figure does not give you the full picture; in a real lens design, the individual bundles would be much larger, so each set of parallel rays would strike every point on the front of the lens.

After calculating all these ray trajectories, the software package then looks at where each ray strikes the image plane, and an example of what this might look like is shown in Figure 9. If you look at on-axis rays (bottom row in the Figure), you see about what you would expect: the rays meet at a point when the image plane is at the lens focus (focus shift = 0), and the rays are blurred when the image plane is not positioned correctly at the lens focus. More importantly, if you look at off-axis rays, then there is no focus shift where the rays come to a sharp point. What this means is that you get a beautifully sharp image at the center of the field of view, while the image is a bit fuzzy off center. And if you look at rays at different wavelengths, the ray-tracing results will look a bit different.

The whole point of adding lots of lens elements (see Figure 2) is to make these spot diagrams look better. With a good lens design, the spots will all have about the same amount of smear, regardless of the light wavelength and the position across the field of view. And of course the focal spots from the different ray bundles should be as small as possible.

The challenge with ray-tracing software is that it only solves the forward problem: if you input a lens design, then the calculations will tell you all about the imaging properties of that lens. But the software will not tell you how to build a better lens. To this day, lens designers are still discovering (and patenting) new and better lens designs, using different lens elements arranged into different lens groups, incorporating non-spherical lens surfaces and new types of glasses with novel refraction properties. As new imaging sensors are developed (e.g. using tiny sensors in smartphones), new lens designs follow. This kind of optical design work is an area of much current research in many areas of science and technology.

What lenses to use?

If you stay with commercial optics, the question then arises as to what optics to buy. Unfortunately, as with many commercial products, there is no easy answer to this question. Here are a few guidelines: 1) Stay away from really cheap optics. These tend to have too few optical elements, poor surface quality, etc., and often the resulting images are poor. 2) At the other extreme, you probably do not need the most expensive optics you can find. Yes, high-end lenses will deliver high-quality images, but at a high cost. As a general rule, go with premium optics only if your needs are so great that you will be directly measuring the optical performance of these parts (which means that you are serious and you know what you are doing). Otherwise… 3) Intermediate-price-range optics are usually good enough for most applications, especially if they come from reputable manufacturers.

Depth of field (Optional)

Another factor that often matters in imaging systems is the *Depth of Field (DoF)*, equal to the range of focal distances that give an acceptably sharp focus on the sensor. Only one focal distance is exactly in focus, of course, so how you define the DoF depends on what you mean by an "acceptably sharp" focus. Glossing over this detail for now, the standard (approximate) formula is

$$DoF \approx \frac{2u^2Nc}{f^2} \tag{8}$$

where DoF is the distance between the nearest and farthest acceptable focal positions, u is the object distance, N is the f-number, f is the focal length of the lens, and c is the "circle of confusion," equal to the acceptable size of the in-focus spot.

Exercise 9. (Optional) To see the Depth of Field in practice, use the 18-55mm lens to measure the image resolution as you move the USAF target through the focal point. Set the lens focal length to 55mm and set the focus distance to 500 mm. Leave the focus distance unchanged as you go through this exercise. Start with the largest lens aperture (f/5.6).

Place the resolution target against the meter stick on the optical table, so you can easily slide it along the ruler (using the sliding mount). Now you can quickly take a series of images at different object distances, by moving the target and hitting the remote shutter button.

To analyze these images, open each and estimate the set of bars that are just barely separated from one another. (A simple visual inspection of the images is fine; no need to add them all to your e-notebook.) Convert these observations to a measurement of the effective image resolution in microns, using the table in Figure 6. Plot your data and you should get something like what you see in Figure 10. Do the same exercise at f/20.

To compare your results with Equation (8), assume that the circle of confusion is equal to the bar spacing divided by six (for the smallest barely resolved bars, as usual).

The main takeaway from this exercise is that the depth-of-field is greater when the f-number is higher, as expected. Here again, you see another compromise you must face in imaging: a higher f-number gives a better depth-of-field and reduced aberrations, while a lower f-number gives you greater light-gathering and a sharper diffraction limit. If you are using a point-and-shoot camera, you needn't worry about any of this; you just point and shoot. But if you are designing an imaging system for any kind of scientific application, then you may have to consider these issues. There is no one "answer" for what kind of lens or imaging system is best in a particular application; there are many compromises to make, and each specific circumstance is different.

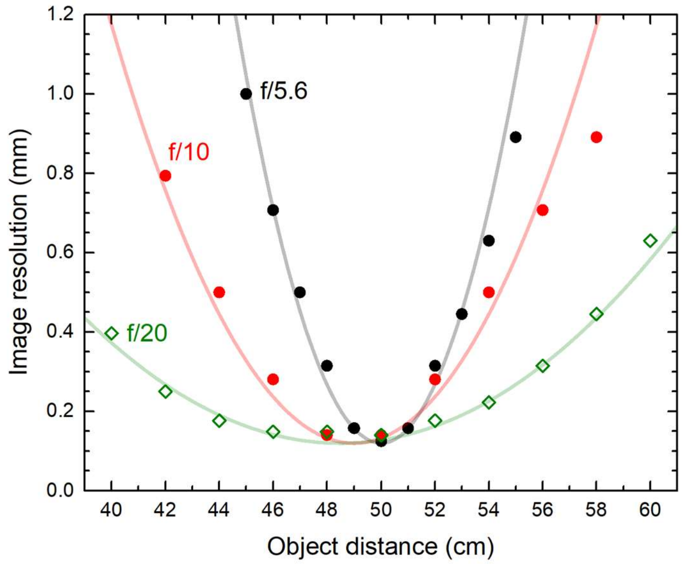

Figure 10. A plot of the measured image resolution limit as a function of the distance from the 18-55mm lens to the resolution target. Set to a focal length of 55m and f/5.6, the image is in focus at an object distance of D = 50 cm, and the focus degrades quickly away from 50 cm. Note how the depth-of-field is larger at higher f numbers.

Part 2. Interferometry

Our next topic deals with interference effects in optical systems. You are probably already familiar with laser interferometry at some level (e.g. from Ph3), and this technology is exploited in many areas of modern science (one prominent example being the detection of gravitational radiation using interferometric detectors). By far the most common application of optical interference effects, however, is dielectric coatings, so let us begin the discussion there.

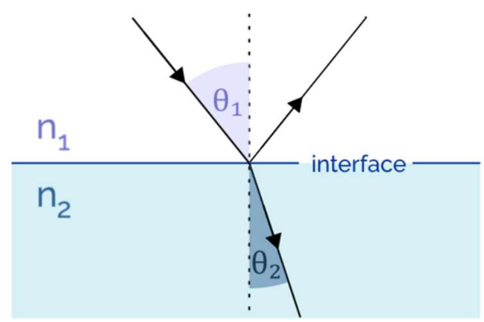

Figure 11. An illustration of Snell's law of refraction.

