

BACKWARDS ~ FORWARDS

WHAT A PHOTOGRAPH IS

A MEASUREMENT OF LIGHT




Learning Photography Course – 1. What a Photograph Is

Watch on YouTube: <https://www.youtube.com/watch?v=IdYFpe9F7Cc>

Chapters

1. 1 Light, and what is being measured
2. 2 The camera obscura
3. 3 Holding the image still
4. 4 Two inventions, two futures
5. 5 The latent image
6. 6 From silver to silicon
7. 7 The first digital camera
8. 8 The chronology, 1851 to 1963
9. 9 The idea the course rests on

1. 1 Light, and what is being measured



Before the cameras, before the settings, before any of the choices this course is going to teach you to make, one thing is happening, and it is worth being exact about what it is.

Light leaves a source. It strikes the world. Some of it bounces off a face, a wall, a wet road, and a small part of that travels towards you. Put a light-sensitive surface where your eye is, let the light fall on it for a while, and something on that surface changes. That change is the photograph.

Notice what is missing from the description. There is no image travelling through the air. There is no picture sitting in the scene waiting to be captured. There is light arriving, and there is a surface keeping a record of how much of it arrived, and where.

So the first thing to fix in your head is that light is a quantity. Not a mood, not a quality. Something there can be more or less of, and the amount can be counted.

And it can be counted literally rather than loosely, because light does not arrive as a smooth flow. It arrives in indivisible packets. That was Albert Einstein's proposal in 1905, and the Nobel Foundation's citation for the 1921 prize in physics honours him, in its own wording, "for his services to Theoretical Physics, and especially for his discovery of the law of the photoelectric effect." A packet of light is either absorbed or it is not. There is no half.

So how many packets are we talking about? The International System of Units lets us work it out, and the working is short enough to follow. Since the revised SI definitions took effect on the twentieth of May 2019, two of the constants involved are exact by definition rather than measured: the Planck constant, fixed at six point six two six oh seven oh one five times ten to the minus thirty-four joule seconds, and the luminous efficacy of green light at five hundred and forty terahertz, fixed at six hundred and eighty-three lumens per watt. Divide the one by the other and a single lux of that green

light works out at roughly four times ten to the fifteenth photons falling on every square metre, every second. Four thousand million million arrivals, on a patch the size of a large tabletop, in one second, at one lux.

That is arithmetic from the SI's own defining constants rather than a measurement anybody took, and I have told you the assumption it rests on – one wavelength of green rather than a full spectrum – so you can discount it accordingly.

Two things follow from a number that size. The first is that counting is the honest verb: there are genuinely discrete arrivals to be counted, and a camera is genuinely counting them. The second is that in good light they arrive so thickly that the graininess of the process never shows. Take the light away and they thin out, and a photograph made at the edge of darkness looks speckled because you are watching the arithmetic run short.

The habit of mind matters as much as the figure, though. Think of light the way you would think of rain – not as a condition outside the window, but as a volume collecting in a gauge. Leave the gauge out longer and it holds more. Stand it under a heavier downpour and it fills faster. Everything a camera does is a way of managing how much collects.

Here is where ordinary language starts working against us, because English uses one word – bright – for three separate ideas, and photography needs all three kept apart.

The first is a rate: how much light is falling on a surface at this moment, per second. A sunlit pavement has a high rate. The same pavement at dusk has a low one. Nothing has accumulated yet. This is only the speed at which it would.

The second is a quantity: what you get when that rate runs for a length of time. The total that actually landed.

The third is neither of those. It is how light the finished picture looks – an outcome, decided partly by the scene and partly by every decision taken afterwards. It is also the one beginners usually mean when they say a photograph came out too bright.

Keep the three apart and most of the confusion in the nine modules after this one never arises. Muddle them, and you can spend years adjusting the wrong one.

Photography has a unit for the second idea, and it comes from a published international standard rather than from a convention someone invented for a video.

ISO 12232 – *Photography: Digital still cameras. Determination of exposure index, ISO speed ratings, standard output sensitivity, and recommended exposure index* – is in its third edition, published by ISO on the 1st of February 2019 and confirmed current on systematic review in June 2024. It carries an amendment, ISO 12232:2019/Amd 1:2020, on the determination of encoding-relative sensitivity.

Inside that standard, formula one reads: exposure index equals K , divided by $H_{\text{sub } m}$. K is a constant, and ISO 12232:2019 fixes it at ten lux-seconds. $H_{\text{sub } m}$ is the average focal plane exposure, and the standard gives its unit as lux-seconds too. Exposure index itself is defined there as,

in ISO's words, "a numerical value that is inversely proportional to the exposure provided to an image sensor to obtain an image."

Read the unit slowly, because it is this whole course in three words. A lux-second is an illuminance multiplied by a time. Rate, times duration.

The first half of it has a definition at the very top of the measurement system. The ninth edition of the SI Brochure, issued by the International Bureau of Weights and Measures in 2019, lists the lux, symbol lx, as the coherent derived unit of illuminance, and expresses it in other SI units as lumens per square metre. Luminous flux, spread over area. That is the rate – how much visible light is landing on each square metre of a surface at this instant, with nothing said about how long it has been landing.

Bolt a second onto the end of it and you have said the other half.

A dose, then. Not a level. When a pharmacist gives you a dose of something, nobody asks how strong it was without also asking how much, and exposure behaves the same way. The international standard that governs it says so in its units.

Which brings us to the last idea here, and it is the one people find hardest to give up.

