Forwards or Backwards

https://forwardorbackwards.com



Learning Photography Course – 5. Lenses, and How Focus Works

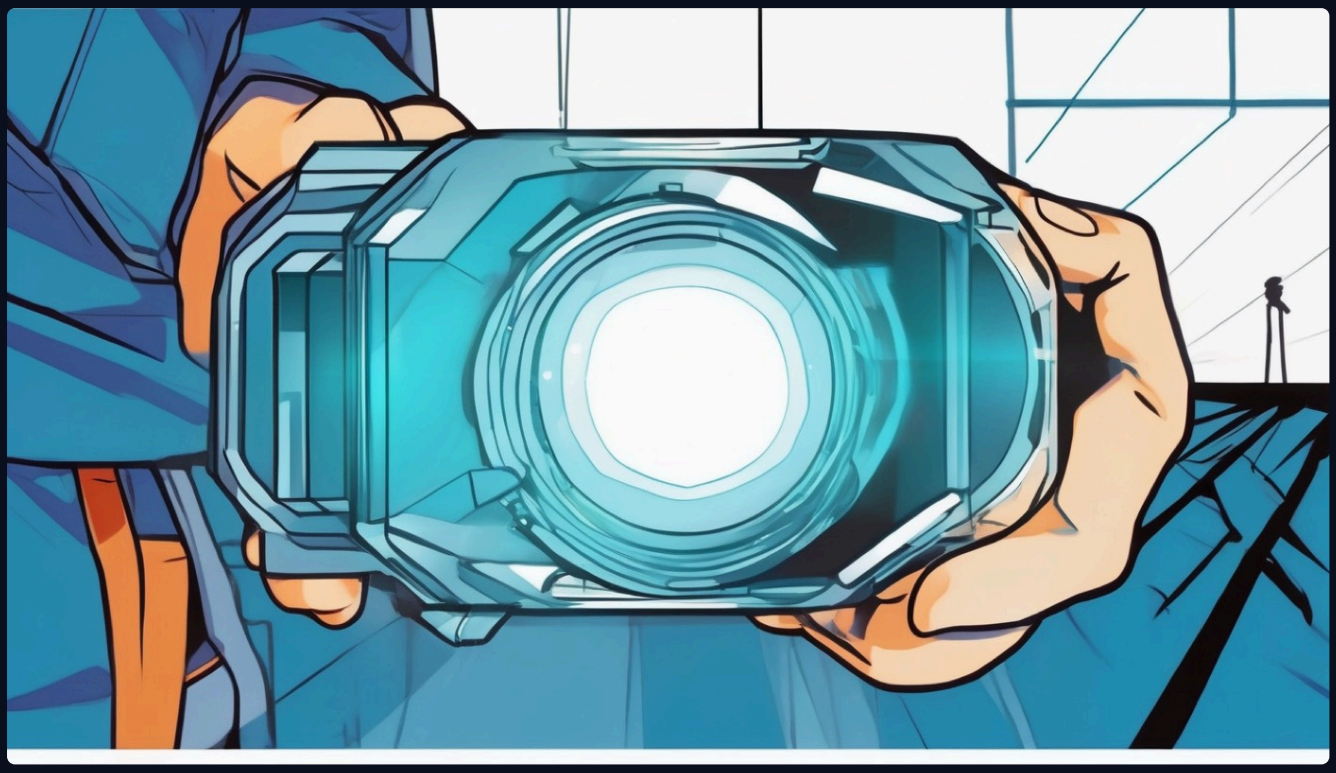
Watch on YouTube: https://www.youtube.com/watch?v=I34nbf-6AX4

Chapters

1. 1 What a lens is doing
2. 2 Focal length and field of view
3. 3 Prime against zoom
4. 4 Maximum aperture, and what it costs
5. 5 Elements, groups and coatings
6. 6 Aberrations
7. 7 Stabilisation
8. 8 Special lenses

9. 9 How autofocus works
10. 10 Focusing in practice
11. 11 Depth of field, bokeh and compression

1. 1 What a lens is doing



Module one said that a photograph is a measurement of light over an interval, and that the apparatus a measurement like that requires is a hole, a shutter and a recording surface. That is true, and it is not quite enough, because a hole on its own makes a very poor picture.

Module one showed why. A pinhole makes an image because it restricts which rays from the scene can reach any given point on the surface — and the smaller you make it, the more sharply it restricts, and the less light gets through. That is the first trade in photography and the pinhole cannot escape it. Make the hole small enough to be sharp and the exposure runs into minutes. Open it up to shorten the exposure and the picture turns to mush.

A lens is the way out of that trade, and the way it escapes is worth understanding properly rather than taking on faith.

Light bends when it passes between materials of different density. Glass is denser than air, so a ray entering a curved piece of glass changes direction, and how much it changes depends on the angle at which it arrives. Shape the glass correctly and all the rays leaving a single point on your subject — the ones spreading out in every direction, that a pinhole would have had to throw away — are bent back so that they meet again at a single point on the other side.

That is the whole idea. The pinhole discards almost all the light to get sharpness. A lens **collects** the light and puts it back together.

And because it puts the rays back together at a point, the image it forms is a real thing in space, not a metaphor. Put a piece of tracing paper where those rays converge and you will see the picture on it, upside down, in the air, before any sensor is involved. That is what "forming an image" means, and it is why a large format photographer focuses on ground glass: they are looking at the actual image, where it actually is.

Which leads to the first thing on a specification sheet that is worth knowing about and almost never explained.

A lens does not project a rectangle. It projects a circle.

The glass is round, the image it forms is round, and somewhere inside that circle a rectangular sensor is cut out of it. The circle is called the image circle, and it has to be large enough to cover the sensor corner to corner or the corners of your photograph are simply dark.

This explains several things that otherwise look arbitrary. It explains why a lens made for a smaller format may physically mount on a larger-format body and then vignette horribly – the circle is too small and the corners fall outside it. It explains why those lenses are smaller and lighter: less glass is needed to throw a smaller circle. And it explains why a tilt-shift lens is so large, because it throws a deliberately oversized circle so the sensor can be moved around inside it, which is a trick chapter 5.8 comes back to.