In the realm of optics at visible wavelengths, a dielectric refers to any of a wide selection of (usually transparent) materials like glasses, crystals, plastics, etc., each being characterized by the index of refraction as a function of wavelength, $n(\lambda)$. And dielectric materials generally have low electrical conductivities. A glass plate with an aluminum coating makes a good mirror because aluminum has a high electrical conductivity, as this pulls the electric field to zero at the metal surface. Dielectric coatings are used to make mirrors as well, but via a different physical mechanism. Both types of mirrors are in widespread use, and hybrid varieties are popular also.

When a light ray strikes a flat interface between dielectric materials, the angle of the transmitted ray is described by Snell's Law

$$n_1 sin(\theta_1) = n_2 sin(\theta_2) \tag{9}$$

while the reflected ray has $\theta_1' = \theta_1$, called *specular reflection*. The intensity of the reflected ray is described by the *Fresnel equations*

$$R_s = \left|\frac{n_1 cos(\theta_1) - n_2 cos(\theta_2)}{n_1 cos(\theta_1) + n_2 cos(\theta_2)}\right|^2 \tag{10}$$

and

$$R_p = \left|\frac{n_1 cos(\theta_2) - n_2 cos(\theta_1)}{n_1 cos(\theta_2) + n_2 cos(\theta_1)}\right|^2 \tag{11}$$

where R_s and R_p are the "skipping" and "plunging" polarizations, respectively. The skipping polarization can be visualized by a flat rock skipping on the surface of a lake – the electric field vector, like the thin rock, is mostly horizontal relative to the surface. If you rotate the rock 90 degrees, it no longer skips over the water, but immediately plunges into it – and this gives you a nice visualization of the plunging polarization. For the special case of normal incidence, the Fresnel equations reduce to

$$R = \left|\frac{n_1 - n_2}{n_1 + n_2}\right|^2 \tag{12}$$

independent of polarization. In all cases, the transmitted intensity is simply $T = 1 - R$, as required by energy conservation. These expressions assume no absorption in an ideal dielectric material, and this approximation is often quite good for real optical-quality materials. If absorption cannot be ignored, then of course the equations become more complex. In practice, the heavy lifting is invariably done by software, just as we saw above with lens design.

One thing to remember from this is that the reflection coefficient for normal incidence is proportional to $(\Delta n)^2$. A typical optical glass like BK7 has an index of refraction of about 1.5 (wavelength dependent), so in air this gives $R \approx 0.04$. If light passes through a pane of glass, then the two air/glass interfaces reflect about 8% total. Diamond has an exceptionally high index of refraction of about 2.4, giving $R \approx 0.17$. [Diamonds are especially popular gemstones because: 1) the large R gives them their sparkle, called "brilliance" in the jewelry business; 2) the index of refraction changes rapidly with wavelength, yielding colorful refraction effects that jewelers call "fire"; and 3) the material is extremely hard and therefore scratch resistant. The facets added to diamond gemstones are designed to enhance these effects.]

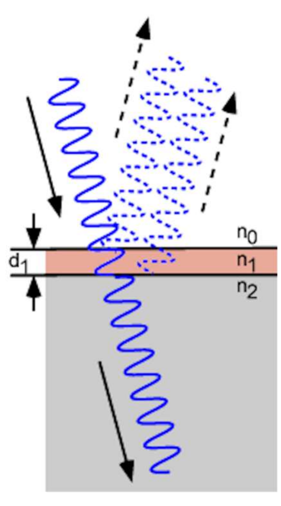

Figure 12. The simplest form of an antireflection coating consists of a single dielectric layer on the surface of a glass substrate with $n_0 < n_1 < n_2$. For a normally incident beam (tilted here for illustration purposes), reflections from the n_0/n_1 and n_1/n_2 interfaces will destructively interfere if the phase shift between the two is 180 degrees. This happen if the thickness d_1 is equal to one-quarter of the light wavelength in the n_1 medium – thus this is called a "quarter-wave" antireflection (AR) coating. Note that the total reflection goes to zero only if the two reflected beams have the same amplitude, which happens only for a specific value of n_1.

Antireflection coatings

If you wear eyeglasses, then you want lenses that have a high index of refraction, as this allows thinner lenses and less weight resting on your nose. Diamond would be a great choice, but high-index plastics are certainly more economical. However, a high air/glass Δn means a lot of unwanted reflections, so most modern eyeglass lenses come with *antireflection coatings* to mitigate this problem. The simplest way to make an antireflection (AR) coating is to add a single "$\lambda/4$" coating as shown in Figure 12.

Multilayer AR coatings can yield extremely low surface reflectivities for single-frequency light sources like lasers, or modest reflectivities over broad wavelength bands, as illustrated in Figure 13. What cannot be achieved, however, is extremely low reflectivities over broad wavelength bands. For most broadband applications like eyeglasses, camera lenses, and other optical elements, AR coatings with 0.5-1 percent reflectivities are typical. For demanding laser applications, multilayer dielectric AR coatings can give reflectivities below R = 0.001%, while multilayer dielectric mirror coatings can give reflectivities above R = 99.999% (for a price).

Beamsplitters

Optical beamsplitters are quite common optical elements, used to split a single laser beam into two beams, or perhaps to split light for multiple imaging devices. An ideal 50:50 beamsplitter would simply split a beam into two equal parts with no loss or unwanted reflections, and much effort has been expended in various attempts to get as close as possible to this ideal. One example is the *wedge beamsplitter* illustrated in Figure 14.

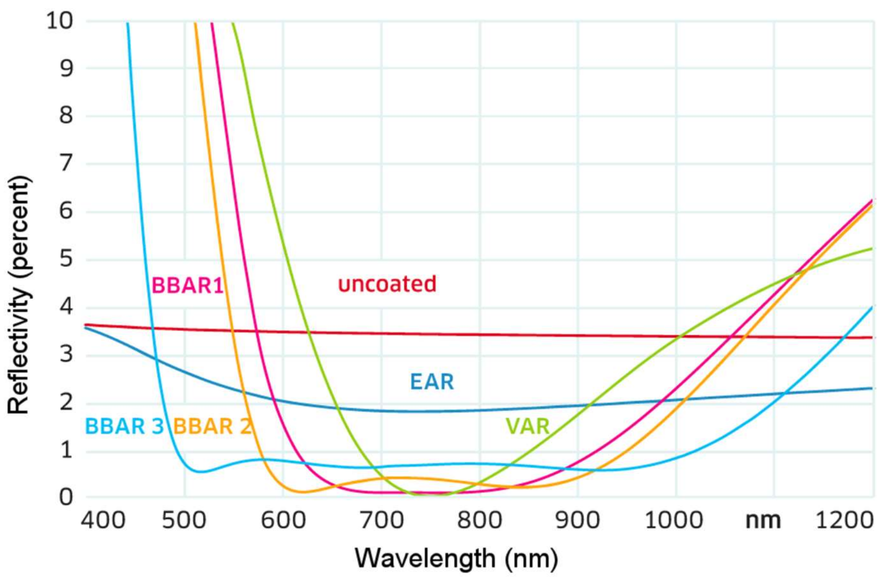

Figure 13. Antireflection coatings can include one or many dielectric layers. The above graph shows the reflection spectrum at normal incidence for:
1) (uncoated) a bare fused silica substrate with n = 1.46,
2) (EAR) a single-layer coating of Magnesium Floride with n = 1.38,
3) (VAR) a double-layer coating optimized to produce zero reflection at 750 nm,
4) (BBAR) several examples of multilayer coatings optimized for good broadband AR performance.
With a finite number of layers, making good broadband coatings requires trade-offs between reflection amplitudes and wavelength coverage. Specialized narrow-band AR coatings (usually associated with laser or atomic transitions) can be made with very low reflection coefficients. Similarly, multilayer dielectric mirrors with extremely high reflection coefficients can fabricated for narrow-band light sources.