There is no such thing as an instantaneous photograph. Every exposure has a duration. It might be a four-thousandth of a second, it might be four hours, but it is never zero, because a zero-length interval collects a zero-length dose and records nothing whatsoever.

A photograph is therefore not a slice of time. It is a smear of it – a summary of everything that happened in front of the lens between the shutter opening and the shutter closing, added together. A face that moved is a blurred face because the light from that face landed in more than one place while the interval ran.

For someone holding a camera, that pays for itself immediately. When you choose how long the shutter stays open, you are not choosing a moment. You are choosing the length of the thing you are describing.

2. 2 The camera obscura



Now take that dose of light and ask a harder question. Light from every point of a scene is flying in every direction at once. Hold a blank sheet of paper up in a lit room and nothing appears on it, because light from the window, the lamp, your hand and the far wall all land on every part of the paper together. The information is there. It is simply hopelessly mixed.

The camera obscura is the oldest fix for that, and it is almost insultingly simple. Put the paper inside a dark box. Make one small hole in the opposite side.

Now each point in the scene outside can only reach the paper through that one hole, which means it can only send a single narrow bundle of rays, arriving at a single small patch. The bundle from the top of a tree lands low. The bundle from the bottom lands high. Because the bundles cannot overlap, they cannot mix, and what appears on the paper is the scene – upside down, the right way round for the geometry, and often startlingly bright for something made by nothing more than an absence of wall.

Britannica, in its account of twelve key dates in the history of photography, records that the Arab astronomer and mathematician Ibn al-Haytham is thought to have been the first to build a camera obscura, in the 10th or 11th century of the common era. That is roughly a thousand years ago, and note that it is given as a span of two centuries, not a date. The uncertainty is real and this course is going to keep saying so.

Then comes the trade, and it is the first genuine trade in photography.

How small should the hole be? Make it smaller and each bundle of rays gets narrower, each patch on the paper gets tighter, and the image gets sharper. But a smaller hole also admits less light, so the dose arriving per second falls, and the image gets dimmer. Make the hole bigger and you get the

opposite: brighter, and softer, because each point of the scene is now spreading across a wider patch and the patches have started to overlap again.

Sharper and dimmer. Brighter and softer. You cannot have both from a hole.

That is the aperture trade-off – the thing module two spends a chapter on, and module five returns to – and here it is, arriving about five hundred years before anyone had a lens to put it in. When you stop a modern lens down to $f/11$ for depth of field and then find you need a longer shutter speed, you are paying exactly the bill Ibn al-Haytham's hole presented.

There is also a floor under the hole, and it is not a limit of workmanship. It is physics, and it works against you in the direction you would least expect.

Light passing the rim of an opening spreads sideways. The narrower the opening, the more it spreads. So as you keep shrinking the hole, a second kind of blur is growing while the first kind shrinks, and at some width the growing one overtakes the shrinking one. Past that point, making the hole smaller stops improving the picture and starts wrecking it.

The width where the two balance carries Lord Rayleigh's name. In the form given in Costantino Sigismondi's paper *Introduction to Pinhole Astronomy*, the best diameter equals one point nine, multiplied by the square root of the wavelength times the distance from hole to screen – with the wavelength taken as five point five times ten to the minus seven metres, which is green light near the middle of what the eye handles best. Run that for a box a hundred millimetres deep and the best hole is about nought point four five of a millimetre across. Open it wider and the picture softens because the patches overlap. Pinch it narrower and the picture softens again, for the opposite reason.

So it is not one trade but two limits leaning in from either side, with a single width standing between them.

And exactly the same physics is waiting at the small end of a modern lens, which is why photographers talk about a lens going soft when it is stopped right down. The engraving on the ring says $f/22$, and it is easy to assume that must be the sharpest setting the lens owns. It is very often the least sharp one, for a reason that was already true of a hole cut in a wall.

The obscura also travelled. It began as architecture – a darkened room, sometimes a whole building, with people standing inside it looking at a wall – and over time it shrank into furniture, and then into something a person could carry and set up in a field, with a lens in the hole and a ground-glass screen to look at.

The National Science and Media Museum in Bradford dates those changes in centuries rather than in years, and that vagueness is the museum being careful rather than idle. By the middle of the sixteenth century, its account says, lenses had begun to be used to increase the brightness and sharpness of the image – which is the first moment anyone put glass in the opening and got both of the things a bare hole refuses to hand over together. By the seventeenth century, the same account

continues, smaller portable versions had appeared, and it names the shapes they took: tents, sedan chairs and pocket models.

Individual instances can be pinned down more tightly than that. Johannes Kepler, the man who gave the device its Latin name, was working with a mobile one in the sixteen-twenties, built as a small black tent with room inside for a single person. Robert Boyle described a portable darkened room of his own in 1669, and is sometimes credited with the first handheld version.

Two named men, roughly half a century apart, and between their two dates a piece of architecture has become a piece of luggage. Note how long the whole business took: a thousand years from Ibn al-Haytham to a tent one man could pitch, and the optics had not changed at all. Only the carrying had.

And in that portable form it was, for practical purposes, already a camera. It had a hole, it had a chamber, it had a surface where the image formed. An artist could sit in front of it and trace.

Trace. That is the word that gives away the missing piece.


Because the one thing the camera obscura could not do is the one thing this module is about. The image on that screen exists only while the light is arriving. Look away and it is still there. Close the hole and it is gone, completely, leaving nothing behind. There is no record, because nothing on that screen has changed.

A camera obscura measures light, in the sense that it sorts the light arriving now into a pattern. What it does not do is hold the measurement. Everything in the next three chapters is people trying to make a surface that keeps a mark.