It also explains why medium format lenses are enormous. Bigger surface, bigger circle, more glass.

Now, focal length, which is the number on the front of every lens and is more often described than defined.

The focal length of a lens is the distance from its rear principal point to the sensor when the lens is focused at infinity. That is the definition, and every part of it is doing work.

Focused at infinity matters because when you focus closer, the glass moves away from the sensor, so the actual distance changes. The number on the barrel is the infinity case.

Rear principal point matters because it is not the front of the lens, or the middle, or the mount. It is an optical location worked out from the design, and in a complicated modern lens it may not lie inside the physical lens at all. That sounds like a technicality and it is the reason a 300mm telephoto is not 300 millimetres long, and a 24mm wide-angle does not sit 24 millimetres from the sensor. The designers move that point where they need it.

So: focal length is a distance, measured in millimetres, and it is a property of the lens and nothing else. It does not change when you put the lens on a different camera. Module three was firm about this and it is worth repeating, because "this 50mm becomes a 75mm on APS-C" is the single most repeated wrong sentence in photography. The lens is a 50mm lens. The smaller sensor crops into its

image circle, so the FIELD OF VIEW matches what a 75mm would have given on full frame. The focal length did not move. The frame got smaller.

Finally, what is actually inside.

A modern lens is not one piece of glass. It is a dozen or more, in groups, and they are there because a single curved piece of glass makes a bad picture in several specific and predictable ways.

At the front, elements gathering light over a wide area. In the middle, an iris — the adjustable hole from module two, made of overlapping blades. Behind or among them, a group that moves to focus, and in a zoom, another group that moves to change the focal length. At the back, elements doing the final correction before the light lands.

Each of those groups exists to fix a particular fault, and chapter 5.6 is the list of faults. For now the useful picture is this: a lens is not a window. It is a stack of deliberate compromises, each one cancelling an error introduced by the one before it.

2. 2 Focal length and field of view



Focal length is a distance. Field of view is an angle. They are related, and keeping them separate in your head is most of the battle.

Here is the relationship, and it is worth seeing rather than being told. The sensor has a fixed size. The lens forms its image at the focal distance behind the principal point. Draw a line from the edge of the sensor, through the lens, out into the world, and do the same from the other edge: the angle between those two lines is the field of view. Move the lens further from the sensor – a longer focal length – and the angle narrows. Bring it closer – a shorter focal length – and the angle widens.

That is all it is. A long lens is a narrow angle. A wide lens is a wide angle. Everything else follows.

The conventional divisions are worth knowing, though they are conventions rather than physics.

A wide-angle lens takes in more than the eye takes in comfortably. It exaggerates the difference in size between near and far things, and it is the lens that makes small rooms photographable and large landscapes look smaller than they were.

A normal lens approximates what the eye does. Its perspective looks unremarkable, which is either its great strength or its great dullness depending on who you ask.

A telephoto lens takes in less. It isolates. It makes distant things large and stacks them against each other in a way that looks unlike ordinary seeing.

Now, why "normal" is around fifty millimetres on full frame, because the usual explanation — that it matches the human eye — is not really right.

The convention is that a normal lens has a focal length roughly equal to the **diagonal of the format**. Full frame is 36 by 24 millimetres, and the diagonal of that rectangle is about 43 millimetres. The classic normal lens is a 50, which is a bit longer than the diagonal, and that is largely historical.

Do the same arithmetic on any other format and you get its own normal. A smaller sensor has a shorter diagonal, so its normal lens is shorter. Medium format has a longer diagonal, so its normal lens is longer — an 80mm on some medium format systems is a normal lens, not a portrait telephoto, which surprises people coming from 35mm.

The reason the diagonal is the right basis is that it gives roughly the same **angle** on any format, and the angle is what the viewer sees. The number is a means; the angle is the end.

Which brings us back to crop factor, from a different direction than module three took.

Module three treated sensor size as three separate questions and warned that merging them is what makes the argument heated. Field of view is the first and easiest of the three, and it is pure geometry: a smaller sensor cuts a smaller rectangle out of the same image circle, so it sees a narrower angle, and the ratio of the diagonals tells you how much narrower.

The honest way to say it, and the way this course will keep saying it, is: *this 50mm lens, on this camera, frames like a 75mm would on full frame*. That sentence is true, it is useful in a shop, and it does not contain the claim that the lens changed.

And finally the thing that is most often got backwards, which is what focal length does to perspective.

Here is the claim you will hear: wide lenses distort perspective and telephotos compress it.

Here is what is actually true: **focal length does nothing to perspective at all.** Perspective — the relative size of near and far things — is determined entirely by **where you are standing**. Nothing else. Move the camera and perspective changes. Change the lens without moving, and it does not.

The reason the myth is so durable is that it describes a real correlation. You *use* a wide lens up close, because that is what it is for, and standing close is what exaggerates near-to-far size differences. You *use* a telephoto from far away, and standing far away is what flattens them. The lens is associated with the effect because it is associated with the distance.

The test settles it in one move, and chapter 5.11 shows it: photograph the same subject at the same size in the frame from two different distances with two different lenses, and the backgrounds will be completely different. Then photograph from one spot with both lenses and crop the wide one to match — and the perspective will be identical.

Distance makes perspective. The lens only decides how much of the result you keep.

3. 3 Prime against zoom



A prime lens has one focal length. A zoom has a range of them. That is the entire difference, and almost everything said about the two is downstream of one fact: a zoom has to be good at several things, and a prime only has to be good at one.

Start with what a prime gives up, because it is not nothing and people who evangelise primes tend to skip it.

A prime cannot reframe. If the subject is the wrong size in the frame, you move, and sometimes you cannot move — there is a wall behind you, or a fence in front of you, or the thing is happening now and will not wait while you walk backwards. A zoom solves that in a wrist movement. For events, for wildlife, for anything where the distance is not yours to choose, that is not a convenience, it is the difference between having the picture and not.

And what it gains, which is real and concentrated in three places.

It can be faster. Because the design only has to work at one focal length, a prime can open wider — f/1.8, f/1.4 and beyond — where a zoom covering the same range would be enormous and ruinously expensive.

