Forwards or Backwards

https://forwardorbackwards.com



Learning Photography Course — 3. Every Kind of Camera

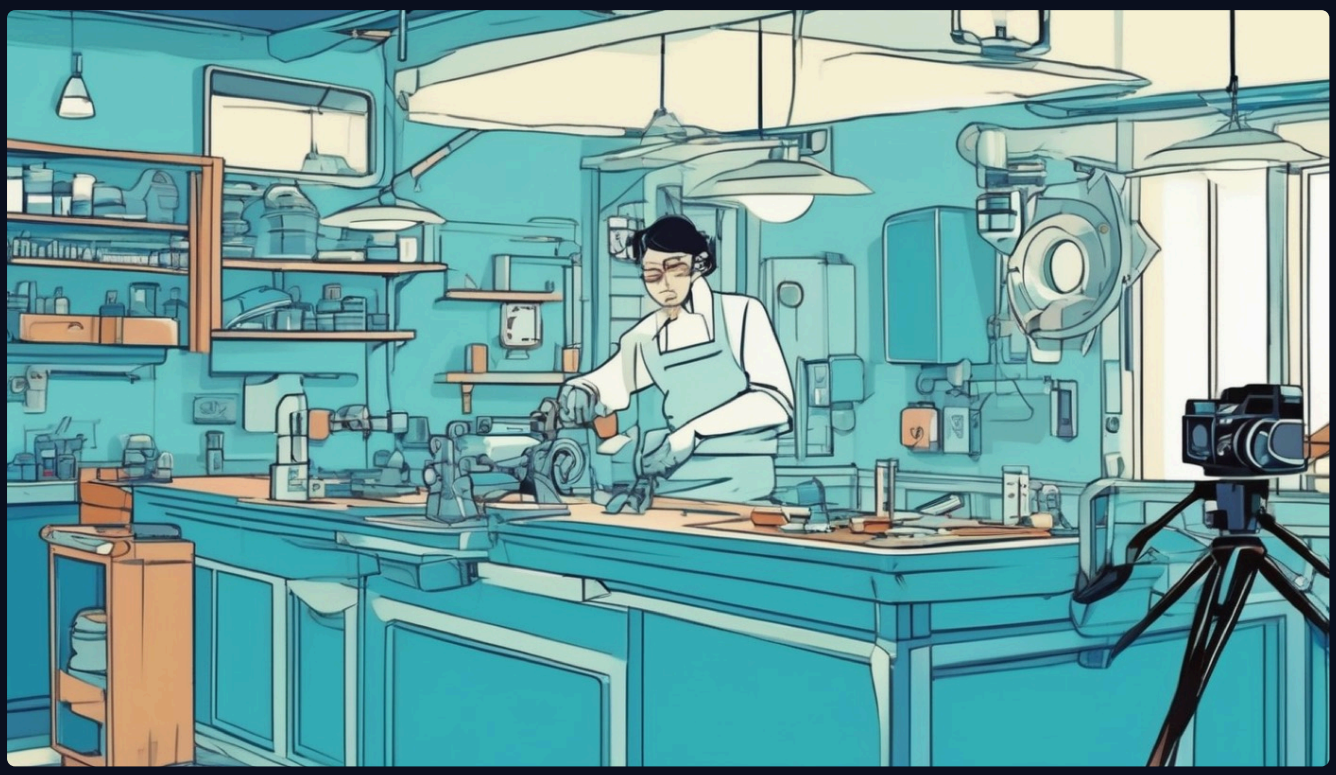
Watch on YouTube: https://www.youtube.com/watch?v=dGAWMpxLHhc

Chapters

1. 1 What every camera has in common
2. 2 Disposable and film compacts
3. 3 Digital compacts
4. 4 Bridge cameras
5. 5 The SLR
6. 6 Mirrorless
7. 7 Sensor size, done properly
8. 8 Medium and large format

9. 9 The phone
10. 10 The market in real numbers

1. 1 What every camera has in common



If you have been shopping for a camera recently, you will have been given the impression that the choice is enormous, that it matters a great deal, and that getting it wrong will hold your photography back for years. Two of those three are untrue, and this module is about which.

So let us start from the other end. Not with what separates cameras, but with what every camera that has ever worked has in common. Strip away the badge, the price and the marketing, and there are three things.

There is a hole, to let light in. There is a shutter, to decide for how long. And there is a recording surface — a sensitised plate, a strip of film, a silicon sensor — to keep what arrives. A disposable camera from a petrol station has all three. A large format field camera has all three. The camera in your pocket has all three. Module one called a photograph a measurement of light over an interval, and those three parts are simply the apparatus that measurement requires. The hole sets the rate, the shutter sets the interval, and the surface holds the result.

Two more things are needed before it is useful rather than merely functional. Something has to focus the image, which in practice means a lens that can move relative to the surface. And something has to let you see what you are about to take, which is the viewfinder in whatever form it happens to arrive.

That is a complete camera. Everything else — every feature on every spec sheet you will ever read — is a variation on one of those five parts. And the reason that matters is not philosophical. It is that it gives you a place to stand when someone tries to sell you something.

Because once the common parts are clear, the genuine differences become countable. There are four of them.

The first is the size of the recording surface. This is the one that changes the most about the resulting photograph, and the one most often described wrongly, which is why it gets a chapter of its own later in this module.

The second is whether the lens comes off. That single mechanical decision drives the size of the body, the price of the system, the range of what the camera can do, and how much money you will end up spending over the following decade.

The third is how you see the picture before you take it. Through the lens by mirror and prism, through a small screen, on the back of the camera, or not at all. This is the difference that people argue about most and that changes the photograph least — but it changes the experience of taking it enormously, and that is not nothing.