Before we leave the hole, one exercise, because it costs nothing and it makes the rest of the course physical rather than theoretical. Curl a finger against your thumb until you are looking through a gap smaller than your pupil, and read something across the room through it. If you normally need glasses for distance, the text will sharpen. Now widen the gap slightly and watch it soften again, and notice, as you do, that the whole view has got brighter.

You have just run the entire trade by hand. Nothing about that changes when there is glass involved and a number engraved on a ring. The lens on a modern camera is a device for gathering more light than a bare hole would and bending it to a point, but the hole is still in there, still adjustable, and still charging you the same price.

3.3 Holding the image still



The oldest photograph that still exists is a view out of an upstairs window, and it is not much to look at.

It is called *View from the Window at Le Gras*, and it was made by Joseph Nicéphore Niépce at Saint-Loup-de-Varennes, in France. Britannica dates the work to 1826 or 1827. The Wikipedia article on the plate gives a tighter window – between the 4th of June and the 18th of July 1827 – and both are worth carrying, because they are answering slightly different questions about the same object.

The thing itself is a sheet of pewter, sixteen point two centimetres by twenty point two, thinly coated with a substance called bitumen of Judea. It lives now at the Harry Ransom Center in Austin, Texas, inside a continuously monitored, stabilised, oxygen-free case designed by the Getty Conservation Institute. A photograph two centuries old is kept in an atmosphere with the oxygen taken out of it, which tells you something about how fragile the first solution to this problem turned out to be.

Then there is the exposure time, and this is where the course does something you should expect it to keep doing.

The traditional account says the exposure lasted about eight hours. That figure has been repeated in accounts of the plate for generations.

But a modern researcher, who studied Niépce's own notes and recreated his processes, concluded that the exposure must have continued for several days.

Eight hours, or several days. This course is not going to pick one for you.

That is a deliberate policy, so let me say why, because it will save you from a lot of confident nonsense elsewhere. When two reputable sources give different figures for the same thing, there are

only a few possibilities. One of them is wrong. Or they are measuring different things and both are right. Or the underlying evidence genuinely does not settle it. Quietly choosing the more convenient number hides which of those is the case, and the choice usually gets repeated by everyone downstream until it hardens into a fact nobody can trace.

There is a second reason here, and it is that the disagreement is more interesting than either number. Whether the answer is eight hours or several days, the shape of the finding is identical: at the beginning of photography, the interval in "light over an interval" was measured in hours or in days. Not in fractions of a second. The first photograph is proof of the whole idea, because you can see the length of it.

So what was actually happening on that plate for all that time?

Niépce called the process heliography, and it works the opposite way round to everything that came after it. Bitumen of Judea hardens where light strikes it. Sit the coated plate at the back of a camera obscura for long enough and the bright parts of the scene harden the bitumen underneath them, while the dark parts leave it soft.

At that point there is still nothing to see. The image is a pattern of hardness, not a pattern of tone.

To bring it out, Niépce washed the plate with a mixture of lavender oil and white petroleum, which dissolved away the bitumen that the light had never hardened. What that wash removes is the shadows. What it leaves behind is hardened bitumen standing where the light was, against bare pewter where it was not.

Two details in that are worth carrying forward for the rest of the course. The first is that the plate had to sit in the dark box for hours or days, entirely passively, doing nothing a bystander could have noticed. The second is that a separate chemical step afterwards was what made the result visible at all.

Light does the recording. Something else does the revealing. Hold on to that, because the next three chapters are variations on it.

There is a practical reading of this chapter as well, and it is about scale. The interval in a photograph runs from something like days, on that pewter plate, to the fastest setting on the camera in front of you. Your shutter dial is a short stretch of an enormously long line, and the reason a night scene needs thirty seconds while the same lens outdoors at noon needs a five-hundredth is not a difference in kind. It is the same measurement, arriving at wildly different rates, reaching the same total by running for wildly different lengths.

Get comfortable with that and long exposures stop feeling like a specialist technique. A thirty-second frame is simply an ordinary photograph of a scene supplying very little light per second, which is precisely what Niépce was taking.

4. 4 Two inventions, two futures



Twelve years or so after the view from Le Gras, the problem was solved twice, in two different countries, in two incompatible ways. The difference between them decided what photography would be for the next century and a half, and it is still deciding what happens inside your camera today.

The first answer came from Louis Daguerre, and it was announced in 1839. Britannica describes the daguerreotype as a polished copper plate coated with light-sensitive silver iodide. Of the two processes that arrived in this decade, it was the sharper.

It had one property, though, that is easy to skim past and impossible to design around. A daguerreotype is a single plate, and that plate is the photograph. There is no master, no original behind it, nothing to reprint from. One exposure produces exactly one object. If you wanted two, you took two.

The second answer came from William Henry Fox Talbot, whose calotype process Britannica records as patented in 1841. Its images were softer than Daguerre's. Set the two side by side in 1841 and the daguerreotype wins on sheer quality without an argument.

But the calotype did something the daguerreotype structurally could not. It produced a negative – and from that single negative, as Britannica puts it, multiple positive images could be made.

Read those two sentences again as a pair, because this is the fork.

Daguerre's process gives you an object. Talbot's gives you a master, and an unlimited number of prints from it. One is a photograph. The other is a system for producing photographs.

Photography went the negative way, and stayed there. Through wet plates, through dry plates, through roll film, through every 35mm camera anyone has ever loaded, the working pattern has

been the same: expose a thing that is not the picture, then produce the picture from it, as many times as required and as differently as required.