It can be better corrected. Chapter 5.6 covers the aberrations; the point here is that a zoom has to cancel them across its whole range, and a prime only has to cancel them at one point. One problem is easier than twenty.

It is smaller and lighter, which sounds like the least important of the three and is often the one that actually changes what you photograph, because a camera you carry takes more pictures than one you leave at home.

There is a fourth thing that is not an optical property at all, and I think it matters more than the other three. **A prime makes the framing decision a decision.** You cannot fiddle. You stand somewhere, and standing somewhere is a choice about perspective, which chapter 5.2 just established is the only thing that changes perspective. Photographers who use one prime for a while usually report that they start seeing in that frame — and that is a skill, learned by not having the alternative.

Now the specification-sheet detail that catches everyone out once.

Look at a zoom and you will see either one maximum aperture or two. A lens marked f/2.8 across its range is a **constant-aperture** zoom: wide or long, it opens to f/2.8. A lens marked f/3.5 to f/5.6 is **variable-aperture**: it opens to f/3.5 at the wide end and only f/5.6 at the long end.

The variable one is not the manufacturer being stingy, it is geometry. Module two defined the f-number as focal length divided by the diameter of the opening. Zoom to a longer focal length and, if the physical opening stays the same size, the f-number must rise. To hold f/2.8 at the long end, the opening has to grow in proportion — which means more glass, more weight and more money. That is why constant-aperture zooms are the heavy, expensive ones. You are paying for the front element to be big enough at the long end.

The practical consequence is worth internalising: a variable-aperture zoom loses light exactly when you zoom in, which is exactly when you also need a faster shutter speed to hold it steady. The two costs arrive together.

And on where the compromise concentrates in a zoom: it is usually at the extremes. A lens is generally at its best somewhere in the middle of its range and at its weakest fully wide or fully long — and often weakest of all fully long and wide open, which is, inconveniently, the setting people buy it for.

That tells you something useful. A zoom with a modest range is easier to make good than one with an enormous range, which is why a three-times zoom can be excellent and a twenty-times zoom is a series of apologies. Module three made the same point about bridge cameras and it is the same physics.

So, a first three-lens kit, and I will give an actual answer rather than hedging.

A normal-ish prime, fast, cheap, small. It teaches you to see, and it works in low light. Whatever "normal" is for your format — chapter 5.2's diagonal rule tells you.

A standard zoom, covering moderately wide to moderately long. This is the lens that will be on the camera most of the time and there is no shame in that.

One lens that does something the other two cannot. Long, if you photograph things you cannot approach. Wide, if you photograph interiors or landscapes. Macro, if the small world is what interests you. Chapter 5.8 covers the choices.

What I would not do is buy a fourth lens before finding out which of the first three is actually living on the camera. That answer tells you more about what you photograph than any amount of planning, and it is worth having before you spend again.

4. 4 Maximum aperture, and what it costs



Every lens has a widest setting, and it is printed on the front because manufacturers know people buy on it. It is worth understanding exactly what you get and exactly what you pay.

What you get is straightforward and comes in three parts. More light, so a faster shutter or a lower ISO — module two's triangle, with one more stop to spend. Less depth of field at a given framing, so more control over what is sharp — chapter 5.11. And a brighter view through the finder, which on an optical finder is a genuine working advantage and on an electronic one much less so.

What you pay is where it gets interesting, because the cost is not linear and people are surprised by how steeply it climbs.

Go from f/2.8 to f/2 — one stop — and the opening must be about one and a half times the diameter, which means the glass has to be about twice the AREA. Go another stop to f/1.4 and it doubles again. The elements get thicker, heavier and much harder to make to the required tolerance, and every aberration in chapter 5.6 gets worse towards the edges of a large piece of glass, so more correcting elements are needed to cancel them, which adds yet more glass.

That is why the price curve is so brutal. A modest 50mm at f/1.8 is one of the cheapest lenses you can buy. The same focal length at f/1.2 can cost ten or fifteen times as much, and it is not because the maker is greedy. You are buying a much larger, much more difficult object.

Whether that is worth it is a real question with a real answer, and the answer depends on whether you will actually use the widest setting. Which brings us to the thing nobody mentions in the shop.

Most fast lenses are not at their best wide open. At the widest aperture you typically get some softness away from the centre, some vignetting — darkening in the corners — and whatever aberrations the design did not fully cancel. Stop down one or two stops and all of it improves markedly.

So a photographer who buys an f/1.2 lens and then finds they shoot it mostly at f/2 has bought a very expensive f/2 lens — though a very good one, because it is being used inside its comfortable range.

The honest framing is this. You are not buying "an f/1.2 picture". You are buying **the option** of f/1.2 when you need it, and an excellent f/2 the rest of the time. If the option is one you will genuinely take — dark rooms, available light, a look you want — it is worth it. If it is aspirational, the money buys more elsewhere.

Finally, the marking that tells the truth and that most photographers never encounter.

Cinema lenses are not marked in f-stops. They are marked in **T-stops**.