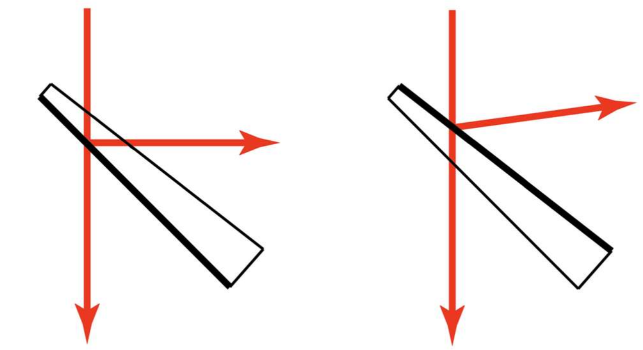

Figure 14. A wedge beamsplitter used in two orientations. The heavy black line indicates a 50:50 dielectric coating, while the other surface has a weakly reflecting AR coating.

Exercise 10. Assume that the main beamsplitter coating in Figure 14 produces a lossless 50:50 split of whatever hits this surface, and assume that all other surfaces reflect 1% and transmit 99% (so R:T = 1:99, again with no loss). Add the weaker reflected beams to the two sketches, ignoring all beams with intensities below 0.02% of the incoming beam. Label the various beam intensities in your sketches, and make sure the outgoing light adds up to 100% of the incoming light (because we are assuming no optical losses).

In the lab, find the wedge beamsplitter on the optical table and verify that your sketches are correct using the Helium-Neon laser. The laser power switch is just below the monitor, and the timer buttons will turn the laser off after the selected time. To help you develop good optical practices, leave the He-Ne laser where it sits and use a mirror (also found on the table) to send the beam in a westerly direction. Keep all laser beams at a constant height (4 inches above the table surface is our standard) and keep the beam paths moving east-west or north-south as much as possible. Following these rules will keep all beams in a single horizontal plane (very convenient!), and many optical elements (like this beamsplitter and mirror) are designed for use with 45-degree incident beams.

From this exercise you can see why the wedge beamsplitter is a popular optical element. If you are not too close to the beamsplitter, then you can easily block the unwanted beams and thereby keep only the desired 50:50 beams. If you used a flat glass plate and the lesser beams were not deflected, then you could introduce a variety of unwanted interference effects that adversely affect whatever scientific goal you are trying to achieve.

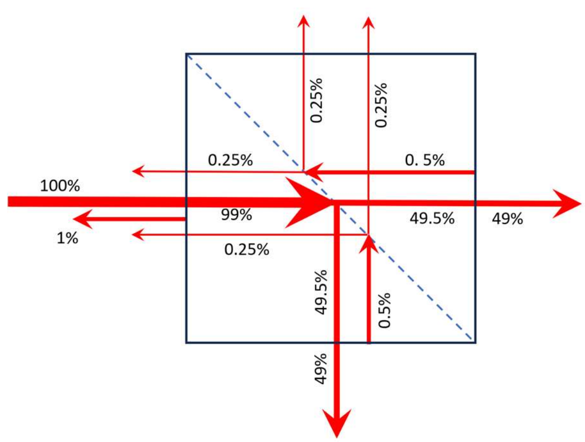

Figure 15. An "exploded" ray tracing through a beamsplitter cube, showing the main 50:50 rays with various weaker rays. In reality, most of these rays would overlap one another along the principal axes. Note that the two main output rays are not accompanied by any unwanted rays with intensities greater than 0.005%.

Another popular alternative is the *beamsplitter cube*, and again you can find an example on the optical table. As illustrated in Figure 15, a multilayer dielectric coating inside the cube splits the beam, while all other surfaces have AR coatings. As you can see in this "exploded" view, the main output beams are not contaminated with unwanted beams (to lowest order), without the need for any wedges. If you place the beamsplitter cube in the He-Ne laser, you will see the unwanted reflections (by rotating the cube a bit), while the main beams remain uncontaminated.

Pellicle beamsplitters are a good option for imaging applications, as these are basically thin-plate beamsplitters with a thickness of only a few microns. Imaging usually uses broad-band light, so interference effects are not so important, while "ghost" images can be problematic with thick-plate beamsplitters. The pellicles are so thin that the images from the front- and back-surface reflections are so close together that they basically merge. The thin membrane is also quite fragile, unfortunately.

Polarization

No introduction to optical physics would be complete without some introduction to polarization, as much attention is given to polarization effects in many scientific studies. You are likely already familiar with linear polarization, and you can find a pair of basic "economy" linear polarizers in the lab. If you place them in the (unpolarized) He-Ne beam, you can observe the expected effects when you rotate the polarizers in their rotatable mounts (try it!).

These two polarizers (Thorlabs *lpvise100-a*) are made using polymer sheets with the long molecules oriented along a particular axis, so the aligned molecules absorb one linear polarization while transmitting the other. These polarizers are restricted to rather low input powers (so the absorption does not damage the polymers), and they are especially inexpensive to make.

Exercise 11. For the Thorlabs *lpvise100-a* beamsplitters, the manufacturer promises an extinction ratio greater than 5000:1 over the range 530-690 nm, which includes the 633 nm He-Ne beam. Specifically, if you send the laser beam through both polarizers, then the intensity ratio of aligned-to-crossed polarizers should be >5000:1. Measure this ratio using the battery-powered photodetector on the table and record your results.

When making your measurement, first do a battery-check on the photodetector and replace the 9V battery if needed (your TA can find you a battery). A voltmeter is available in the lab to measure the photodetector output. The photodetector has a diameter of 8mm, so a laser beam will hit the sensor if the beam goes straight in through the center of the input hole. This input hole also blocks the room lights quite well, but you should not assume that the background signal is always negligible. Block the laser beam occasionally to see if room lights make a significant contribution. Finally, remember that this detector is powered by a single 9V battery, so the output voltage cannot go above this value. To avoid saturation issues, *keep the output voltage below 5V*. Fortunately, there are three gain settings, and you can assume that the gain ratios are accurately listed.