And here is the part that matters to you, holding a digital camera, in the present day. A raw file is a negative. Not by analogy, and not as a figure of speech, but in the precise sense that has mattered since 1841: it is not the finished picture. It is a record you make prints from. Every adjustment you apply to it is a printing decision, the original is untouched underneath, and you can go back and print it differently tomorrow.

People who shoot raw and then talk about "the photo" as though the file were the final object have quietly given up the advantage Talbot won, and are working, in structure, like a man in 1839 with one copper plate.

One more thing about 1839, and then we will get to how any of this chemistry works.

The year is remembered as photography's founding date, and it was also the year the arguments started – about priority, about who had done what first, and about whose claim deserved which credit. Those arguments are a genuine part of the story rather than a footnote to it.

Here is the sequence, with the dates the institutions themselves hold.

On the seventh of January 1839, François Arago presented the daguerreotype to the Académie des Sciences in Paris. That is the date the Musée d'Orsay gives, and notice who is speaking: not Daguerre, but his advocate – an astronomer and politician with the ear of the Académie.

The news reached England within days, and it flushed out a man who had been working quietly and publishing nothing at all. On the twenty-fifth of January 1839, William Henry Fox Talbot exhibited his photogenic drawings in London at the Royal Institution, and that is the date under which the Bodleian Libraries' Talbot project files the event. Six days after that, on the thirty-first of January 1839, his paper went before the Royal Society under a title that is a small manifesto by itself: *Some account of the art of photogenic drawing, or the process by which natural objects may be made to delineate themselves without the aid of the artist's pencil*. The Royal Society still holds the manuscript, catalogued as A P, twenty-three, nineteen.

Three weeks and a bit separate the two claims. What turns 1839 from a race into a grievance is the third man.

Hippolyte Bayard had a process of his own – a direct positive, yielding one unique print on paper, with no negative and no copies. His first noted success came in March of 1839, and he mounted what is described as the first public exhibition of photographs on the twenty-fourth of June that year. He was also, the usual account says, persuaded to delay announcing his work to the Académie by Arago, a friend of Daguerre's.

That last detail is a serious charge against the conduct of the man who ran the announcement, and it deserves a health warning. The strongest version of it this course could reach is an encyclopaedia entry rather than a minute from the Académie or a letter in an archive. Carry it as the widely told account. Do not carry it as proven.


Bayard's answer came in October 1840, and it is the first staged photograph anybody made: a self-portrait posed as a drowned man, captioned with a complaint that the government had given everything to Daguerre and nothing to him. A man inventing photography and inventing the fake news photograph inside eighteen months of each other.

Three claimants, three incompatible processes, one year. The lesson for someone learning photography is not which of them deserved the credit. It is that a technology with a single obvious inventor and a clean founding date has almost never had either, and the tidy version you were taught is usually the version whose author had the better advocate in the room.

What you do with all of this is a habit rather than a technique. Treat whatever your camera writes as the negative and treat everything after it as printing. Keep the original file untouched, work on copies of it, and expect to return in a year and print it differently because you have changed your mind or learned something.

That discipline has one more consequence worth naming now. Because the negative is not the picture, it does not have to look like one. A raw file that appears flat and lifeless on the back of the camera is not a failed photograph, any more than a strip of developed film held up to a window is. It is a record holding more than the preview can show, and the preview is only ever the first print anyone has pulled from it.

5.5 The latent image



Silver is the material that made photography practical, and to understand why, you have to look at a single crystal of silver halide – a compound of silver with one of the halogens – suspended in gelatin, one of countless millions spread across a film or a plate.

When light strikes that crystal, the account given by chemists runs roughly as follows. A photon lifts an electron out of the crystal's valence band and into its conduction band, where it is free to move. The crystal contains silver ions that are also free to move. A freed electron gets trapped somewhere in the structure, a mobile silver ion arrives, the two react, and what they leave behind is a single atom of metallic silver.

One atom. In a crystal. Out of a scene containing a cathedral and a sky.

What makes the process work is what happens next. In the words of the chemistry account this course is drawing on, metallic silver in its turn becomes an electron trap and the process is repeated – "one atom of silver begets two, then four and so on." The first atom makes the next arrival more likely at the same spot, and the effect compounds. That is why the sensitivity of a photographic emulsion depends, as the same source puts it, on the mobility of electrons, of interstitial silver ions, and of the holes in the valence band. All three have to be able to move for the specks to build.

Now the crucial fact, and it is the one this chapter is named for.

After the exposure, the film looks exactly as it did before. Nothing has changed that any eye or instrument in an ordinary darkroom could see. The scene is recorded – the silver specks are sitting there in a pattern that corresponds to where the light fell – but the quantity of silver involved is far too small to be visible. That invisible record is called the latent image, and it can sit in the cassette for weeks while you finish the roll.

Which reframes what development is. Development is not the moment the picture is created, and it is not a process of revealing something that was hidden. It is amplification.

The developer's job is to distinguish crystals that carry one of those specks from crystals that do not, and to convert the whole of a marked crystal into metallic silver while leaving unmarked ones alone. A vanishingly small signal selects which crystals get converted; the mass of silver that makes the image visible comes from the crystal itself.

That is why an underexposed film cannot be rescued by longer development. Development amplifies what the light already marked. Where the light never marked anything, there is nothing there to multiply, and pushing the process further cannot invent a signal that was never recorded.

If you have ever wondered why photographers are relaxed about overexposing film and frightened of underexposing it, that is the mechanism underneath the habit.