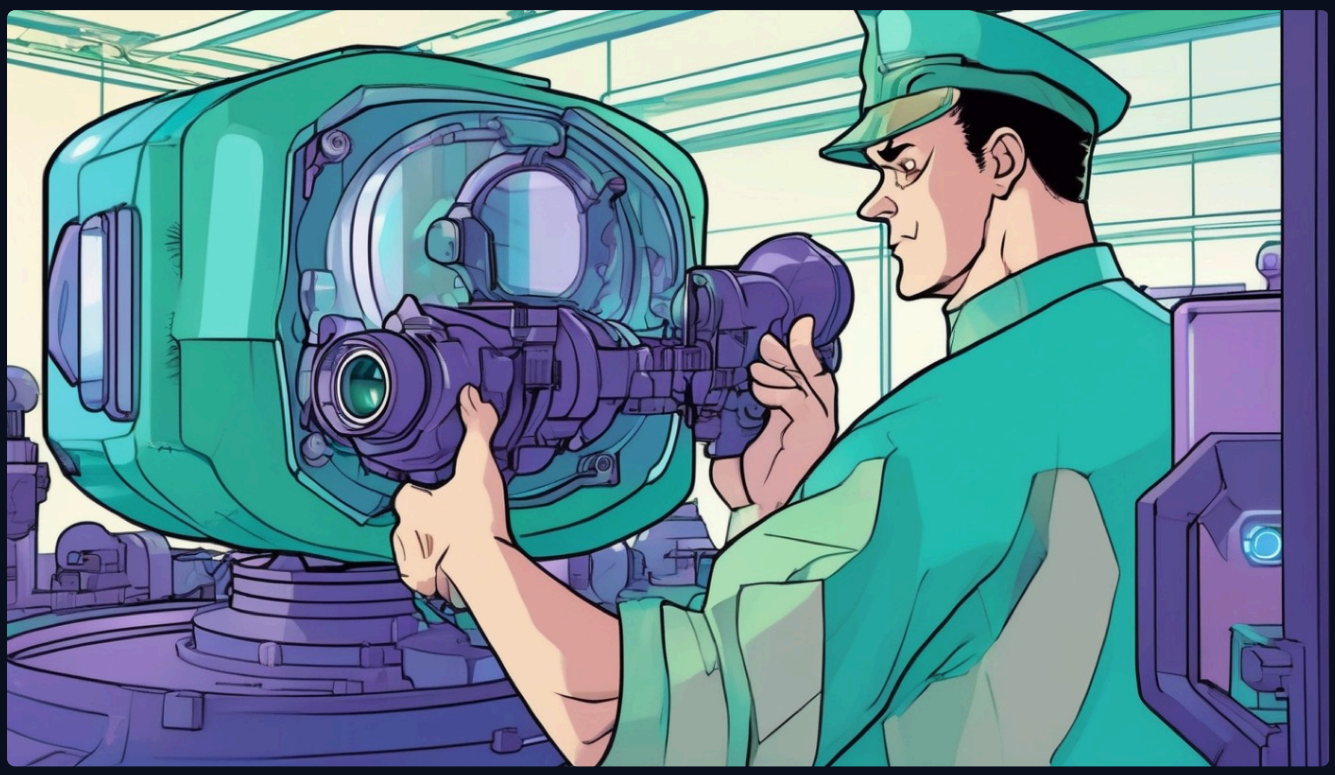
The difference matters. The f-number is a **geometric** figure: focal length divided by opening diameter, calculated from the design. It says nothing about whether the light actually gets through. Real glass absorbs a little, every air-to-glass surface reflects a little, and a lens with fifteen elements loses more than one with six.

A T-stop is the **measured transmission**. T/2 means this lens passes as much light as a theoretically perfect f/2 lens would. So a lens might be f/2 and T/2.3 — the geometry says f/2, the measurement says it behaves like f/2.3.

Cinema marks T-stops because a scene is shot on several lenses and cut together, and a brightness jump at every cut is unacceptable. Two lenses set to the same T-stop give the same exposure; two set to the same f-stop might not quite.

Stills photographers get away with f-numbers because the camera meters through the lens and absorbs the difference automatically. It is still worth knowing, because it explains why two lenses at the same f-number occasionally give slightly different exposures — and because it is one of the few places where photography admits that a number on a barrel is a calculation rather than a measurement.

5. 5 Elements, groups and coatings



Open a lens diagram — the cross-section in the brochure — and it looks like overkill. Fifteen pieces of glass, in ten groups, to take a photograph. A pinhole needs none. Why all this?

Because chapter 5.1's single piece of curved glass fails in several distinct, predictable and independently annoying ways, and every element after the first is there to cancel one of them.

That is the honest summary of lens design: **a modern lens is a stack of deliberate errors arranged so that they cancel.** Not a window. A negotiation.

Two terms first, because they are used loosely. An **element** is one piece of glass. A **group** is one or more elements cemented together and functioning as a unit. "Fifteen elements in ten groups" means fifteen pieces of glass, some of them glued into shared assemblies. The group count matters because every air-to-glass surface is a chance for reflection, so cementing two elements together removes two such surfaces.

Now the two kinds of glass that do most of the heavy lifting.

An aspherical element is not a section of a sphere. Ordinary lens surfaces are spherical because spheres are easy to grind accurately — but a spherical surface does not bring all rays to the same point: the ones through the edge focus slightly differently from the ones through the middle. That is spherical aberration, and chapter 5.6 comes back to it. An aspherical surface is shaped so that it does bring them together. It is harder and more expensive to make, and it lets one element do the work that several would otherwise do — so an aspherical element often makes a lens *smaller*, not larger.

A low-dispersion element attacks a different problem. Glass bends different wavelengths by different amounts — that is what a prism does, and a lens is a prism that has been persuaded to behave. The result in a photograph is colour fringing, which chapter 5.6 covers as chromatic aberration. Special glass formulations bend the colours much more nearly equally, so the fringing is reduced at the source rather than corrected afterwards.

Then coatings, which are invisible and matter enormously.

Every time light passes from air into glass, a few per cent reflects instead of transmitting. That sounds trivial until you count the surfaces: fifteen elements might mean twenty air-to-glass surfaces, and a few per cent lost at each one is a great deal of light not reaching the sensor.

Worse, that reflected light does not vanish. It bounces around inside the barrel and some of it lands back on the sensor where it does not belong, which is where **flare** and **ghosting** come from — the veiling haze across a frame shot towards the sun, and the little coloured shapes marching in a line away from a bright light.

A coating is an extremely thin layer, a fraction of a wavelength thick, applied so that the light reflecting off its outer surface and the light reflecting off its inner surface cancel each other. Modern lenses use several such layers, tuned to different wavelengths.

You can see the result by tilting a lens in the light: the faint coloured sheen on the front element is the coating, and the colour you see is the wavelengths it is least good at cancelling.