When your optical layout looks good and is locked down, measure the beam intensity with the polarizers aligned and with them crossed. With care, you should be able to measure an intensity ratio above 5000. The difficult bit is adjusting the rotation angle to minimize the transmitted beam. The trick is to tweak the rotation while watching the voltmeter (with the optics bolted down).

You may find several other types of polarizers in research labs, chosen for their specific properties. Some are better suited for IR, visible, and UV applications, some can tolerate high laser powers, and some have remarkably high extinction ratios. Many polarizers are made from birefringent crystals (such as calcite) that split laser light into two beams that propagate in slightly different directions. With names like Wollaston prisms, Glan-Taylor polarizers, and Glan-Thompson polarizers, these can provide extinction ratios greater than 100,000:1. As you might expect, these high-performance polarizers are quite expensive, so we will not be investigating them in this lab. As you just witnessed, however, the "economy" polymer polarizers are quite good.

Polarizing beamsplitter cubes

Another type of linear polarizer is based on reflection at Brewster's angle, which is illustrated in Figure 16. This angle is unique in that the amplitude of the reflected plunging polarization goes to zero.

Exercise 12. Using Snell's law and the Fresnel equations, show that Brewster's angle is equal to

$$\theta_B = \arctan\left(\frac{n_2}{n_1}\right) \tag{13}$$

At Brewster's angle, the electric field vector of the reflected plunging beam would be along the ray propagation direction, and this is not possible for electromagnetic radiation (which provides a simple physical explanation for why $R_p \to 0$ at Brewster's angle).

Although $R_p \to 0$ at Brewster's angle, R_s remains less than unity, and this means that the refracted ray contains both polarizations. However, one can use some multilayer coating trickery to enhance the S-polarization reflection, thus yielding a *polarizing beamsplitter cube* illustrated in Figure 17. You can find an example in the lab (Thorlabs ccm1-pbs251) but note that the extinction ratios are relatively modest: the transmitted beam extinction ratio is specified at >1000:1, while the reflected beam extinction ratio is only specified to be >100:1.

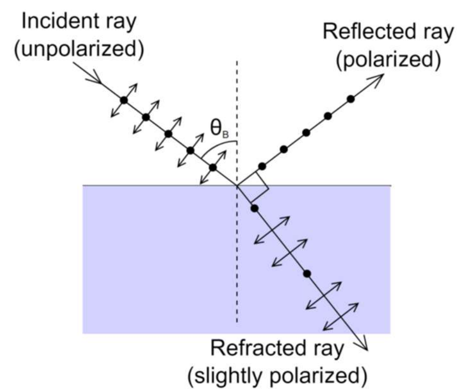

Figure 16. An illustration of ray reflections at Brewster's angle.

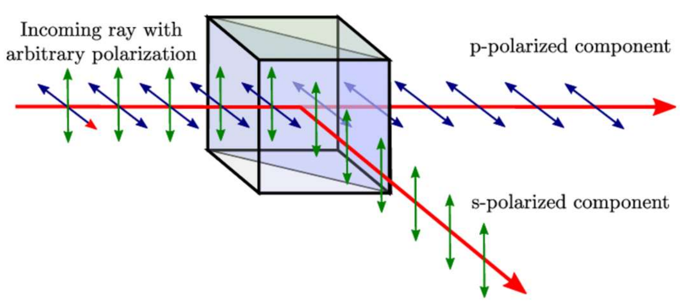

Figure 17. A polarizing beamsplitter cube separates the s- and p-polarization components of an incoming light beam.

Exercise 13. Measure these two extinction ratios using your He-Ne laser, like you did above, using one of the *lpvise100-a* polarizers to provide a polarized input beam. (Make sure you use the polarizing beamsplitter cube, not the unpolarizing cube.) You should find that the transmitted beam meets spec easily, while the reflected beam might not. [Turns out you need to rotate the cube a few degrees (about the vertical axis, so the reflected beam is not quite at 90 degrees relative to the input beam) to optimize the reflected extinction ratio. The easiest way to do this is first minimize the reflected intensity by rotating the economy polarizer, then rotate the cube while observing the reflected beam by eye to obtain a lower minimum. For best results, iterate these steps. The amount of rotation necessary is likely a wavelength-dependent number, with only one wavelength giving you a minimum reflection at 90 degrees.]

One take-away message from this exercise is that not all polarizers are created equal, and specifications matter. Another take-away is that manufacturers tend to list best-case specifications, and sometimes you had better make your own measurements if you want the device to meet spec.

Next locate the half-wave plate retarder on the optical table and send the He-Ne beam through the economy polarizer → half-wave plate → polarizing beamsplitter cube. Observe what happens when you rotate the half-wave plate. This arrangement can be used as a *variable beamsplitter*, using the half-wave plate rotation to tune the relative intensities of the reflected and transmitted beams. Observe the two beam intensities as a function of the rotation angle of the half-wave plate. The best way to explain this behavior is with the *Jones calculus*, which lets you calculate the cumulative effects of polarizing optical elements….

Jones calculus

Begin by representing the polarization of a plane wave propagating in the z direction by an electric-field vector (called the Jones vector)

$$\vec{E} = \begin{bmatrix} \hat{E}_x \\ \hat{E}_y \end{bmatrix} = \begin{bmatrix} E_x e^{i\varphi_x} \\ E_y e^{i\varphi_y} \end{bmatrix} \tag{14}$$

where E_x and φ_x represent the amplitude and phase of the electric field at some point in space and time (dropping the usual e^{ikz} and $e^{i\omega t}$ factors associated with traveling waves, for clarity). For example, the vectors

$$|H\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \qquad |V\rangle = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \qquad |L\rangle = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ i \end{bmatrix}, \qquad |R\rangle = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -i \end{bmatrix} \tag{15}$$

represent: 1) light polarized along the x axis (horizontal polarization), 2) light polarized along the y axis (vertical polarization), 3) left-circular-polarized light, and 4) right-circular-polarized light. States with arbitrary $\hat{E}_x$ and $\hat{E}_y$ are often called "elliptically" polarized light. Unpolarized light (like sunlight, where the electric field vector varies randomly over space and time) is not well represented using the Jones formalism.

Polarizing optical elements are then represented by matrices that act on Jones vectors. For example, horizonal and vertical polarizers are expressed as

$$M_H = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \qquad M_V = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \tag{16}$$

and you can verify that multiplying a matrix and a vector (for example $M_H|H\rangle$ or $M_H|V\rangle$) gives the expected results.

To see how useful this formalism can be, let us apply it to the case of the half-wave plate above. This optical element introduces a half-wave retardation along one polarization axis, so it can be represented by (think about it)

$$M_{Half} = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \tag{17}$$

If we rotate this optical element (about the z axis) by an angle θ, then it can be represented by

$$M_{Half,rotated} = R(-\theta) M_{Half} R(\theta) \tag{18}$$

where

$$R(\theta) = \begin{bmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{bmatrix} \quad (19)$$

This rotation action applies to any Jones matrix, and $R(\theta)$ can be applied to Jones vectors as well. Of course, there is much additional information online regarding the Jones calculus and how it is used in polarized-light calculations.