Two numbers usually get quoted at this point: how many silver atoms a speck needs before development will take it, and by what factor development multiplies the original signal. Both are widely repeated. Neither could be traced to a source that actually states it, so this course is not going to give you either, and you should be slightly suspicious of anyone who does without saying where it came from.

Which leads to the honest part, and it is worth saying plainly because it is rare.

The mechanism of the latent image is still not fully understood, and competing views exist. That is the assessment in the chemistry source this chapter has been quoting, about a process that has been in continuous industrial use since the middle of the nineteenth century – roughly a hundred and eighty years.

It worked long before it was explained, and it is not fully explained now. A course that admits that about silver has earned the right to be believed about everything else.

One practical point falls out of the latent image before we leave it. Because the record is chemically stable and invisible, film is patient in a way nothing digital is. A roll half-shot in March and finished in August carries both sets of frames as latent images the whole time, and the pictures from the spring are made on the same afternoon as the pictures from the summer.

There is no equivalent stage in a digital camera. The charge in a photosite is read out within the exposure's own timescale and written immediately, which is why a memory card full of files behaves so differently from a cassette full of unseen ones. Same measurement, very different patience.

6. 6 From silver to silicon



Now we do the same act a third way, and this is the chapter where the course can finally hold both sides of the comparison to the same standard, because the silicon figures are published and the silver ones largely are not. Where an assumption is doing work inside a number, you will hear me name the assumption.

The starting point is that light striking certain materials frees electrons – the same first move as in the silver halide crystal, arriving in a different material. In a digital camera that material is silicon, and the freed charge, instead of seeding a speck of metal, is collected and counted.

Silicon charges a specific and quite low price of admission. The technical account published by Scientific Imaging on the photoelectric effect in image sensors gives the band gap of silicon as one point one two electronvolts. That is the energy an arriving photon has to be carrying before it can lift an electron out of the valence band and up into the conduction band. Arrive with less and nothing happens whatsoever.

Energy and wavelength are two ways of saying the same thing, so a threshold in one is a threshold in the other. The same source runs the conversion and gets eleven hundred and ten nanometres. Anything longer than that, which means anything deeper into the infrared, does not carry enough energy to promote an electron, and silicon is – in that account's word – transparent to it. Anything shorter, which includes every wavelength your eye can see, can be absorbed, and where it is absorbed an electron and a hole are created.

One photon in, one electron-hole pair out. That is the exchange rate across the visible range, and it is the reason a sensor can honestly be said to count rather than merely to respond. The silver crystal took a photon and started a chain reaction that ran away from the original event. Silicon takes a photon and hands back precisely one unit of charge, then waits for the next one.

Three published numbers describe how well a given sensor does that, and one datasheet supplies all three.

Quantum efficiency is the proportion of photons arriving at a photosite that actually yield a charge carrier the camera can count. Full-well capacity is how much charge a single photosite can hold before it can hold no more. Read noise is the uncertainty introduced by the act of reading the well out – an error measured in electrons, present whether or not any light arrived at all.

Hamamatsu's technical note for the ORCA-Quest qCMOS camera, model C15550-20UP, dated May 2022 and revised in April 2025, gives peak quantum efficiency as eighty-five per cent at four hundred and sixty nanometres, full well capacity as seven thousand electrons, and readout noise as nought point four three electrons rms in standard scan, falling to nought point two seven electrons rms in ultra quiet scan. Its photosites are four point six micrometres square.

Read that noise figure twice. Nought point two seven of an electron. The uncertainty in the measurement is a fraction of the smallest thing being measured, which is another way of saying this instrument can tell four photons from five.

Now set it against what a photograph actually delivers. ISO 12232:2019 fixes K at ten lux-seconds, so at an exposure index of one hundred the average focal plane exposure is one tenth of a lux-second. Carry the green-light photon rate from the opening chapter through that figure, and a photosite four point six micrometres on a side collects on the order of eight and a half thousand photons over the exposure – against a well built to hold seven thousand electrons.

Those two quantities being the same size is not luck. It is what a sensor designed for photography looks like: a bucket sized for the rain that actually falls on it. And carry the caveat with the number, because the arithmetic assumes one wavelength of green rather than a real spectrum, and a real scene is never one wavelength.

One limitation worth stating plainly. Those figures come from a scientific instrument, because scientific instrument makers publish datasheets and consumer camera makers do not. Nobody publishes the full-well capacity of the camera in your hands. That is a genuine hole in what this course can tell you, and it means any confident full-well figure for a consumer camera is a measurement somebody made privately or a number somebody guessed, and you are entitled to ask which.

What can be said beyond the figures is structural, and it is the part that changes how you read a picture.

A crystal of silver halide is a threshold device. Either enough light arrived to mark it or it did not, and development then treats the marked and the unmarked completely differently. There is no partial crystal. The smooth range of tones in a photograph comes from millions of crystals of varying sizes, each making its own yes-or-no decision, averaged across an area.

A photosite on a sensor is not a threshold device. It is a well, and charge accumulates in it, and at the end of the exposure the amount in the well is read out as a number. More light, more charge, a

bigger number. Two wells that received different amounts of light return different values, and the tonal range comes from the values themselves.

So the mechanism is different, and the two run into their limits differently. A silver crystal that has been marked cannot be marked any further. A charge well that has filled up cannot hold any more, and once it is full, every extra photon arriving is simply lost – which is the reason blown highlights in a digital file are so often unrecoverable, and why module two spends a whole chapter on the histogram.