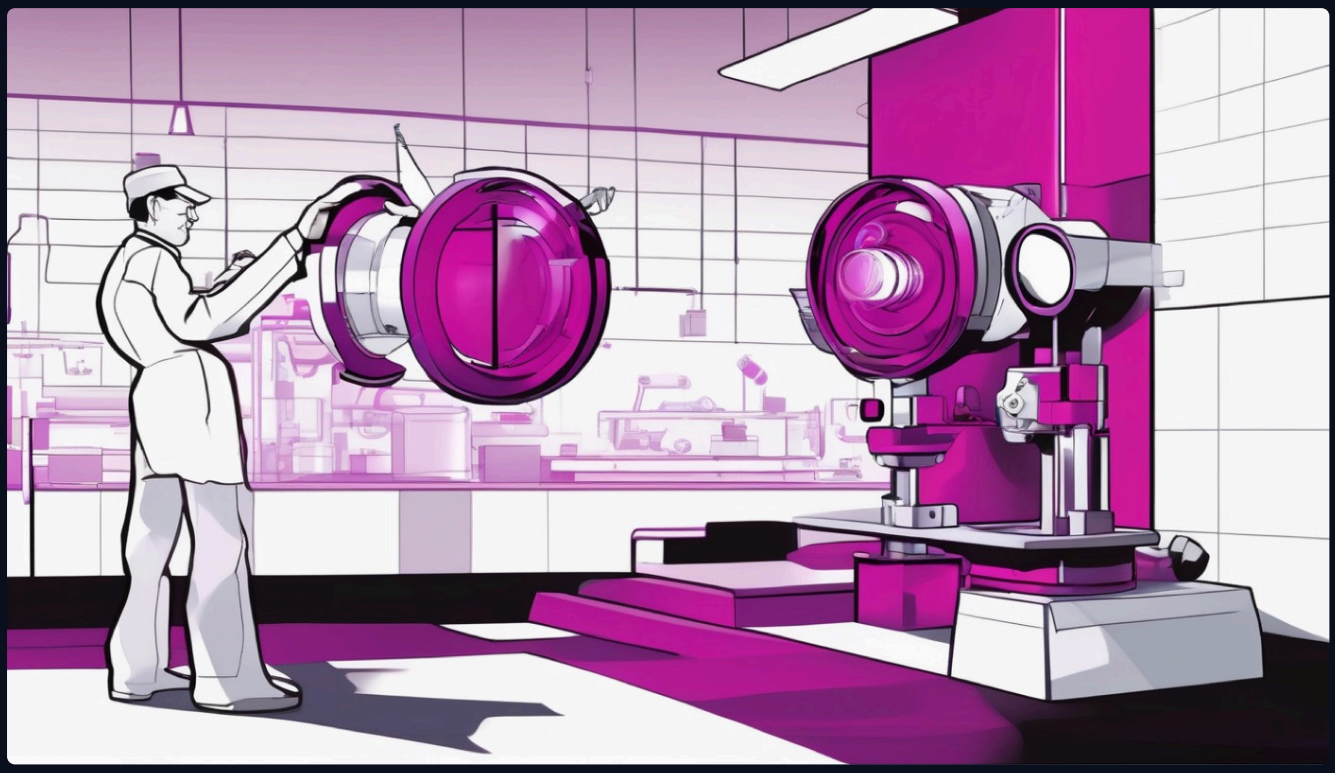
And one last piece of design that is easy to overlook and changes how a lens handles.

In an old lens, focusing moved the whole barrel, which extended and rotated the front element. In an **internally focusing** design, only a group inside moves. The lens does not change length and the front does not rotate.

That matters for two practical reasons. Filters that depend on orientation — a polariser, a graduated neutral density — stay where you set them, instead of being turned as you focus. And a lens that does not extend is easier to weather-seal and better balanced, because its centre of mass is not wandering.

None of this appears in a photograph directly. All of it decides whether the photograph is any good.

6. 6 Aberrations



An aberration is a way in which a real lens fails to do what an ideal lens would. There are a handful of them, they are well understood, and every lens you will ever use is a particular set of decisions about which ones to fight hardest.

Knowing them is not pedantry. It is what lets you look at a photograph that is not quite right and say what is wrong with it, which is the difference between fixing it and buying a new lens.

Chromatic aberration is the one you have probably seen without naming.

Glass bends different wavelengths by different amounts, so red, green and blue do not focus in exactly the same place. There are two flavours and they behave differently.

Lateral chromatic aberration means the colours land at slightly different distances from the centre of the frame – so they are aligned in the middle and increasingly separated towards the corners. It shows as coloured fringes, usually green and magenta, along high-contrast edges near the edge of the picture. It is essentially absent in the centre. It corrects well in software, because the fix is simply to resize the colour channels slightly relative to each other.

Longitudinal chromatic aberration means the colours focus at different DISTANCES – in front of and behind the plane of focus. It shows as colour fringing on out-of-focus edges: typically magenta in front of focus and green behind it, all over the frame including the centre. It does not correct well in software, because the information really is at a different focus distance. Stopping down reduces it. It

is common on fast lenses used wide open and it is one of the things you are paying to avoid in an expensive one.

Spherical aberration is chapter 5.5's reason for aspherical elements. A spherical surface focuses rays through its edge slightly differently from rays through its middle, so instead of a point you get a point with a soft halo around it. It reduces contrast and it is worst wide open, because wide open is when the outer part of the glass is in use. Stop down and it improves quickly. It is also, deliberately, the basis of some portrait lenses' flattering softness — an aberration kept on purpose.

Coma makes points of light near the edges of the frame render as small comet shapes with tails pointing outward. It is mostly irrelevant until you photograph stars, at which point it is the difference between points and a field of tiny birds. Astrophotographers care about coma more than almost any other measure, and module eight comes back to it.

Astigmatism means lines running in one direction focus at a different distance from lines running across them — so at any one focus setting, the radial detail is sharp and the tangential detail is not, or the other way round.

Distortion is the one that bends straight lines. **Barrel** distortion bows them outward, like the sides of a barrel, and is typical at the wide end. **Pincushion** pulls them inward and is typical at the long end. A zoom often does both — barrel at one end, pincushion at the other, and something in between in the middle. Distortion corrects very well in software, because it is a predictable geometric remapping, which is why modern lens designs sometimes allow quite a lot of it and let the camera fix it.