The fourth is how much computation sits between the light and the file. This is the newest of the four and the fastest-moving, and it is the one that makes a modern phone competitive with equipment it has no physical right to compete with.

Four variables. Everything else on the specification sheet is downstream of them.

Which brings us to reading one of those sheets, because a specification is written by someone who wants you to buy the camera, and it is worth knowing which of its numbers are load-bearing.

Some numbers are measurements. The sensor dimensions in millimetres. The maximum aperture of the lens. The flash sync speed. The weight. These describe physical facts, they are comparable between manufacturers, and they do not change when the marketing department has a meeting.

Some numbers are measurements taken under conditions the manufacturer chose. Battery life to a standard test, burst rate with certain settings, autofocus low-light sensitivity at a stated aperture. These are honest, but they are honest about a specific situation that may not be yours. Read the footnote; the footnote is where the condition lives.

And some numbers are not measurements at all. Megapixels above a certain point, once the print is the size it is actually going to be. Digital zoom, which is cropping with a confident name. Autofocus point counts, where a camera claiming a thousand points and one claiming five hundred are usually describing the same sensor divided up differently.

The test for any number on a spec sheet is simple and you can apply it in a shop. Ask what would change in the photograph if this number doubled. If you can answer specifically, the number matters. If you cannot, it is there to be compared against a rival, not to be used by you.

2. 2 Disposable and film compacts



The single-use camera is worth taking seriously, and not for nostalgic reasons. It is worth taking seriously because it is the clearest object lesson in the whole of photography about what a camera actually needs.

Inside the cardboard there is a plastic lens, usually a single element. There is a shutter with one speed, typically somewhere around a hundredth of a second. There is an aperture that does not move, usually around f/11. And there is a roll of film, normally rated four hundred ISO, which was loaded at the factory and which you will never see until it is processed.

Look at what that means. All three exposure controls — the three that module two spent an hour teaching you to balance — are welded in place. There is exactly one exposure the camera can give, and it gives it every single time you press the button.

And it works. It works because of a decision made by whoever specified it: colour negative film has a great deal of latitude, particularly towards overexposure, and a fixed exposure of f/11 at a hundredth on four-hundred-speed film sits comfortably inside that latitude for most daylight. The camera is not metering. It is betting, and the film is forgiving enough that the bet usually pays.

That is a real photographic insight hiding in a disposable object. A sufficiently tolerant recording medium removes the need for a great deal of apparatus.

What it cannot do follows just as directly from the same design, and it is worth being plain about, because the revival has produced a lot of writing that is not.

It cannot work in low light. The flash has a range of two or three metres and beyond that the frame simply goes dark. It cannot control depth of field, because the aperture does not move — although at f/11 on a short plastic lens, almost everything beyond a couple of metres is acceptably sharp anyway, which is the point of choosing f/11. It cannot focus, so anything closer than about a metre is soft. And it cannot be checked, so you will not know whether you got the picture for a fortnight.

The pictures have a look: slightly soft, heavily vignetted at the corners where the single-element lens gives up, contrasty, with the hard direct flash that a fixed tube produces. That look is the cumulative signature of every compromise in the list above.

Which brings us to the revival, and to being honest about what it is.

It is not that the pictures are better. On any measure you could put a number to — resolution, dynamic range, colour accuracy, noise — a ten-year-old phone beats a single-use camera comprehensively and it is not close.

It is that the constraints are the product. You get twenty-seven frames. You cannot see them. You cannot delete them. You cannot take the same picture nine times and pick. Every one of those is a limitation, and together they produce an experience that a device capable of infinite free frames cannot offer, because the value was in the scarcity rather than in the camera.

Film compacts proper — the fixed-lens, autofocus, point-and-shoot cameras of the eighties and nineties — sit one step up. They meter, they focus, many have a small zoom, and some have genuinely good lenses. They are the same bet made with better equipment. And their second-hand prices have risen to a point that is difficult to justify on optical grounds alone, which tells you the same thing: what is being bought is the way of working.

3. 3 Digital compacts



A digital compact is a camera whose lens does not come off, and it is important to understand that this is a design decision rather than a limitation that the manufacturer failed to overcome.

When the lens is fixed, the designer knows exactly what it has to cover, exactly how far it sits from the sensor, and exactly what aberrations it needs to correct. Nothing has to be left adjustable for a lens that might be fitted later. The lens can collapse into the body when it is switched off. The whole optical path can be shorter. And the corrections applied in software can be tailored to one specific lens on one specific sensor, because that pairing is the only one that will ever occur.

That is why a good compact can be pocketable in a way that no interchangeable-lens camera manages. It is not that the parts are smaller. It is that the problem is smaller.

There are two broad classes on sale, and the distinction between them is the sensor.

The larger class is built around what is universally called a one-inch sensor. Since this course does not let a misleading name stand, here is the correction, once: it is not an inch. The diagonal measures about fifteen point eight six millimetres, a little over six tenths of an inch. The name is inherited from the outside diameter of the glass envelope of a nineteen-fifties video camera tube, and it has survived for seventy years because nobody wanted to be the manufacturer who changed the numbering. Call it the type-one-inch class, understand that it is a name and not a measurement, and move on.

The other class uses sensors smaller again, typically around a third of the size in area, and it is the class that smartphones largely replaced. What is left of it now tends to be specialised: waterproof cameras, very long zooms, and cameras built to survive being dropped.

Which brings us to the number that surprised the industry.