But what is the same in both is the thing this module keeps insisting on. In each case the light arriving is converted into something durable, and the amount converted depends on how much light arrived and for how long. Silver, silicon, and for that matter bitumen on a pewter plate – the recording medium changes, and the measurement does not. ISO 12232:2019 defines its exposure index against $H_{\text{sub m}}$, the average focal plane exposure in lux-seconds, and it is a standard for digital still cameras. The unit that described a French courtyard in 1827 is the unit written into the digital standard published in February 2019.


All of which is why "digital film" is a misleading way to talk about a sensor, and it is worth being clear about the reason, because the phrase sounds harmless.

It is misleading about the two ends of the range, where the threshold crystal and the charge well reach their limits in different ways. And it is misleading in the way it invites you to think of ISO as a film speed you have swapped into the camera – a mistake module two opens by dismantling, because the standard defines an index describing the exposure a sensor was given, not a change in what the silicon is able to do.

The sensor is not film. It is the third thing in this module that measures the same quantity.

For you, the working consequence is a question to ask before you reach for a setting. When a digital frame comes back with a white sky containing nothing at all, that is the well, and no amount of processing afterwards will retrieve a value the sensor never held. When it comes back dark and grainy, that is a different failure with a different remedy. Telling those two apart on the back of the camera is worth more than any preset, and it is the whole subject of module two's chapter on the histogram.

7. 7 The first digital camera



In 1975, an engineer at Eastman Kodak called Steven Sasson built the first digital camera, using a charge-coupled device. Britannica records the year, the man and the technology. What the machine could actually do is worth going through slowly, because every one of its limits is a limit you will recognise.

Amateur Photographer's account of the world's first digital camera gives the specification. The image was one hundred pixels by one hundred pixels – ten thousand pixels in total, which is to say a camera of nought point nought one megapixels. It was black and white. It weighed three point six kilograms. You held it with both hands, and it looked like something from a workshop, because that is where it was assembled.

Now the two numbers that carry this chapter, from the same account.

Capturing the image took fifty milliseconds. Writing it to magnetic tape took twenty-three seconds.

Sit with the ratio for a moment. Catching the light took a twentieth of a second. Keeping it took the better part of half a minute – roughly four hundred and sixty times as long as the exposure that produced it. Each cassette of tape held up to thirty images, and getting one back afterwards meant a further thirty seconds before it appeared on a television screen.

That imbalance is not a quirk of 1975. It is the shape of the entire history you have just watched.

Niépce's plate sat in the window for eight hours or several days, and even then the picture did not exist until a wash of lavender oil and white petroleum had been through it. A sheet of film records its scene in a fraction of a second and then needs a developer, a fixer, a wash and a dry before anyone can look at it. Sasson's camera needed fifty milliseconds of light and twenty-three seconds of tape.

Catching light has never been the hard part. Keeping it is. A modern camera has hidden that from you so thoroughly that it is easy to forget the same division is still running: your camera is still writing a file after the shutter has closed, its buffer still fills during a burst, and when it stops responding it is because the keeping has fallen behind the catching. That is the same bottleneck Sasson was watching in 1975, and it is where module ten picks the story up.

Two more facts about the machine, and then the question it always raises.

The patent on it expired in 2007 – a detail that quietly dates the whole affair, because the protection on the first digital camera was still running through most of the period in which film remained the normal way to take a photograph.

And then there is the question everybody asks when they hear this story, which is what Kodak made of it. A company whose business was film had, on its own premises, a working demonstration that a photograph need not involve film at all.

The answer everyone has heard is a single line: that Sasson was told his camera was cute, and told not to tell anybody about it. You will find it in business books, in conference talks, in a thousand articles about disruption. You will almost never find it with a source attached.

So let me be exact about what this course could and could not establish. The line is attributed to Sasson himself in places that are perfectly reputable, but every trail runs back to interviews he gave decades after the event, not to any document written in 1975 and not to any minute of any meeting. It is recollection, retold so often that its wording has very likely drifted in the retelling. It may be entirely true. It is not evidence, and I am not going to read it to you as though it were.

What is documented is duller and considerably more useful. IEEE Spectrum, in an article by Joanna Goodrich published on the sixth of April 2022, reports that after several demonstrations Kodak's executives said they did not see a market for the camera, and that Sasson was not permitted to speak about it publicly or to show the prototype outside the company. Sasson's own estimate, widely reported and consistent across retellings, was that the technology stood fifteen to twenty years away from rivalling a film camera. On that, he was close to right.

Nor did the company bury the work. Kodak filed for a patent on the electronic still camera in 1977 and was granted United States patent four, one three one, nine one nine in 1978, naming Sasson and his supervisor Gareth Lloyd as the inventors, and Kodak went on to sell a professional digital camera system in 1991. Snopes, examining the specific claim that Kodak hid the invention to protect film sales, rates it a mixture rather than true, and on those grounds.


One caution about the dates in this chapter, since the policy has been stated and should be applied to my own paragraphs. A grant year of 1978 and the expiry of 2007 mentioned a moment ago do not sit comfortably together, because a United States patent granted in that era ran seventeen years from its grant and would therefore have lapsed in the mid-nineteen-nineties. Either those two figures describe different patents in the same family, or one of them is wrong, and I have not been able to establish which.

And the number everybody wants – what the decision eventually cost the company – does not exist in any form this course is willing to state. Kodak's later collapse has a long list of causes and no audited line attributing any part of it to a demonstration in 1975. Anybody who quotes you a figure for that has invented it.

