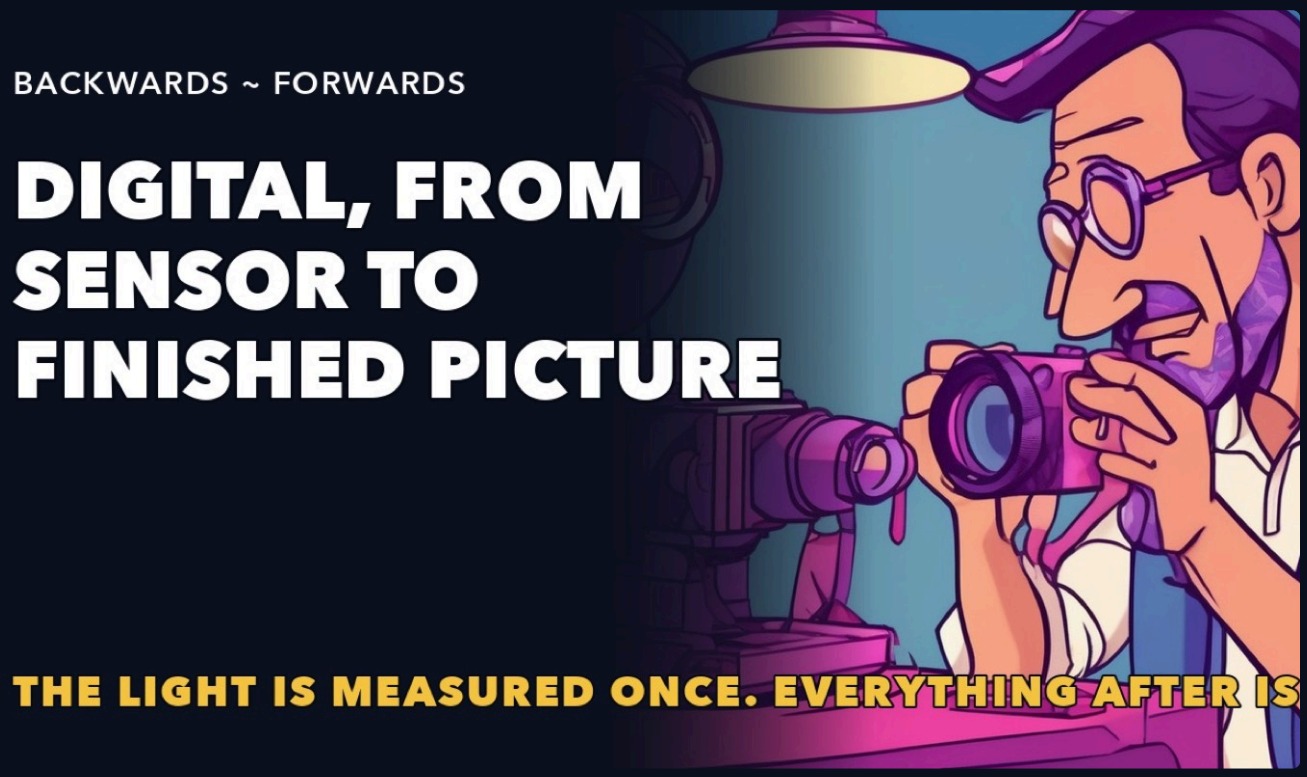
Forwards or Backwards

https://forwardorbackwards.com



Learning Photography Course - 10. Digital, from Sensor to Finished Picture

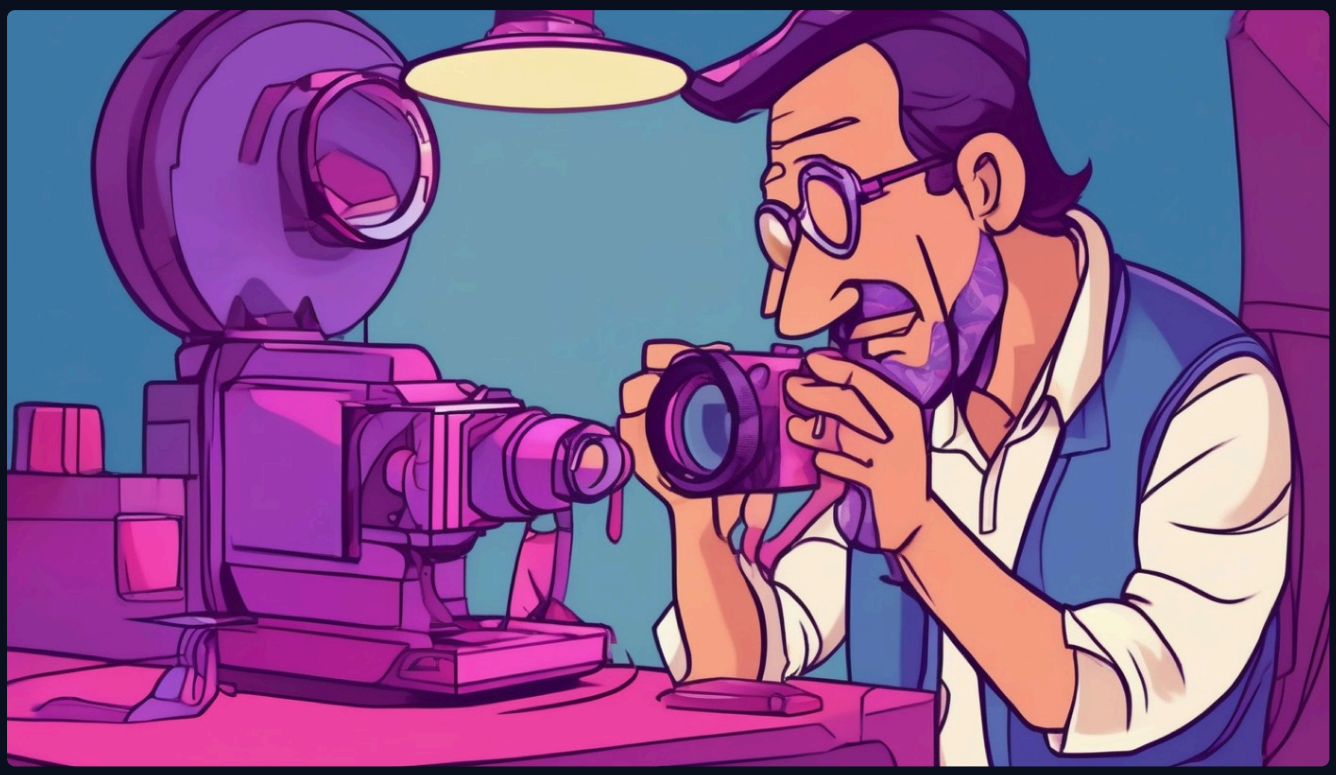
Watch on YouTube: https://www.youtube.com/watch?v=TLt5CYdBV7A