Field curvature is the subtle one, and the one that gets blamed on focus errors.

An ideal lens would focus a flat subject onto a flat plane. A real lens focuses it onto a gently curved surface — and the sensor is flat. So if you focus in the middle, the corners are focused slightly behind the sensor, and if you focus in the corners, the middle is out. On a distant landscape it usually does not matter. Photograph a flat wall, or a page of text, and it does: the corners look soft and no amount of refocusing fixes both at once. Stopping down helps, because more depth of field covers the curvature.

Which brings us to the honest question: what can software fix and what can it not?

Software corrects well: distortion, lateral chromatic aberration, vignetting. All three are geometric or a simple channel-by-channel adjustment, and modern raw converters apply profiles for known lenses automatically — often without telling you, which is worth knowing when you compare a raw file with what you saw.

Software cannot fix: longitudinal chromatic aberration, because the data is genuinely at the wrong focus; field curvature, for the same reason; and any loss of actual resolution. Sharpening increases edge contrast so detail *looks* crisper; it does not restore detail the lens did not deliver. Module ten is firm about that difference.

So when you read that a lens "corrects well in software", ask which aberration. For distortion it is nearly a free lunch. For the ones that live in the third dimension, there is no lunch at all.

7. 7 Stabilisation



Stabilisation exists because you are not a tripod. Your hands shake slightly, all the time, and at slow shutter speeds that shake becomes blur across the whole frame.

There are two places to fight it, and modern cameras often fight it in both.

Optical stabilisation puts a group of elements inside the lens on a movable mount. Sensors detect the camera rotating, and the group shifts to keep the image steady on the sensor. Because it is in the lens, it is designed for that lens's focal length, which matters because a long lens magnifies shake as much as it magnifies everything else.

In-body stabilisation moves the sensor instead. The advantage is that it works with **any** lens you mount, including old manual ones that have never heard of stabilisation, which for anyone adapting lenses — module three's flange-distance story — is a considerable gift.

Together, on systems that support it, they cooperate: the lens handles some axes, the body handles others, and the combination is better than either. There are five axes worth correcting — two shifts and three rotations — and a lens group and a sensor are good at different ones.

Now the number on the box, and what it actually means.

A manufacturer will claim a certain number of **stops** of stabilisation. Module two defined the stop as doubling or halving, and it means the same thing here: a claim of four stops means that a shot you

could hand-hold at one two-hundred-and-fiftieth, you should be able to hand-hold at about one-fifteenth — four halvings slower — and still get an acceptably sharp frame.

Those claims are comparable between manufacturers, and it is worth knowing why. There is a published standard: **CIPA's DC-011, "Measurement and Description Method for Image Stabilization Performance of Digital Cameras"**, whose current edition is dated 2024. Before a standard existed, every maker measured its own way and the numbers meant nothing across brands.

I am not going to quote you a figure, and the reason is a small point of principle this course keeps to. A stabilisation figure belongs to **one body with one lens**, tested a particular way. It is not a property of "stabilisation" as a technology, and quoting one would teach you something true of one camera and false of the next. Module three declined to assert one for the same reason.

What is worth carrying instead is the shape of the thing: stabilisation buys you slower shutter speeds, roughly in stops, and the standard is why the claims can be compared.

And now the limit, which is the single most useful sentence in this chapter.

Stabilisation does nothing whatever about subject motion.

It corrects the camera moving. It cannot correct the world moving. Photograph a running child at one-fifteenth of a second with eight stops of stabilisation and you will get a beautifully steady photograph of a blurred child — the background will be pin-sharp, which somehow makes it worse.

Module two's reciprocal rule was about camera shake, and stabilisation is exactly what relaxed it. The other half of the shutter-speed decision — is the subject moving? — is untouched, and no amount of stabilisation will ever touch it.

Two smaller practical notes. On a tripod, stabilisation can occasionally make things worse: the system looks for motion, finds almost none, and on some designs hunts. Many lenses and bodies detect a tripod now, but if you are getting unexplained softness on a tripod at slow speeds, turn it off and try again. And stabilisation costs battery, because something is being actively held still whenever the system is awake.

8. 8 Special lenses



Most photography happens on ordinary lenses. A few problems cannot be solved with one, and this chapter is the short list of the lenses that exist because of them.

A macro lens is built to focus very close and stay sharp doing it. Ordinary lenses are corrected for subjects a reasonable distance away, and a design optimised for infinity is not usually optimised for ten centimetres.

The number that defines macro is **magnification**, and it is worth being precise because marketing is not. **One to one** means the subject is projected onto the sensor at life size: a ten-millimetre insect occupies ten millimetres of sensor. On a full-frame sensor, thirty-six millimetres wide, that insect fills rather more than a quarter of the frame.

A lens that reaches one to one is a true macro. A lens sold as "macro" that reaches one to two is projecting the subject at half life size — useful, but it is a different thing, and the specification will tell you if you look for the magnification ratio rather than the word.

Macro brings two consequences that module eight puts to work. Depth of field at these distances is measured in millimetres, not metres. And the effective aperture changes as you focus closer — the geometry that defines the f-number shifts when the lens is extended a long way from the sensor, so a lens set to f/8 at one-to-one is passing meaningfully less light than f/8 at infinity. The camera's meter handles it; your shutter speed still has to survive it.

A tilt-shift lens does two separate things, and they are usually conflated.

Shift moves the lens sideways or vertically relative to the sensor, inside that oversized image circle chapter 5.1 mentioned. Point a normal camera up at a building and the verticals converge, because the sensor is no longer parallel to the building. Keep the camera level — so the sensor stays parallel — and *shift* the lens upward to include the top instead. The verticals stay vertical, because nothing was tilted. This is why architectural photographers own these lenses.

Tilt angles the lens relative to the sensor, and it does something that sounds impossible: it moves the plane of focus so that it is no longer parallel to the sensor. Ordinarily what is sharp is a flat slab at one distance. Tilt the lens and that slab tips over — so you can have a sharp plane running away from you along the ground, with things above and below it out of focus.