The honest version is better anyway, because it is not a story about executives being stupid. It is a story about a company that could see the thing perfectly clearly and could not see the when. Being fifteen years early is, at the time it is happening, almost impossible to tell apart from being wrong.

What can be said, without any of that, is that the machine worked. In 1975, in a corporate laboratory, a photograph was made with no chemistry in it anywhere – and the measurement being taken was the one this module started with.

8. 8 The chronology, 1851 to 1963



Between Talbot and Sasson there is a little over a century, and the figures that follow come from Philip Greenspun's History of Photography Timeline unless another source is named. They are not a list of gadgets. Each one changes who is able to take a photograph, and where.

Start in 1851, with Frederick Scott Archer and the wet plate collodion process. Two things about it mattered at once. It was much cheaper than the daguerreotype, and it was a negative-positive process, so it gave unlimited reproductions – Talbot's advantage, now at a price people could meet.

It had a condition attached, and the name gives it away. The plate had to be coated, exposed and developed while it was still wet. That meant the darkroom had to travel with the photographer and be standing before the exposure was made. Working away from a studio in the 1850s meant carrying the chemistry to the subject.

The release from that came in 1871, when Richard Leach Maddox introduced the gelatin dry plate – gelatin and silver bromide on glass. The photographer no longer had to coat and develop while the plate was wet. A plate could be bought ready-made, carried dry, exposed whenever, and developed later somewhere else.

That one change is what makes everything after it possible, because it separates buying materials from making pictures. In 1880 George Eastman founded the Eastman Dry Plate Company and began manufacturing dry plates commercially, and the supply chain existed.

Then, in 1888, the first Kodak. Britannica describes it as George Eastman's first handheld camera, with a hundred-exposure roll of paper negatives. Greenspun's timeline describes the same camera as taking a twenty-foot paper roll for a hundred circular photographs.

Those two descriptions are almost certainly the same roll counted two ways – one by frames, one by length – and the second adds a detail the first omits.

A museum with the object on a shelf settles it. The Smithsonian's National Museum of American History, cataloguing its own Original Kodak camera under accession twenty-three thousand five hundred and ninety-eight, records that the camera took hundred-exposure rolls giving circular images two and five-eighths of an inch in diameter – about six point seven centimetres. Circular, confirmed, and measured.

One wrinkle, because the policy applies to helpful sources as well as unhelpful ones. The Smithsonian's wording says rolls of film, where Britannica and the Greenspun timeline both say paper negatives. The 1888 camera is generally the paper one, with film arriving the year after, so the label is probably using film in its loose everyday sense rather than contradicting anybody. The round frame and its diameter are firm. That one word is not.

Why a circle at all? Because a lens throws a round image, and cutting a rectangle out of a circle throws away everything outside the rectangle. Printing the whole circle wastes none of it – and it quietly avoids the corners, which is where a simple lens of 1888 was performing at its very worst.

Which is a fact about your camera too. Every rectangular photograph you have ever taken is a piece cut out of a circle of light, and the four corners of your frame are the places where the glass is working hardest and succeeding least.

In 1889 Eastman improved the Kodak, using film rather than paper. Glass to paper to film, in eighteen years, and the camera has become something you buy loaded, shoot, and send away.

Colour is where the dates get slippery, and it is another place where two reputable sources point in different directions.

Britannica gives Kodachrome as 1935, invented by Leopold Godowsky Junior and Leopold Mannes, producing colour transparencies suitable for projection and reproduction. Greenspun's timeline gives 1936, describing Kodachrome as the first multi-layered colour film. Both are reputable and they do not agree.

They may well both be right, because they may be dating different events – an invention and a commercial release are not the same moment, and a film can be finished in one year and on sale the next. Until that is resolved from a primary source, the honest form of the sentence is the vaguer one: Kodachrome arrived in the mid-1930s, from Godowsky and Mannes, at Kodak.

Colour also has a chapter that runs three decades earlier than Kodachrome, and it contains the strangest piece of engineering in this entire module.

The National Science and Media Museum's history of the Autochrome names the inventors as the brothers Auguste and Louis Lumière, who presented their research into colour photography to the Académie des Sciences in 1904. Commercial manufacture of autochrome plates began in 1907; the first public demonstration of the process took place on the tenth of June 1907; and the plates

reached buyers in Britain that October. Research, manufacture, demonstration, sale – four separate events, and the reason this chapter keeps insisting you ask which one a bare year refers to.

The process itself is very nearly absurd, and the museum describes it in detail. The plate was covered in microscopic grains of potato starch, sieved to keep only those between ten and fifteen microns across – thousandths of a millimetre – then dyed red, green and violet and spread over a glass plate coated in sticky varnish, at something like four million grains to the square inch. Over the grains went a panchromatic emulsion, sensitive right across the visible spectrum. Light on its way in had to pass through the coloured starch first, so every grain acted as a tiny filter over the patch of emulsion sitting behind it. Process the plate to a positive transparency, hold it up to a window, and the light coming back out through those same dyed grains recombines into a full colour picture.

No colour dyes in the image itself. No three exposures, no filter changes between frames, no registration to get wrong. One plate, one exposure, and a random mosaic of dyed vegetable starch doing the work that three separate negatives had needed before it.

For a course about measurement, the Autochrome makes the cleanest possible point about colour. Colour is not an extra thing a camera does on the side. It is the same measurement of the same quantity, taken separately behind red, green and violet windows, and then put back together. Which – allowing for silicon, and for a tidy printed grid in place of scattered potato starch – is precisely what the sensor in the camera beside you is doing right now. The idea was working in 1907.