In the twelve months of two thousand and twenty-five, worldwide shipments of cameras with a built-in lens were two million, four hundred and thirty-six thousand, nine hundred and eleven units. That is CIPA's own figure, from its December two thousand and twenty-five statistical release, and it covers shipments worldwide for the full calendar year.

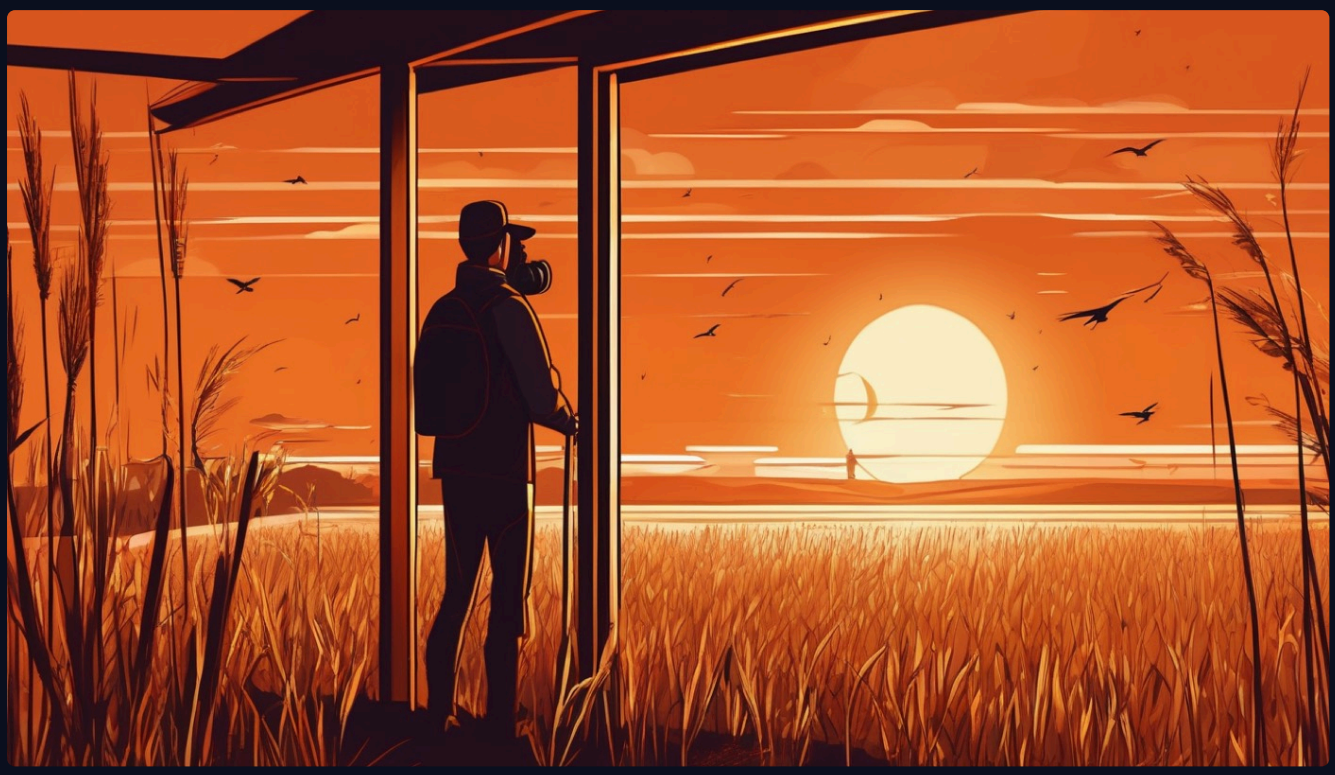
The interesting part is the comparison. That figure is one hundred and twenty-nine point six per cent of the previous year — an increase of twenty-nine point six per cent in a single year, in a product category that had been written off for a decade.

A word about how that percentage is presented, because it trips people up constantly and it will trip you up when you go and read the source. CIPA publishes its comparisons as a percentage of the previous year, not as a change. When the table says one hundred and twenty-nine point six, it means this year was one hundred and twenty-nine point six per cent of last year. The growth is the twenty-nine point six, not the hundred and twenty-nine. Reporting gets this wrong regularly.

The shipped value rose faster still, to one hundred and forty-eight point nine per cent of the previous year. More compacts, and more expensive compacts.

Nobody in the industry forecast this, and the explanations offered afterwards are guesses. What is not a guess is the shape of it: the category that smartphones were supposed to have finished is the fastest-growing category in cameras.

4. 4 Bridge cameras



The bridge camera is the most misunderstood category on sale, largely because it is judged against cameras it was never trying to be.

The proposition is a single number: reach. A bridge camera offers a zoom range that no interchangeable-lens system can match in one body at anything like the weight or the price. Sixty times. Eighty times. In the extreme cases, a field of view equivalent to twenty-four millimetres at the wide end and well over two thousand at the long end, from a camera you can carry on a strap all day.

To see why that is possible only here, hold two facts together. The sensor is small — usually the same small class as a basic compact. A small sensor needs a short actual focal length to give a given field of view. And a short actual focal length means small, light glass. The two thousand millimetre equivalent is not a two thousand millimetre lens; it is perhaps a two hundred and fifty millimetre lens in front of a sensor about an eighth the width of a full frame one. The equivalence is in the framing, not in the physics of the glass.

That is the bargain, and it is a genuine one. Now the price of it.

The cost lands on aperture, and it lands hard. A superzoom typically starts around f/2.8 at the wide end and finishes somewhere around f/6.5 at full extension. On its own that sounds survivable. But remember the sensor. Depth of field and total light collected both scale with sensor size, so f/6.5 on a sensor this small collects, for the same shutter speed, a small fraction of the light that the same f-number would deliver on a larger format.