Exercise 14. Using the Jones calculus, show that a rotating half-wave-plate can rotate the axis of linear polarization. How does the polarization angle vary with the rotation angle of the half-wave plate? Compare your calculations with what you observed in the lab.

White-light interferometry

The topic of dielectric coatings suggests another application of optical interference – making absolute thickness measurements of thin dielectric films. You may have noticed, for example, that soap bubbles often show colorful interference effects, and it is straightforward to quantify this phenomenon and turn it into a useful method for making non-contact measurements. While you may never encounter this technique in your career, it serves as a useful demonstration of several general optical techniques.

The basic idea is illustrated in Figure 18, where a beam of light passes through a dielectric film. Reflections occur at the front and back surfaces of the film, and these two reflections will interfere with one another if they overlap. (For clarity, the figure shows a large angle of incidence; with normal incidence, the rays would overlap.) If we use a broadband light source (often called "white" light), then some wavelengths would constructively interfere, while other wavelengths would destructively interfere. And we can use this spectral information to measure the thickness of the film.

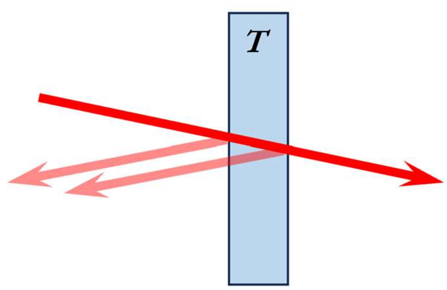

Figure 18. When a beam of light passes through a dielectric film, the front and back surfaces produce reflections of equal magnitude (why?). If these beams overlap and interfere, the spectral information in the reflections can be used to measure the thickness of the film. (For clarity, the figure shows a large angle of incidence; with normal incidence, the rays would overlap.)

Exercise 15. Assume that the incoming beam in Figure 18 strikes the thin film at normal incidence, so the two reflecting beams overlap completely. Also assume equal reflection coefficients $R \ll 1$ for both the front and back interfaces. Show that the intensity of the reflected beam is

$$I = 2I_0R\left[1 - cos\left(\frac{4\pi T}{\lambda}\right)\right]$$

where I_0 is the incident intensity, T is the film thickness, and λ is wavelength.

To make this measurement strategy work, we need to measure the reflected light as a function of wavelength, and the best way to do this is by using a diffraction grating to disperse the colors, as illustrated in Figure 19. Some variant of this geometry is typically used with diffraction gratings, having these notable characteristics: 1) light first passes through a narrow slit that is imaged onto the camera, so the resolution of the spectrograph is limited by the width of this slit. 2) A grating must be illuminated with a broad beam of collimated light, because the resolution of the spectrograph is also limited by the number of "grooves" in the grating that are being used.

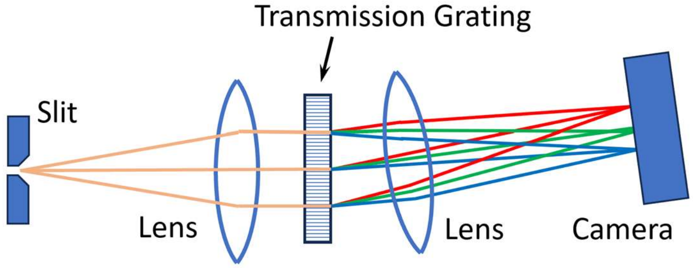

Figure 19. One way to measure the spectral information in white light is to disperse the light using a diffraction grating (in this case a transmission grating). The first lens provides a wide, collimated beam for the grating, while the second lens focuses the slit onto the camera.

The optical layout in Figure 19 is commonly used for grating spectrometers, but it will not work for measuring the reflected light from a thin film. To that end we must use the more complex optical setup illustrated in Figure 20. Here the slit is first imaged onto the thin film at P2, and the reflected light then makes it back through the grating to be dispersed, and the slit is further imaged onto the camera. Convince yourself that 1) this geometry will work for this purpose, and 2) there is no simpler way to accomplish this task.

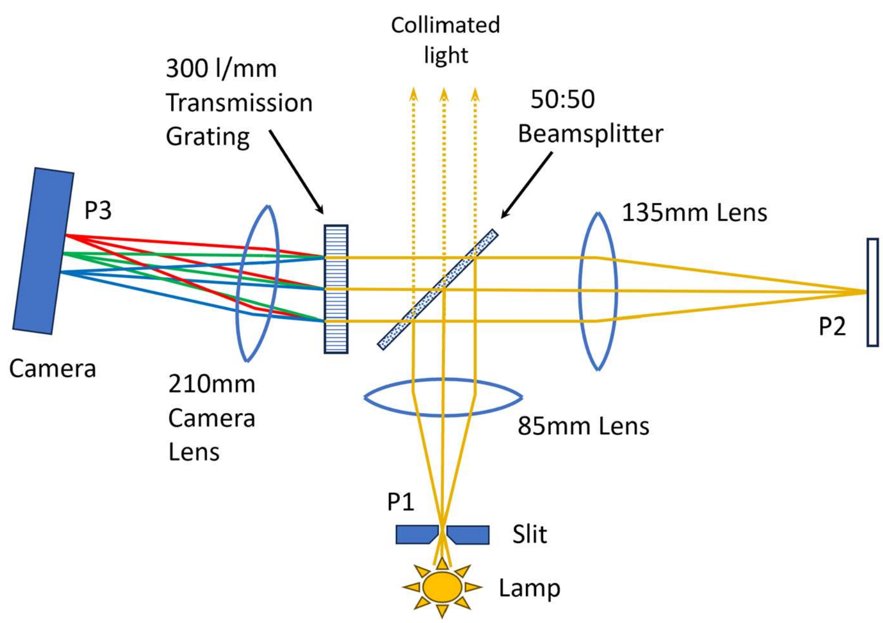

Figure 20. The optical layout for measuring the thickness of a thin film using white-light interferometry.