The rule governing it is the **Scheimpflug condition**: the plane of the subject, the plane of the lens and the plane of the image all meet along a single line. When those three planes share a line, everything on the subject plane is in focus at once. It is a piece of geometry from the large-format era that module three's movements chapter touched, and a tilt-shift lens is that capability bolted onto a modern body.

A teleconverter sits between lens and body and multiplies the focal length — commonly by 1.4 or 2. It is the cheapest reach you can buy and it is not free.

The cost is aperture, and it is exactly predictable. Module two's f-number is focal length divided by opening diameter. A teleconverter multiplies the focal length and does nothing to the diameter, so the f-number rises by the same factor: a 1.4x converter costs one stop, a 2x costs two. An f/2.8 lens becomes an effective f/4 or f/5.6. It also magnifies whatever aberrations the lens has, along with everything else.

Then the specialists, briefly. **A fisheye** does not correct distortion at all — it embraces it, to fit an extremely wide angle onto the frame, with straight lines bending everywhere except through the centre. **A mirror lens** folds the light path with curved mirrors instead of glass, making a very long lens surprisingly short and light, at the cost of a fixed aperture and doughnut-shaped out-of-focus highlights, because the light path has a hole in the middle. And **a soft-focus lens** deliberately leaves spherical aberration uncorrected, for a glow that is genuinely different from blur added afterwards, because it is happening in the optics rather than to the pixels.

9. 9 How autofocus works



Module four covered the autofocus controls: when the camera focuses, where it looks, and what it is looking for. This chapter is the part underneath — how a camera works out that something is in focus at all, which turns out to be two quite different ideas with quite different consequences.

Contrast detection is the simpler of the two, and the more obviously sensible.

An out-of-focus edge is a soft gradient. A focused edge is an abrupt change from light to dark. So the camera can look at the pixels in the focus area and measure how much contrast there is across neighbouring ones — and then move the lens and measure again. More contrast means closer to focus. Less means further away.

The problem is right there in the description. **Measuring contrast tells you how good the focus is, but not which way to go to improve it.** So the camera has to move the lens and see what happened. And because it cannot know it has reached the peak until it has gone past it and contrast has started falling again, it must overshoot and come back.

That is the little hunting motion you see on a camera focusing slowly in dim light. It is not indecision. It is the algorithm doing the only thing it can do.

Contrast detection is very **accurate** — it measures at the sensor, on the actual image, so there is no alignment to go wrong. It is just not fast, and it is poor at following something moving towards you, because by the time it has established which way to go the subject has moved again.

Phase detection solves exactly that problem, and the idea is genuinely elegant.

Take light from two different parts of the lens — opposite sides of the glass — and compare the two views. If the subject is in focus, those two views land in the same place. If it is out of focus, they are offset from each other, and here is the crucial part: **the size of the offset tells you how far out you are, and the direction of the offset tells you which way.**

So a phase-detection camera does not hunt. One measurement tells it "move this far, in this direction", and it drives the lens there in a single movement. That is why it can track a subject coming towards you — it is not searching, it is measuring an error and correcting it.

If you have used a manual-focus film SLR with a split-prism focusing screen, you have used phase detection with your own eyes: the image was split in two, and you turned the ring until the halves lined up.

The catch, in a traditional SLR, was that this happened in a **separate module** in the bottom of the camera, fed by a secondary mirror. The measurement was being made somewhere other than where the picture would be taken, and if those two places did not agree exactly, the camera focused confidently and consistently slightly wrong. That is what "front focus" and "back focus" meant, and why bodies had a menu for calibrating individual lenses. Module three told this story from the mirror's side.

On-sensor phase detection removed the whole problem by moving the measurement onto the imaging sensor itself. Some photosites are masked so that each sees light from only one side of the lens; comparing them gives the phase difference, at the sensor, where the picture is. Two things that had to be aligned became one thing.

That is the change that made mirrorless autofocus better than SLR autofocus rather than merely different, and it is why most cameras now run a **hybrid**: phase detection to get to approximately the right place fast and in the right direction, then contrast detection to confirm the last fraction.

And now what defeats both, because it is the same thing and it is worth knowing why.

Both methods need detail to work with. Contrast detection needs an edge whose contrast it can measure. Phase detection needs a recognisable pattern it can match between two views. Point either at a smooth white wall, a clear sky, or a subject in light so dim that the sensor is mostly noise, and there is nothing to measure.

That is why cameras hunt on blank surfaces and in the dark, and why the fix is always the same: give it an edge. Focus on something at the same distance that has some detail, or use the focus-assist lamp, or switch to manual. The camera is not being stupid. You have asked it to find a pattern in something that has none.

10. 10 Focusing in practice



Knowing how autofocus works is worth having. Using it well is a set of habits, and this chapter is those habits.

Module four covered the controls: single against continuous, the area modes, subject detection and back-button focus. The question here is when to trust it, and what to do when you should not.

Single focus for anything not moving towards or away from you. It locks, and a lock you can rely on is worth a great deal – you can recompose, you can wait, you can take five frames knowing all five are focused where the first one was.

Continuous for anything whose distance is changing. And the important point: continuous does not mean "better". On a static subject it is strictly worse, because the camera keeps re-evaluating and can drift onto something behind your subject at the moment you press.

The exception is the one module four described: with back-button focus, continuous behaves like single whenever you lift your thumb. That is why so many photographers leave the mode alone and control it with the button instead.

Tracking is the mode where the camera decides what the subject is and follows it around the frame, and modern tracking is very good. It is worth saying how good and where it stops.

It is excellent when the subject is distinct from its background, well lit, and of a type the camera recognises – module four's detection lists. It struggles when something crosses in front, when the

subject turns away so the recognisable part disappears, and when two similar subjects are close together, at which point it may confidently switch to the wrong one.

The practical posture is to trust it and watch it. The confirmation is on screen; if the box is on the right thing, take the picture. If it has jumped, take control back. Trusting it blindly is how you get a sharp photograph of the wrong person's face.

Manual focus still matters, in four situations that autofocus genuinely cannot serve: macro, where the depth is so shallow that you want to place the plane by hand; astrophotography, where there is nothing to lock onto; video, where you often want a deliberate, timed focus change rather than the camera's idea of one; and any adapted lens with no electronics.