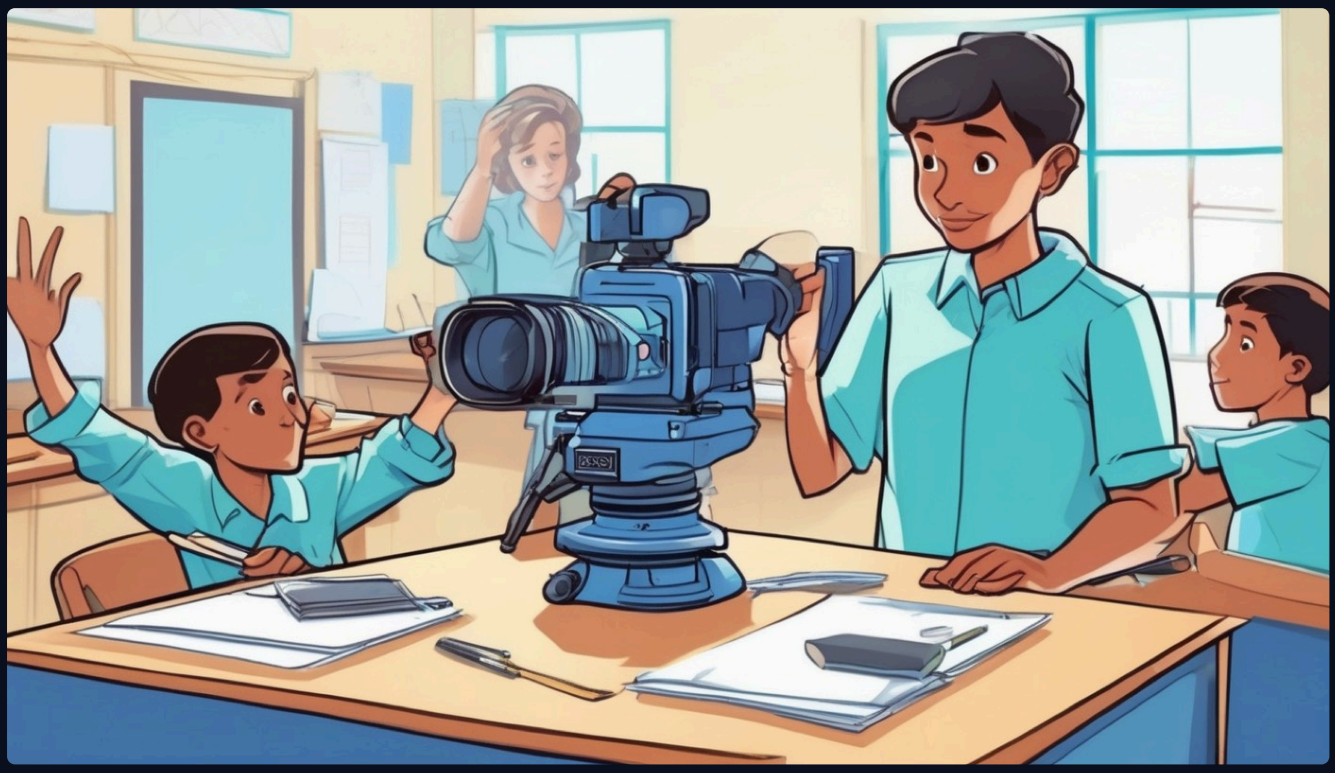
The practical consequence is the one every bridge owner discovers. The camera is transformed in bright light and struggles badly in poor light, and the struggle arrives precisely at full zoom, which is where you wanted to use it. You are asking for a long focal length, a small aperture and a fast enough shutter to hold a long lens steady, all at once, from a sensor that needs generous light to look its best.

So who is it right for? Not nobody, despite what the enthusiast internet implies.

It is right for anyone whose subject is a long way away in daylight and who is not going to carry a second bag. Wildlife at a distance where the alternative is a lens costing more than the holiday. Aviation. Sport from the back of a stand. Travel where the choice is honestly between this one camera and three lenses you will leave in the hotel.

And it is right for a reason that has nothing to do with optics: the picture you take is better than the picture you did not take because the correct equipment was at home. That is not a consolation. It is the most reliable rule in the whole of photography.

5. 5 The SLR



The single-lens reflex is the most elegant mechanism in photography, and it is worth understanding properly even though you may well never own one, because everything that came after it is defined by what it removed.

The problem it solves is old. If the viewfinder is a separate window in the body, you are not seeing what the lens sees. You get parallax — the framing shifts because the window is a few centimetres from the lens. You cannot see the depth of field. You cannot see whether a filter is doing anything. And when you change to a different lens, the viewfinder has no idea.

The reflex answer is to show the photographer the taking lens's own image, and the way it does that is a piece of mechanical ingenuity.

Behind the lens sits a mirror at forty-five degrees. Light coming through the lens hits it and is reflected upward, onto a focusing screen sitting at exactly the same distance from the lens as the film or sensor. So the image on that screen is in focus precisely when the image at the film plane is in focus. That is the whole trick, and everything else is delivery.

The image on that screen, however, is upside down and reversed, because that is what a lens does. Which is where the pentaprism comes in.

A pentaprism is a solid block of glass, five-sided, sitting above the focusing screen — it is the reason an SLR has a hump on top. Light enters it, reflects internally an odd number of times, and leaves

through the eyepiece the right way up and the right way round. Not approximately. Exactly. Move the camera right and the image moves right.