Color dispersion is another widely used element of optical technology, and diffraction gratings are a common tool for dispersing light. The basic phenomenon is described by the *grating equation*, which gives

$$m\lambda = d[\sin(\alpha) + \sin(\beta)] \tag{20}$$

where m is the diffraction *order*, λ is light wavelength, d is the spacing between adjacent lines (a.k.a grooves) of the grating, α is the angle of incidence of the input light (with normal incidence having $\alpha = 0$), and β is the angle of the diffracted light. Of course, you can easily look up the derivation of this expression online. An often-encountered special case is normal incidence with $m = 1$, giving

$$\sin(\beta) = \frac{\lambda}{d} \tag{21}$$

To see all this in action, your next exercise is to set up the white-light interferometry demonstration illustrated in Figure 20. Walking through this sketch, note that:

1) The Lamp consists of an incandescent light bulb that produces a smooth blackbody spectrum. LED bulbs run a lot cooler, but even "white" LEDs put out a complex optical spectrum that tends to confuse the signal you will be looking for.
2) The lamp light next passes through an adjustable slit (P1) to produce a thin "line" of light.
3) The 85mm lens "collimates" the light from the slit – diverging rays enter the lens from the slit, and parallel rays exit the lens.
4) The beamsplitter sends half the light through the 135mm lens that focuses the light at P2. Thus the line of light at P1 is reimaged at P2.
5) The mirror at P2 sends the light back through the beamsplitter and then through a transmission diffraction grating (the beamsplitter also sends some light back toward the lamp, which we can ignore).
6) The transparent grating has something like a "sawtooth" surface structure on one side, essentially an array of linear wedges spaced at 300 lines per millimeter. This structure does two things: 1) like a set of prisms, refraction diverts the rays to one side (upward in the sketch); 2) like a set of slits, diffraction disperses the colors in the light, yielding a spectrum on the camera. (The dispersion is illustrated in the Figure by yellow rays turning into RGB rays.) The diffraction angle as a function of wavelength is given by Equation (21).
7) [Interesting aside about gratings. If the grating were made simply as an array of tiny slits, then shining a laser on the grating would produce three rays: a central ray "diffracted" into the forward direction with $\beta = 0$ and two rays diffracted into the $\beta = \pm\beta_{diff}$ directions. These are the $m = 0$ and $m = \pm 1$ diffraction orders. By using a sawtooth surface structure, the narrow prisms divert most of the light to the desired $\beta = +\beta_{diff}$ direction, which increases the efficiency of the grating. The sawtooth pattern is called the "blaze" of the grating. Don't be confused into thinking that the narrow "prisms" disperse light into colors. This prismatic dispersion is tiny in comparison to the grating dispersion. The prisms here mainly change the direction of the light via refraction, like a lens.]
8) Note that the slit at P1 is reimaged at P2 and then again on the camera sensor at P3. Making the slit thinner increases the spectral resolution of this spectrometer.
9) Replacing the mirror at P2 with a thin film, light then reflects off both the front and back surfaces of the film. These two reflections interfere with one another, and you should see this clearly in the spectral image. Moreover, you can use the spectral information to measure the thickness of the film without making any direct physical contact with the film.

Setting all this up in the lab is helpful for understanding the physics, plus the exercise gives you some experience working with real optics in a laboratory setting. As is often the case, getting good results often comes down to aligning the optics well, so here are some steps to help guide you toward developing good optics habits:

1) When working on an optical table, you usually want to confine the layout to a single plane, just as Figure 20 shows the main light rays in the plane of the page. To this end, some of the optical elements you will be using in this exercise have already been set up with a 4-inch height; others you may have to adjust.
2) If you take a close look at the lamp/slit assembly (before you place it on the table), you can see an adjustable slit in front of a small incandescent bulb. Send 2-3 volts to the lamp and you will see the filament (a wire coil) glowing dimly. Please *be gentle* when you adjust the slit; it should take very little torque to turn the micrometer. But if you apply too much torque, you can break it; it is a delicate (and surprisingly expensive) piece of equipment. You can also see a small wire glued to the slit; this is a focusing aid, as you will see later. (Please do not touch the wire; that tiny dab of glue is not holding it on very securely.)
3) Another good alignment practice is to keep the various ray paths on right angles (mostly), using the holes on the optical table to guide your placements.
4) To begin your setup, place just the lamp/slit assembly, the first lens, and the beamsplitter cube on the table. The lens should be close to the cube, with the back of the lens (the side normally connected to a camera) facing the slit. You want this lens orientation because that is how the lens was designed to work – camera lenses normally take collimated light in the front and focus it onto a camera sensor; this setup is the same but in reverse. Note that the beamsplitter cube includes a switchable magnetic base; when the lever is set to ON, the base grips the table quite firmly. The lamp and lens are on weaker magnetic bases; the magnets are useful for making rough placements and adjustment, and then you can clamp them in place when you are ready.
5) To aid with the alignment, place a white card at P2 in the diagram (roughly) and move the slit forward and backward. Send 8V to the lamp, and you should be able to focus the lamp filament (a wire coil) on the white card. At a different distance you can focus the slit on the card. The small wire glued to the slit puts a dark gap in the slit of light, indicating that you are indeed focusing on the slit.
6) Now move the slit so the light after the beamsplitter is collimated – neither converging nor diverging. Move the card back-and-forth to observe the collimation, and move the slit position back-and-forth so the light is collimated as best you can. Adjust these three elements until you achieve a good alignment. Your goal is to make several things happen simultaneously: a) the pieces should look aligned relative to the holes in the table, not crooked; b) you should see that the light is indeed collimated after the beamsplitter, neither diverging nor converging (move a white card around to check this); c) the height of the collimated beam should be at the 4-inch standard. It may take a few iterations (and some patience) to get everything looking good. When you are satisfied, bolt everything down, and you should not have to move these elements again. (However, you may need to redo this step later; rest assured that it goes much faster after you have done it a few times.)
7) Continue by adding the 135mm lens and the mirror, again placing the front of the lens near the beamsplitter cube. Check that the light beam actually hits the mirror, and that the mirror sends the beam back through the beamsplitter cube. Place a white card about 12 inches after the cube and move the mirror around until you see an image of the slit on the card (with the gap from the wire visible). You will probably have to adjust the mirror left/right pointing to get a nice bright image, and adjust the up/down pointing to keep the beam horizontal. Then move the mirror so the light exiting the beamsplitter cube is collimated and place the grating. Use a white card to see that the dispersed beam is deflected to the right, as shown in Figure 20.

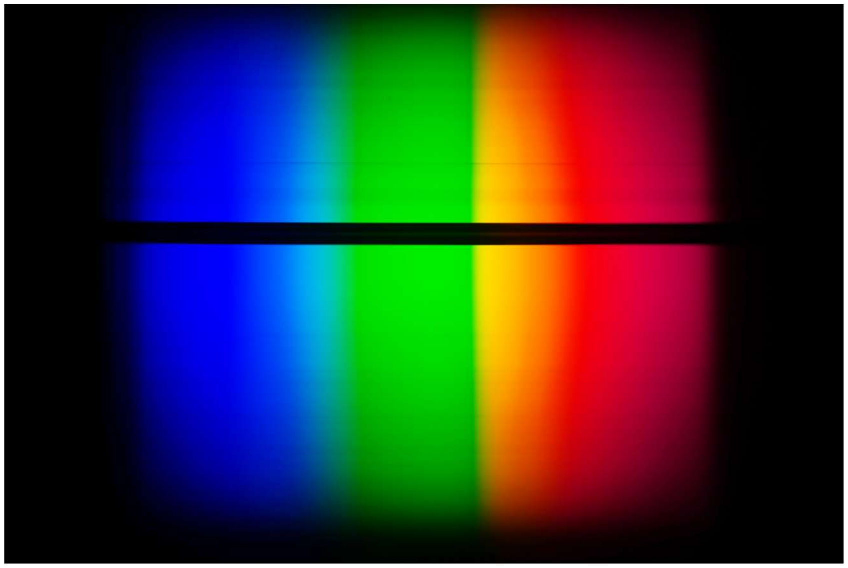

Figure 21. An uncropped screen shot showing the optical spectrum produced by the apparatus in Figure 20. The central line is caused by a wire blocking a small section of the slit.