Two aids make manual focus on a modern camera far better than it was on film.

Focus peaking highlights the edges that are currently sharp, in a colour, live. You turn the ring and watch the highlight sweep through the scene. It is fast and approximate — it tells you where focus is, not precisely how sharp.

Magnification zooms the live view to a hundred per cent or more so you can see the actual detail. It is precise and slow. Peaking to get there, magnification to confirm, is the sequence.

Finally, focus-and-recompose, which almost everyone does and which introduces an error worth understanding rather than fearing.

You put the centre point on the subject's eye, half-press to lock, then swing the camera to frame properly. The problem is that swinging the camera rotates it about the camera, not about the subject — so the subject is now slightly further away than it was when you locked, because it has moved from the centre of an arc to its edge.

How much this matters depends entirely on depth of field. At f/8 on a distant subject, it is nothing. At f/1.4 on a face a metre away, where the depth might be a couple of centimetres, it is the difference between the eye and the ear.

Three fixes, in order of preference. **Move the focus point instead of the camera** — module four's joystick, which is what it is for. **Recompose by moving your body** rather than pivoting, keeping the distance constant. Or **stop down** far enough that the error disappears into the depth of field.

Knowing the error exists is most of it. It explains the mysterious case where the eye is soft and the ear is perfect, and it stops you blaming the lens.

11. 11 Depth of field, bokeh and compression



Depth of field is the range of distances that look acceptably sharp. Everything in that sentence is doing work, and the word doing the most is *acceptably*.

Here is the thing that is rarely said plainly: **strictly, only one distance is in focus.** A lens focuses a plane. Everything nearer and everything further is projecting not a point but a small disc — and as you move away from the focused plane, those discs grow.

So there is no line where sharpness stops. There is only a point at which the disc becomes big enough to notice, and *that* is a judgement about viewing.

The disc has a name: the **circle of confusion**. The value used for it encodes assumptions — how big the picture will be, how far away the viewer stands, and how good their eyesight is. Change any of those and depth of field changes without a single thing about the lens changing.

That is why two depth-of-field calculators disagree, and why a picture that looked acceptably sharp as a postcard falls apart printed large. The lens did not lie. The convention did.

Four things change depth of field, and it is worth being able to list them.

Aperture. Smaller opening, more depth. This is the one everybody knows.

Focal length. Longer lens, less depth, at the same subject distance.

Distance. Closer subject, less depth — and this one is far stronger than people expect. Depth of field collapses at close range, which is why macro is measured in millimetres, and expands enormously towards infinity.

The format's circle of confusion, which is the convention above. This is where module three's second equivalence question comes from: compare formats *at the same framing* and the smaller one gives more depth of field.

Then a tool that falls straight out of the above, and is the most useful piece of arithmetic in landscape photography.

Hyperfocal distance is the focus distance at which depth of field extends from half that distance all the way to infinity. Focus there and everything from half-way to the horizon is acceptably sharp — the largest total depth any given aperture can give you.

Two practical warnings. It is built on the circle of confusion, so it inherits every assumption in it — a hyperfocal landscape can disappoint printed large. And focusing at infinity is *not* the same thing: it wastes all the depth that would have extended beyond infinity, where there is nothing to be sharp. Focusing slightly closer than infinity buys you the foreground for free.

Now bokeh, which is a quality and not a quantity.

Bokeh is not "how much blur". That is depth of field. Bokeh is the *character* of the blur — what an out-of-focus point of light actually looks like.

The shape comes mostly from the iris. Module two described it as overlapping blades; with a few straight blades, a point of light renders as a visible polygon. With more blades, and with curved ones, it renders closer to a circle. That is why lens specifications list the number of aperture blades, a number that looks like trivia and is a direct statement about out-of-focus highlights.

Wide open, the iris is out of the way and highlights are round regardless. Stop down and the blade shape appears. So a lens with pentagonal bokeh at f/4 may have perfectly round bokeh at f/1.4.

Other characters show up too: a bright ring around the edge of each highlight, from spherical aberration deliberately left in; the flat-sided highlights module four's electronic first curtain produces at very fast speeds; and the doughnuts a mirror lens gives, from the hole in the middle of its light path.

Whether any of it is good is a matter of taste, and photographers argue about it with real passion. What is not a matter of taste is the distinction: *depth of field is how much is blurred; bokeh is what the blur looks like.*

And finally, the demonstration this module has been promising since chapter 5.2.

"Compression" is the flattening you see in a telephoto photograph — the background looming large behind the subject, distances stacked up. It is routinely attributed to the lens, and the lens has nothing to do with it.

Here is the test, and it settles the argument completely.

First, photograph someone with a wide lens from close, filling the frame with their head. Then walk back a long way, fit a telephoto, and fill the frame with their head again. Same subject, same size in the frame, two very different photographs — the wide one has a vast background far behind them, the telephoto one has a background that looks enormous and close. This is the effect everybody calls compression.

Now, stand in one spot. Take one photograph with the wide lens and one with the telephoto, without moving. Then crop the wide one to match the telephoto's framing.

They will be identical. Same perspective, same relationship between subject and background, same apparent compression. The only difference is resolution, because you threw pixels away in the crop.

The conclusion is unavoidable. **Perspective is made by where you stand.** The focal length decides how much of the scene you keep — nothing more. "Compression" is the look of standing far away, and a long lens is simply the tool that makes standing far away practical.

Which is why the advice at the end of this module is not about lenses at all. If the photograph is not working, the first thing to change is not the lens. It is your feet.

Module six leaves the camera behind entirely and looks at the light.

Sources and credits

Primary sources

- CIPA, 'DC-011 Measurement and Description Method for Image Stabilization Performance of Digital Cameras', standard index at https://www.cipa.jp/e/std/image-stabilization.html and the 2024 edition document CIPA_DC-X011-2024_E.pdf at https://www.cipa.jp/e/std/image-stabilization/documents/CIPA_DC-X011-2024_E.pdf –, TITLE AND EDITION ONLY.

Not regulated financial advice.