Then the exposure itself. The mirror has to get out of the way, because it is sitting directly in the light path to the film. So at the moment of exposure, the mirror swings up, the shutter opens, the exposure is made, the shutter closes, and the mirror drops back.

That sequence produces the two things every SLR photographer knows in their hands. Blackout: for the duration of the exposure the viewfinder goes dark, because the mirror that was feeding it is elsewhere. And mirror slap: a physical component accelerating, stopping and returning, which makes a noise and puts a small amount of vibration into the camera at the exact moment you would rather it did not. Mirror lock-up exists because of this, and so does the sound that people who like SLRs genuinely like.

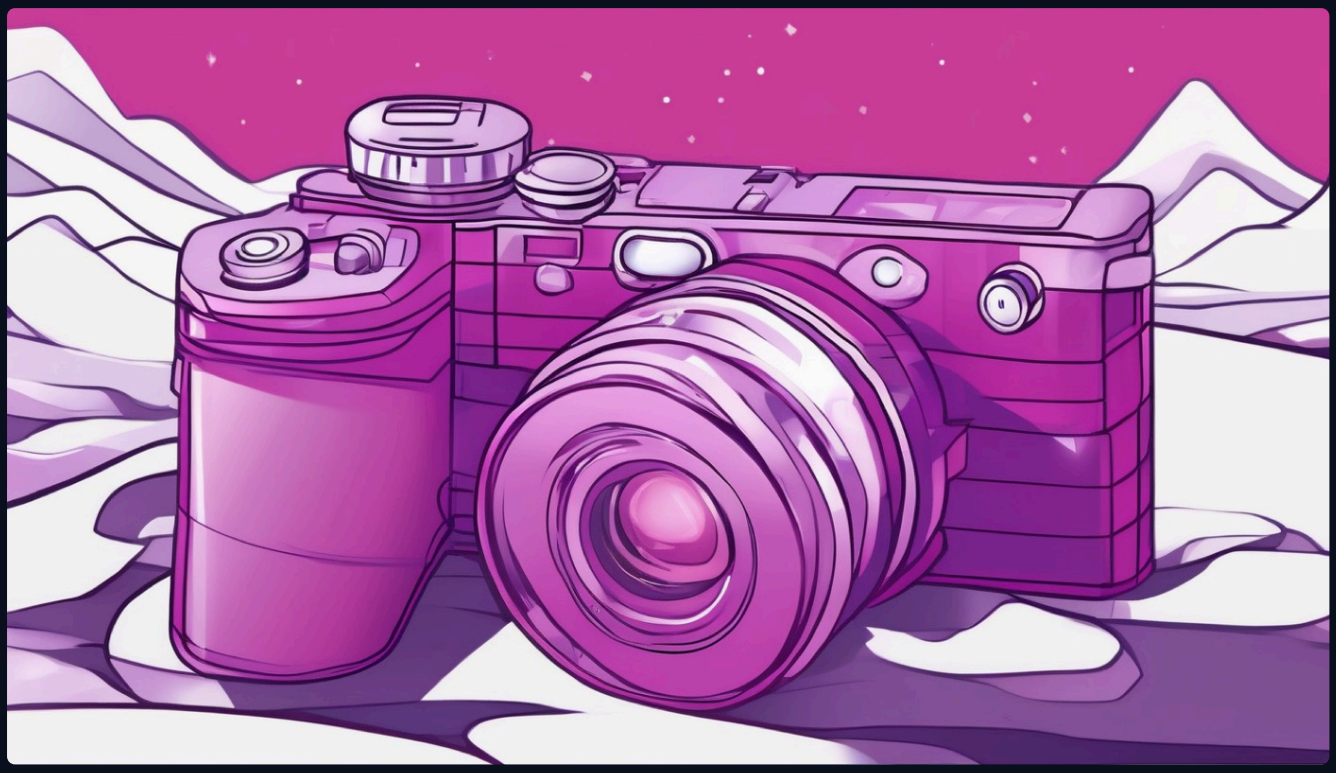
But here is what the mechanism buys, and it is the one thing no successor can offer.

What you see is light. Not a reconstruction, not a video feed, not a sensor read out and redrawn at some number of frames per second. Photons that came through the lens, bounced off a mirror and a piece of glass, and landed on your retina. There is no lag because there is no processing. In very low light the image gets dim, exactly as the scene is dim, rather than getting noisy or starting to stutter.

That is the SLR's irreducible advantage, and it is worth stating plainly because the category's decline sometimes gets told as though the design were flawed. It was not flawed. It was superseded on everything except this.

As for the transition from film SLR to digital SLR: less changed than you would think. The mirror, the prism, the focusing screen, the shutter, the mount and in many cases the lenses carried straight across. What changed was behind the shutter — a sensor and its electronics in place of a film gate. The viewfinder a photographer looked through in nineteen eighty-five and the one they looked through in two thousand and fifteen worked in precisely the same way. Which is why the lenses lasted through the change, and why they are still being adapted onto cameras today.

6. 6 Mirrorless



Take the mirror out.

That is the whole of it, and the consequences run right through the camera, which is why a sentence that short reorganised the entire industry in about fifteen years.

Start with what goes when the mirror goes. No blackout, because nothing has to move out of the light path. No mirror slap, so no vibration and considerably less noise. No mirror box, which is a substantial volume inside the body doing nothing but housing a swinging component. And no separate autofocus module.

That last one deserves unpacking, because it fixed a problem that used to be a routine irritation.