8) With the 55-210mm lens attached to the camera, manually set the lens focal length to 210mm, and set the focal distance to infinity (as indicated on the camera display). Use aperture-priority mode (A) and set the aperture to f/6.3. Place the camera right behind the grating and look for the spectrum. Adjust the mirror position so the wire gap is in focus in the camera image. If needed, set the ISO to give you a reasonable shutter time. To get a nice spectral image, you may need to adjust the camera position, camera pointing, mirror position, and mirror pointing. Also, the grating is in a rotatable mount. It's pretty easy to get an okay spectrum, but you can usually make it better if you play around with the alignment. In laboratory optics, you always have to deal with alignment issues… it's not like just buying something!
9) When everything looks good, bolt it all down to produce a stable spectrometer, ready to make some measurements. (But again, don't be afraid to tweak and/or rebuild the setup if needed. In any experimental field, things rarely work well on the first try. That is why you are practicing now with just a few optical elements… to get some useful experience.)

Exercise 16. Capture a nice image of the bare blackbody spectrum and add it to your e-notebook, which should look something like that in Figure 21. Note that the colors in this image are not entirely accurate. First the camera interprets (in software) the camera sensor data, consisting of pixels covered with an interleaved array of microscopic color filters. Subsequent image processing distorts the colors further, and yet more color distortion happens in your computer display. The basic R/G/B segments come out okay, but the intermediate color structures seem to include some digital artifacts.

Exercise 17. Locate the red, green, and blue laser pointers on the table, and use these to add known spectral lines to the spectrum. To fire up each laser in succession, first fit the laser in the mount so the button is squeezed on. Then apply 3VDC (no higher!) using the other DC power supply and clip leads (the center terminal on the laser is negative). Next point the laser beam through the lamp using the back hole in the lamp assembly. Some of the beam will pass through the wire coil and add a line to your spectrum. Close the slit down so this line is fairly sharp (but again, be gentle when closing the slit). Figure 22 shows a composite image made three separate images – one for each laser. You can

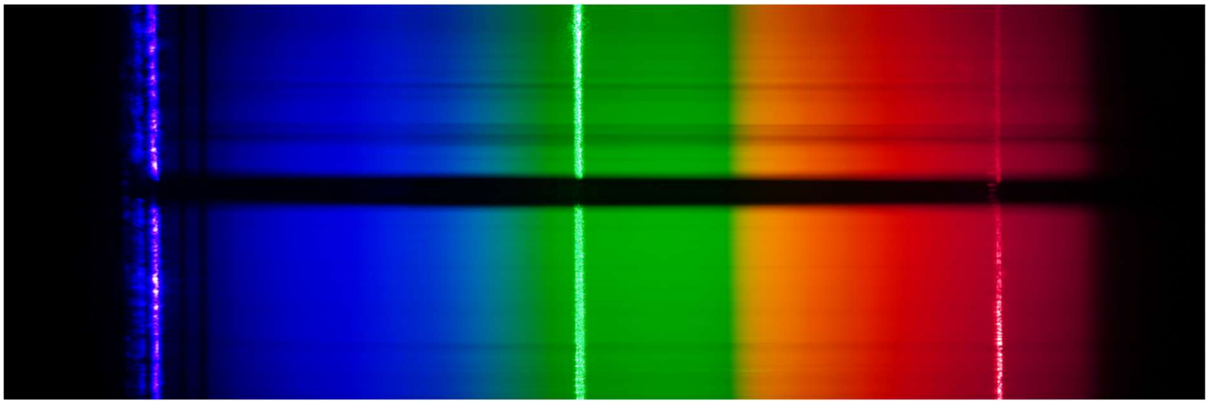

Figure 22. A composite image showing the blue, green, and red laser lines superimposed on the lamp spectrum at wavelengths of 405, 532, and 650 nm.

simply add the three separate images to your e-notebook, but make sure you do not change the optical alignment of the spectrometer as you change lasers. If you bump something that moves the spectrum on the camera, you just have to start over.

Exercise 18. The three lasers have wavelengths of 405, 532, and 650 nm, and you can assume that each of these numbers is accurate to ± 1 nm. Measure the line positions on your images and plot wavelength as a function of x-pixel number on the images. Next model this data using the grating equation, camera properties, and other features of your setup. Your model should include only one adjustable parameter – the wavelength of the center x-pixel on your images. Plot your model line together with your three data points. Using this setup, what do you measure for $\partial\lambda/\partial x$?

Exercise 19. Calculate $\partial\lambda/\partial x$ from the grating equation. You will need to know the grating lines/mm, the camera lens focal length, and the size of the camera pixels (all available in this handout). If your answer is not the same as above (within uncertainties), then obviously something is wrong.

Exercise 20. Without otherwise changing the spectrometer alignment, replace the mirror in Figure 20 with the pellicle beamsplitter in a "cube" mount. Each pellicle surface reflects about 4% of the incident light, and the two reflections interfere to produce a series of dark bands on the spectrum. Include an image of your new spectrum in your e-notebook, and use it to calculate the thickness of the pellicle from the spectrum. The manufacturer specifies that the pellicle material has a thickness of 2 μm and an index of refraction of $n = 1.5$ at 550 nm (green in the spectrum). This exercise demonstrates how you can use interferometric effects to measure the thickness of a transparent material to high accuracy without contacting the object being measured.

Exercise 21. Replace the pellicle beamsplitter with the cellophane sheet mounted in a second cube mount. The cellophane is much thicker than the pellicle, plus the surface quality is not great. So you will have to move the sheet around to find a good spot, and the dark bands will be closely spaced and considerably fainter than you saw with the pellicle beamsplitter. Nevertheless, move the cellophane around until you get a faint spectrum you can use, and use your image to measure the thickness of the cellophane film.