The slippery-date problem attaches to the camera that defined the following century as well. The timeline used here gives the Leica as marketed in 1924, from Oskar Barnack at Leitz, and calls it the first high quality 35mm camera. It is very commonly dated 1925 instead.

Announcement against general sale is the likely explanation, and that is the lesson worth taking from every one of these disputes. A date in photographic history is often the date of an announcement, or of a patent, or of the day a thing reached a shop shelf – and those are routinely three different years. When you see a bare year with no event attached to it, you do not yet know which one you have been given.


The chronology closes with instant photography: Polaroid selling instant black-and-white film commercially in 1948, and the first colour instant film in 1963. A finished picture, with no darkroom and no laboratory anywhere in the transaction – the end point of everything since Archer's wet tent, and the last stop before Sasson.

Read the century as a single movement and it has one direction. In 1851 a photograph required a chemist, a vehicle and a tolerance for wet glass. By 1871 the materials would wait for you. By 1888 the camera came loaded and went back in the post. By 1948 the picture arrived in your hand where you were standing.

Every step removes a condition on where and when a photograph can be made, and none of them changes what a photograph is. That is the test to apply to anything sold to you as a revolution in photography, including the ones arriving this year: ask which condition it removes. If the honest

answer is none, it is a convenience, and conveniences are worth having – but they are not a change in the thing itself.

9.9 The idea the course rests on



Here is the sentence this module has been building, and it is the sentence the remaining nine are built on.

A photograph is a measurement of light over an interval.

Every word in it is doing work. Measurement, because a quantity is being determined, and ISO 12232:2019 states that quantity's unit: lux-seconds, an illuminance multiplied by a time. Light, because that is the only thing a camera ever responds to – never the subject, never the scene, only the light that has left it. Over, because the measurement accumulates. And interval, because it has a beginning and an end and a length, and can never have a length of zero.

We have now watched that same act performed three ways. Bitumen of Judea hardening on pewter in a French window in 1826 or 1827. Silver halide crystals gathering atoms of metallic silver in a pattern too faint to see. Charge collecting in silicon wells on a bench at Eastman Kodak in 1975. Three materials, three centuries, three completely unrelated chemistries and physics – and one measurement, in the same unit, of the same thing.

Which brings us to the second half of the idea, and to why this module had to come first.

Everything else in photography is control of that measurement.

That is not a slogan. It is a genuine claim about what the controls on your camera are, and it makes a prediction you can test on any camera you pick up. The aperture sets how much light per second the lens admits, which is the rate. The shutter speed sets how long the collection runs, which is the interval. The ISO setting governs how the measurement that resulted is indexed and rendered. Focus and lens choice determine where each bundle of light lands – the question Ibn al-Haytham's hole

answered in the 10th or 11th century. Lighting changes what is there to be measured in the first place.

There is nothing else. Every button, every dial, every menu in every camera made since 1839 falls into one of those groups, and once you can sort them that way, an unfamiliar camera stops being a wall of controls and becomes a small number of questions with answers in different places.

So that is how the rest of this course is laid out, and each module is one part of that sentence taken seriously.

Module two is the exposure triangle: aperture, shutter and ISO as a single system, including the point at which the arithmetic stops holding and a metered ten seconds is not ten seconds any more.

Module three is every kind of camera, which is to say what genuinely differs between them and what only appears to. Module four is every control on the camera, one at a time. Module five is lenses and how focus works – where the light lands, which is the hole question with glass in it.

Module six is lighting: changing the measurement by changing what is being measured. Module seven is composition and seeing, the only module about the part the camera cannot do for you.

Module eight is photographing the world, which is where all of it meets actual subjects in actual conditions. Module nine is film, and module ten is digital from sensor to finished picture – silver and silicon, each given the full treatment this module could only sketch.

Nine modules, and one sentence underneath all of them.

You will not need to remember much of this chapter's history. You will need the sentence. When a picture comes out wrong and you cannot think why, the question to ask is not what setting to change. It is which part of the measurement went wrong – the rate, the interval, or what was there to measure – because that tells you which control you are looking for, and there are only ever three places to look.

Sources and credits

Primary sources

- ISO 12232:2019 – the exposure index, and $K = 10$ lux-seconds.
- BIPM, SI Brochure, 9th edition (2019), Table 4 – the lux as a coherent derived unit, lumens per square metre; and the exact defining constants fixed on 20 May 2019.
- Nobel Prize in Physics 1921 citation, nobelprize.org – Einstein, for the discovery of the law of the photoelectric effect.
- National Science and Media Museum – the portable camera obscura, and Autochrome: Auguste and Louis Lumiere, presented to the Academie des Sciences in 1904, manufactured from 1907, first publicly demonstrated 10 June 1907.
- Smithsonian National Museum of American History, accession 23598 – the 1888 Kodak's circular images, 2 5/8 inches in diameter.
- Hamamatsu ORCA-Quest qCMOS C15550-20UP technical note (May 2022, rev. April 2025) – 85% quantum efficiency at 460 nm, 7,000 electron full-well capacity, 0.43 and 0.27 electrons rms read noise, 4.6 micrometre pixels.
- IEEE Spectrum, Joanna Goodrich, 6 April 2022 – Steven Sasson's 1975 camera and Kodak's response; US patent 4,131,919, filed 1977, granted 1978.
- Musee d'Orsay – Arago's announcement to the Academie des Sciences, 7 January 1839. Bodleian Talbot project and Royal Society archive AP/23/19 – Talbot at the Royal Institution, 25 January 1839, and his paper of 31 January 1839.

Not regulated financial advice.