In an SLR, the autofocus sensor cannot be at the film plane, because the film is there. So part of the light is diverted, usually through a secondary mirror, down to a dedicated phase-detect module in the bottom of the body. That module has to be aligned so that when it reports focus, the image at the film plane is also in focus. When the alignment drifts, the camera focuses slightly in front of the subject, or slightly behind it, consistently. Front focus and back focus: real, common, and the reason cameras acquired micro-adjustment menus.

A mirrorless camera measures focus using the imaging sensor itself. The thing judging focus and the thing recording the picture are the same piece of silicon. There is no alignment to drift, because there are not two things to align. That is why the complaint largely disappeared, and it is a better explanation than "autofocus got better", though autofocus also got better.

Then the viewfinder, which is the real trade.

With no mirror, the finder is a small screen showing a live feed from the sensor. What you gain is substantial and easy to undersell. You see the exposure before you take it — change the shutter speed and the finder gets darker. You see the white balance. You see a histogram, a level, focus peaking, and magnification for manual focus. In very dark conditions the finder can amplify, showing you more than your eye would see through glass.

What you lose is that it is a screen. It refreshes at some rate, it has some latency, and it has a dynamic range and a resolution. Modern ones are extremely good and for most photography the losses are invisible. Following an unpredictable fast subject is where people still notice, and in very bright light an optical finder still has a directness that a panel does not reproduce.

Both halves of that are true at once. Anyone telling you only one half is selling something.

The last consequence is mechanical, and it is the one that quietly matters most.

Removing the mirror box shortens the flange focal distance — the gap between the lens mount and the sensor. And a shorter gap is a strictly more capable gap, because an adapter can add distance but cannot remove it. A lens designed for a long flange distance can be spaced correctly in front of a short one with a hollow ring of the right thickness. The reverse is impossible: you cannot move a sensor further back inside a camera that has already been built.

That single geometric fact is why a mirrorless camera will mount almost anything. Decades of SLR lenses, rangefinder lenses, cinema lenses, obscure mounts from companies long closed — all of it becomes adaptable. It is why there is a second-hand lens market of a size that would have astonished anybody in two thousand and five, and it is why lenses you already own are rarely an argument for staying where you are.

7. 7 Sensor size, done properly



This chapter exists because sensor size is where photography's most reliable argument lives, and the argument is only heated because three separate questions are being answered as though they were one.

Here are the formats, from the smallest upward. A phone sensor, typically somewhere between about five and ten millimetres on the diagonal. Four Thirds, seventeen point three by thirteen millimetres. APS-C, around twenty-three point five by fifteen point six, with small variations between manufacturers. Full frame, thirty-six by twenty-four millimetres — the dimensions of a frame of thirty-five millimetre film, which is the only reason that number exists. And medium format, which in digital terms usually means around forty-four by thirty-three millimetres, smaller than most medium format film ever was.

Now the three questions.

The first is field of view, and it is the easy one. A sensor of a given size, behind a lens of a given focal length, sees a certain angle of the world. Halve the sensor's dimensions and you see half the angle — you are cropping into the middle of the image the lens is projecting. That is all crop factor is. A twenty-five millimetre lens on Four Thirds frames like a fifty on full frame, because Four Thirds is half the linear size. It is pure geometry, there is nothing subtle in it, and it is the only part most explanations get to.

The second question is depth of field at the same framing, and here the answer changes.

Be careful with the phrase "at the same framing", because it is doing all the work. If you put the same lens at the same aperture at the same distance on two different sensors, depth of field is identical — the lens has no idea what is behind it. But you will not have the same picture. To get the same picture on the smaller sensor, you must either use a shorter lens or move further away, and both of those increase depth of field.

So the honest statement is this. A smaller format, framed the same way, at the same f-number, gives more depth of field. Not because f/4 means something different — f/4 is f/4 everywhere, and the brightness is the same — but because you had to change something else to get the same composition. The f-number did not change. What it does to the picture did.

Whether that is good or bad depends entirely on what you are photographing. More depth of field is a gift to a landscape photographer and an obstacle to a portrait photographer, and the same sensor delivers both.

The third question is noise, and this one is about light rather than about equipment.

Light arrives as discrete photons, and they arrive randomly. Count photons over a period and you do not get the same number every time — you get a fluctuation around an average. That fluctuation is shot noise, and it is a property of light itself. No sensor, however good, can remove it, because it was there before the light reached the camera.

The size of that fluctuation, relative to the signal, falls as the square root of the number of photons counted. Collect four times as much light and the relative noise halves.

Now put a larger sensor at the same f-number and the same shutter speed. The f-number fixes the brightness per unit area; the larger area therefore collects more total light. More total light, less relative noise. That is the whole mechanism, and notice that nothing in it is about the manufacturer, the sensor generation or the price. It is about how much light you caught.

Which is why equivalence arguments get so bad-tempered. Somebody says a smaller sensor is noisier and means the third question. Somebody else says f/4 is f/4 and means the first and second, and is also correct. Somebody says modern small sensors are excellent, which is true and is a statement about read noise and processing, not about shot noise. Three people, three questions, one argument nobody can win.

What is actually true is short. Field of view is geometry. Depth of field at the same framing changes because you changed something to get the framing. Noise follows total light collected. Hold those three apart and the whole topic goes quiet.

And then keep this in proportion. In two thousand and twenty-five, cameras with sensors smaller than thirty-five millimetre outsold larger ones. Four million, four hundred and fifty-eight thousand, six hundred and ninety-six interchangeable-lens bodies with sub-full-frame sensors, against two million, five hundred and forty-three thousand, two hundred and sixty-nine at full frame or larger — that

second figure derived by subtraction from CIPA's own totals for the year. The format people argue about most is not the format people buy most.