Appendix I: An Abridged User Guide to the Sony A6400 Camera

These images show you the names and locations of various buttons, dials, and wheels:

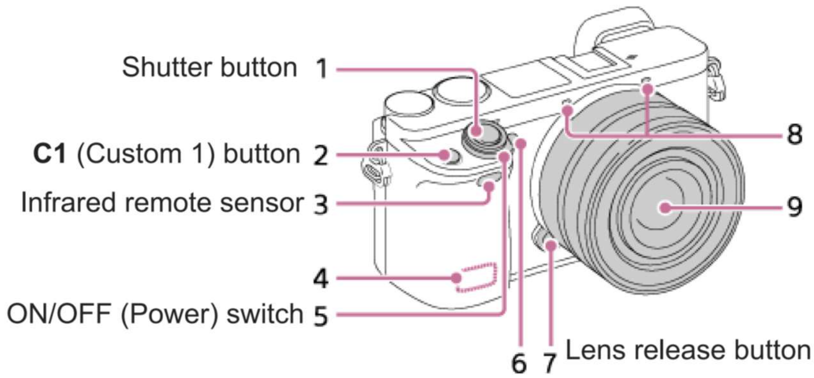

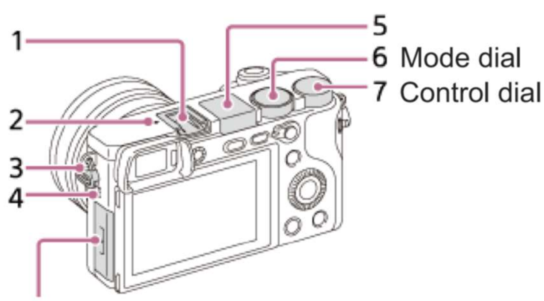

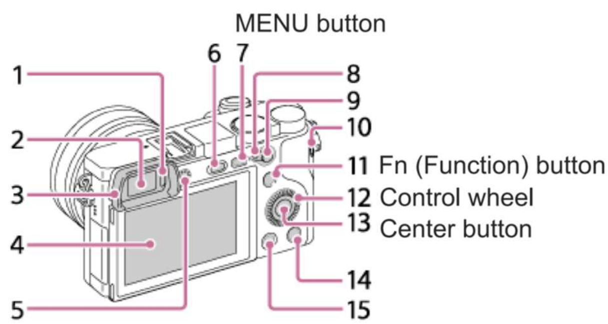

This camera has an APS-C size CMOS sensor with a physical size of 23.5mm × 15.6mm, and it saves images with 6000 × 4000 pixels, giving a **pixel size** of 3.92 μm (square).

What follows is a brief introduction to the various settings you will need in the Ph77 Optics Track, along with some general camera terminology. The full User Guide for this camera is 659 pages long, so this Appendix describes only a small subset of what this camera can do.

Manual and AF Lenses. A manual lens has no electronic connection to the camera – focus and aperture must be set by hand. With an AF (auto-focus) lens, the focus and aperture can be changed electronically by the camera. If you set the camera to Manual Mode, then an AF lens behaves like a manual lens.

Camera Modes

These are set using the Mode Dial on top of the camera, and the Mode sets the overall behavior of the camera during operation.

Auto Mode. For general photography, setting the Mode Dial to AUTO gives acceptable results. Then you have a "point-and-shoot" camera, where the camera decides what settings are best. Many of the camera's user adjustments are not available in this mode. Unfortunately, auto mode can become quite confused in scientific situations.

Program Auto (P) Mode. Close to full Auto, but you can set the shutter speed or aperture, while the camera keeps an appropriate exposure level (which depends on a combination of shutter speed, aperture, and ISO number). This mode is useful in portrait photography, but we will not use it much in Ph77.

Aperture Priority (A) Mode. Here the user sets the aperture (F number) while the camera does the rest. Unlike Auto Mode, some user adjustments are allowed.

Shutter Priority (S) Mode. Here the user sets the shutter speed while the camera does the rest. Like Aperture Mode, some user adjustments are allowed.

Manual (M) Mode. No automatic adjustments. Everything must be done manually.

Some Important Settings

Pressing the Menu button reveals a plethora of camera settings. Please do not change most of these unless you take notes on what you are changing. And then be sure to change the settings back when you are finished. Randomly changing settings can cause quite a bit of confusion down the road. You are welcome to change all of the settings listed here without hesitation, and you need not set them back to their initial state. These settings are all you need to do the lab exercises.

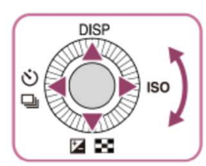

Display: Press the top (DISP) side of the Control Wheel (not the same as the Control Dial) to toggle through the display modes.

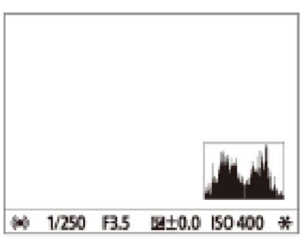

It's best to keep the display relatively uncluttered, and the image on the left works well in practice. Note that the shutter speed, f-number, exposure parameter, and ISO number are all visible, along with a histogram of the light intensity. With a manual lens, the f-number is blank. In manual mode, you see a "metering mode (M.M.)" that helps you adjust the exposure level.

ISO: Stands for International Organization for Standardization. Historically, the ISO number indicated the sensitivity of photographic film, which was a fixed number for each specific film. In modern cameras, the ISO number represents the gain of the amplifier reading out the camera sensor, which is adjustable. A high ISO number is useful for low-light situations or fast shutter speeds, but the amplifier noise may give a grainy image at high ISO. A low ISO number gives an image with less noise, but more light is required (usually by using a slower shutter speed).

The ISO number is linear in light intensity. For example, an ISO number of 100 with a shutter speed of 1/1000th second will give essentially the same image as an ISO number of 200 with a shutter speed of 1/2000th second.

With the Sony A6400 camera, the ISO setting ranges from 100 to 102400, giving an impressive dynamic range of 1000x. This allows you to take photos in both low-light and high-light conditions.

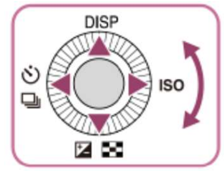

To change the ISO, press the ISO side of the control wheel, then press the top/bottom sides of the control wheel to adjust, then press the center button to lock it in.

Shutter Speed: This can be adjusted simply by rotating the control wheel (in most modes).
With the Sony A6400 camera, the shutter time ranges from 30 seconds to 1/4000th second, again providing a lot of dynamic range. With shutter times longer than 1/100th second, image shake can be a problem when pressing the main shutter button. Alternatively, you can use an IR remote shutter without touching the camera. (The IR sensor is located on the front of the camera, as seen in the sketch above.) With a low ISO number and a fast shutter speed, you can capture images with exceptionally high light levels (sometimes useful when using lasers).

Focus Magnification: Press the C1 button to get a magnified view of the live screen, which you can use when manually focusing. You can move the magnified region using the side buttons on the control wheel. Keep pressing C1 to cycle through the different magnification levels.

Exposure Compensation: Press the bottom side of the Control Wheel to change the exposure setting relative to what the camera software decides is the likely best image exposure.

White Balance: If you press the **Fn** button, you can see a quick table of other settings, including the White Balance. This usually works well in Auto, but you may want to adjust this if the colors look off.

Advanced Settings

Set and forget. Do not change these
Image (1-2) = Jpeg, Fin, L, 3:2
Focus mode (1-5) = Manual Focus
Focus magnification time (1-13) = No limit
File Format (2-1) XAVC S HD
Custom Button 1 (2-8-3) = Focus Magnifier
Release w/o Lens, Card (2-4) = Enable
Zoom setting (2-5) = Optical zoom only
Function menu set (2-8)