8. 8 Medium and large format



If a larger sensor collects more light and gives less depth of field at a given framing, the obvious question is why anyone stops at full frame. The answer is that the advantages continue, and so do the costs, and at a certain point the costs win for all but a few kinds of work.

What the extra area buys is the same three things, more of them. More total light for a given exposure, so cleaner files and more dynamic range in the shadows. Less depth of field at a given framing, so subject separation that is difficult to achieve otherwise. And more area to put photosites on, so more resolution without making each photosite smaller.

There is also a quality that is harder to put a number to and that ought to be described carefully rather than mystified. Because you need a longer lens to frame the same subject, the relationship between focal length, aperture and subject distance is different, and the transition from sharp to unsharp happens over a different distance. People call this the medium format look. It is a real consequence of real geometry, not a property of the sensor being expensive.

The mechanics differ too. Many medium format cameras use leaf shutters built into the lens rather than a focal-plane shutter in the body. A leaf shutter opens from the centre outward and closes back, exposing the whole frame at once at every speed – which means flash synchronises at any shutter speed the lens offers, rather than at the sixtieth to two-hundred-and-fiftieth a focal-plane shutter imposes. For anyone balancing flash against bright sunlight, that is not a detail. It is the reason the format is used.

Large format adds something else entirely: movements. Front and rear standards that tilt, shift, swing and rise independently. Shift keeps the film plane parallel to a building so that verticals stay vertical. Tilt swings the plane of focus so that it runs along a subject rather than across it, which is how a field can be sharp from the flowers at your feet to the hills behind without stopping down to nothing. These are controls over geometry that no amount of sensor quality substitutes for, and they are why architectural and product photographers kept using view cameras long after everyone else stopped.

So why did it stay professional? Cost is the obvious answer and the least interesting one. The real reasons are operational. The cameras are slower to use and slower to focus. Frame rates are low. Lens ranges are narrower, with nothing at the extreme wide or long ends. Depth of field is so shallow that focus accuracy becomes a constant discipline rather than an occasional worry. And the files are large enough to change how you work after the shoot.

Every one of those is acceptable if the picture justifies it and the client is paying for it. Almost none of them is acceptable if you are carrying the camera up a hill for your own pleasure. That is the whole of the explanation, and it has not changed in seventy years.

9. 9 The phone



The camera in your pocket has a sensor that, by every argument in the last two chapters, should not be able to compete. It is a fraction of the area of a Four Thirds sensor, behind a lens of a few millimetres' actual focal length, in a device whose primary job is something else.

And it competes anyway, in most conditions, for most people. Understanding why is the most useful thing in this module, because the mechanism is where photography is going.

The physics has not been repealed. The sensor still collects the light that falls on its area, and shot noise still behaves as it did in chapter seven. What changed is everything that happens afterwards.

A modern phone does not take a photograph when you press the button. It has already been taking them — a rolling buffer of frames captured before your thumb moved. When you press, it selects from that buffer, aligns a stack of frames to compensate for your hand moving between them, and combines them. Combining frames adds their light together, which is the one thing that genuinely attacks shot noise: more total photons counted, less relative fluctuation. A phone with a tiny sensor can, by stacking, collect more total light than a single exposure on a much larger sensor.

On top of that sits everything else. Exposure fusion for scenes with bright sky and dark foreground, blending frames at different exposures. Semantic understanding, where the processor identifies faces, skin, sky and foliage and treats each differently. Depth estimation to produce a blurred background by computation rather than by optics. Noise reduction and sharpening trained on enormous quantities of photographs.

The result is a picture that is not a measurement of light over an interval in the sense module one described. It is a reconstruction, informed by many measurements and by a great deal of prior knowledge about what photographs of this kind usually look like. That is not a criticism. It is a description, and it is the honest one.

So where does it win? Anywhere the scene is the ordinary difficult case — a person against a bright window, a street at dusk, a room lit by lamps. Anywhere the picture would not have been taken at all because the camera was at home. And anywhere the picture is going to be seen on a screen of about the same size as the one you took it with.

Where does it not? When you need real reach, because a long lens is a physical object and computation cannot manufacture detail that never arrived. When you want to control the look rather than accept a good one — the phone's decisions are made for you and are made well, but they are made. When the light is genuinely difficult in a way that stacking cannot fix, because the subject is moving. And when the file has to survive being worked on, because a great deal of the processing is baked in before you get to it.

And then there is the matter of scale, which is the sentence that puts the whole of this module in proportion.

In the twelve months of two thousand and twenty-five, worldwide smartphone shipments were one point two six billion units, according to IDC's figures released in March two thousand and twenty-six, which IDC labels as preliminary. In the same twelve months, CIPA recorded nine point four four million digital still cameras shipped worldwide — a figure the last chapter of this module takes apart properly.

Divide one by the other and you get about a hundred and thirty-three. For every dedicated camera shipped in two thousand and twenty-five, roughly a hundred and thirty-three smartphones were shipped.

That figure is a derivation, from two different organisations counting two different things, and it is rounded because the smartphone total is itself given to three significant figures. But it does not need to be precise to make its point. Almost every photograph taken in the world this year will be taken on a telephone, and the camera industry this module has been describing is, in volume terms, a rounding error beside it.

10. 10 The market in real numbers



The last chapter of this module is the industry itself, in figures, because almost everything written about the camera market is written from impressions and the actual numbers tell a stranger story than the impressions do.

Everything that follows comes from one document: the Camera and Imaging Products Association's statistical release for December and the full year of two thousand and twenty-five. CIPA is the trade body through which the Japanese camera manufacturers report, and since essentially every camera in this module is made by one of its members, its totals are the industry's totals. The figures are worldwide shipments, in units, for the calendar year.

Start with the interchangeable-lens cameras, because that is where the story usually gets told.

Mirrorless: six million, three hundred and eleven thousand and fifty-four units. That is one hundred and twelve point five per cent of the previous year, a rise of twelve and a half per cent. Of all interchangeable-lens cameras shipped, mirrorless accounted for just over ninety per cent — ninety point one, derived from CIPA's own totals.

Now the single-lens reflex. Six hundred and ninety thousand, nine hundred and eleven units. Sixty-nine point three per cent of the previous year, which is a fall of thirty point seven per cent in one year.

It is worth sitting with that for a moment. The SLR is the mechanism this module spent a chapter admiring, the design that defined serious photography for sixty years, and it is now just under ten per cent of interchangeable-lens shipments and about seven per cent of all cameras shipped. The

shipped value fell faster than the units did — to sixty-four point three per cent of the previous year, thirty point four billion yen — which tells you that the DSLRs still being made are also, on average, the cheaper ones. That is what the end of a product category looks like in the numbers.

Put the whole industry together. Compacts at two million, four hundred and thirty-six thousand, nine hundred and eleven. Mirrorless at six million, three hundred and eleven thousand and fifty-four. SLRs at six hundred and ninety thousand, nine hundred and eleven. Total: nine million, four hundred and thirty-eight thousand, eight hundred and seventy-six digital still cameras shipped worldwide in two thousand and twenty-five.

Those three figures add exactly to that total, which is a thing worth checking and worth saying you have checked. A widely read technology site reported the same release with figures that summed to about three hundred and fifty thousand short of the total it printed in the same article. The numbers in this film come from CIPA's own table, and they reconcile.

Nine point four million was one hundred and eleven point two per cent of the previous year. The industry grew by eleven per cent.

And here is where the impression and the number part company. The camera market's peak year was two thousand and ten, when shipments reached about a hundred and twenty-one million. That figure is CIPA's own annual series as it is generally reported; the association's monthly documents from that era are no longer published at their old addresses, so this course carries it as a reported figure rather than one re-read from the source, and will not put a decimal place on it.

Even with that caution, the comparison is unambiguous. Nine point four million against about a hundred and twenty-one million. The industry today ships under a tenth of what it shipped at its peak. Not a decline — a collapse, of well over ninety per cent, inside fifteen years.

So what does a collapse on that scale do to the companies that survive it?

It changes what they make. When the volume goes, the cheap end goes first, because the cheap end was competing directly with the thing that took the volume. What is left is a market where the average camera sold is far more expensive and far more capable than the average camera sold in two thousand and ten, and where the shipped value has fallen by nothing like as much as the unit count.

It changes who they sell to. A market of nine million is not a mass market. It is a market of people who have decided they want a camera specifically, which is a different customer with different expectations, and it explains a great deal about how cameras are designed and priced now.

And it changes the rate of change. Development budgets follow revenue, and a category shipping six hundred thousand units on falling value does not attract them. That is the mechanism behind the SLR's position, more than any judgement about the mechanism itself.

But finish on the part that nobody predicted, because a course that only told you about the collapse would be describing two thousand and twenty-three rather than today.

The industry grew in two thousand and twenty-five, and it grew having also grown the year before. The fastest-growing category was compacts, up twenty-nine point six per cent — the category everyone had finished writing obituaries for. Shipments outside Japan rose by nearly thirteen per cent; shipments inside Japan fell slightly, at ninety-eight point four per cent of the previous year, making it the one market that did not join in.

The first four months of two thousand and twenty-six continued in the same direction. CIPA's cumulative figures to April recorded two million, forty-eight thousand, nine hundred and twenty-five mirrorless bodies, one hundred and sixty thousand and eighty-seven SLRs, and seven hundred and ninety thousand, one hundred and sixty-four cameras with built-in lenses.

Which leaves a sensible place to end.

The camera industry is a tenth of the size it was, it is growing, and the growth is coming from the cheapest cameras it makes rather than the most expensive. None of that was forecast. Any confident claim you read about where photography is going next should be weighed against the fact that the people whose business it is did not see the last two years coming.

And for you, the practical conclusion of this whole module is smaller and more durable than any of these numbers. There are four things that genuinely differ between cameras: the size of the recording surface, whether the lens comes off, how you see the picture before you take it, and how much computation sits in between. Decide those four, and you have chosen your camera. Everything else on the box is a number waiting for you to ask what would change in the photograph if it doubled.

The next module opens up the body you have chosen and goes through every control on it, one at a time.

Figures in this module are worldwide shipments for the full year of two thousand and twenty-five, from CIPA's statistical release of December two thousand and twenty-five and from IDC's smartphone figures released in March two thousand and twenty-six, with CIPA's cumulative data to April two thousand and twenty-six where stated. This module was assembled in September two thousand and twenty-six.

Sources and credits

Primary sources

- CIPA, Production and Shipment of Digital Still Camera, December and January-December 2025 — worldwide shipments, full calendar year: total 9,438,876 units (111.2% of 2024); built-in lens 2,436,911 (129.6%); interchangeable lens 7,001,965 (105.9%); mirrorless 6,311,054 (112.5%); single lens reflex 690,911 (69.3%); ILC under 35mm 4,458,696.
- CIPA, cumulative January-April 2026: mirrorless 2,048,925; single lens reflex 160,087; built-in lens 790,164.
- IDC, worldwide smartphone shipments, full year 2025: 1.26 billion units, up 1.9%; released 5 March 2026 and labelled preliminary.

Not regulated financial advice.