Chapters

1. 1 The sensor
2. 2 Making colour
3. 3 Noise and dynamic range
4. 4 Raw and JPEG
5. 5 Getting the pictures off
6. 6 Not losing them
7. 7 The catalogue
8. 8 A sane editing order

9. 9 Colour management and export

1. 1 The sensor



Module nine ended with a sentence worth keeping: film is a chemical process with a fixed response; digital is a measurement with an adjustable one. This module follows that measurement all the way, from light arriving on a sensor to a picture leaving your computer.

A sensor is a grid of millions of tiny light-sensitive sites called **photosites**. Each one works like a bucket in the rain. During the exposure, arriving light frees electrons in the silicon, and the photosite collects them in a small store of charge, often called its **well**. More light, more electrons. At the end of the exposure, each well is measured, and the count becomes a number.

A well can only hold so much. The most it can collect is its **full-well capacity**, and when it is full it simply stops counting. Any extra light is lost. That is what a clipped highlight is on a digital camera: **not a bright value, but a bucket that overflowed**. Film rolls off gently into its highlights; a sensor hits a hard ceiling.

There have been two main ways of reading those wells. In a **CCD** sensor, charge is passed along from site to site and measured at the edge. In a **CMOS** sensor, each photosite has its own small circuit that turns charge into a voltage where it sits. CMOS came to dominate cameras because that local circuitry can be built with ordinary chip-making processes, draws less power and can be read out faster.

Two refinements matter today. A **back-side illuminated** sensor flips the silicon so the wiring sits behind the light-collecting layer instead of in front of it, so less light is blocked. A **stacked** sensor

goes further, putting processing circuits on a separate layer underneath, which makes readout much faster.

And readout speed is not a detail. Most sensors are read one row at a time, top to bottom. If something moves while that happens, each row records it at a slightly different moment. A spinning propeller comes out bent into curved blades; a fast pan leans every upright. That is **rolling shutter**, and faster readout is how modern sensors shrink it.

2. 2 Making colour



Here is something most photographers never think about. **A sensor is colour-blind.** A photosite counts electrons. It has no idea whether they were freed by red light or blue.

So colour has to be added with filters. In front of the photosites sits a mosaic of tiny coloured filters, one per site, so that each photosite only receives red, green or blue light. The pattern almost every camera uses was patented by **Bryce Bayer of Eastman Kodak**. His patent, granted in nineteen seventy-six, describes it.

In the **Bayer array**, **half of the filters are green**, and the other half are split between red and blue. Why so much green? The patent's own reasoning is about the eye: we see fine detail mainly through brightness, and green carries more of the brightness signal than red or blue. The patent even notes our relatively limited ability to see detail in blue.

That leaves a problem. Each photosite measured only one colour, but every pixel in the finished picture needs all three. The missing two are **worked out from the neighbours**. This step is called **demosaicing**, and it is worth being plain about what it means: **two thirds of the colour information in every pixel was estimated, not measured.** Good demosaicing is so accurate you never notice. Poor demosaicing shows up as false colour on fine patterns, such as the weave of a shirt.

There are alternatives. Fujifilm's X-Trans sensors use a different, less regular filter pattern. Sigma's Foveon sensors stack three layers that respond to different colours at every site, so no neighbour-guessing is needed. And some cameras are made as **true monochrome**, with no colour filters at all;

every photosite measures brightness directly, which gives more detail and more light to each site, in black and white only.

Finally, the **anti-aliasing filter**. Fine repeating patterns can clash with the regular grid of photosites and produce moiré, a false rippling pattern. A thin filter in front of the sensor blurred the image very slightly to prevent it. As sensors gained more photosites, the grid became fine enough that many makers removed the filter and kept the extra sharpness, accepting the occasional moiré.

3. 3 Noise and dynamic range



Digital noise has several sources, and the first one is not the camera's fault at all. **The light itself is noisy.**

Photons do not arrive in a steady stream. They arrive at random, like raindrops, and if you count them over a short time the count wobbles. Physics tells us how much: the wobble is about the **square root of the count**. That is called **shot noise**. Here is a worked example. A photosite that collects ten thousand photons has a wobble of about a hundred, so the signal is about a hundred times the noise. A photosite in deep shadow that collects only a hundred photons has a wobble of about ten, and the signal is only ten times the noise. **A hundred times less light, ten times worse.** That is why shadows are noisy before any electronics are involved.

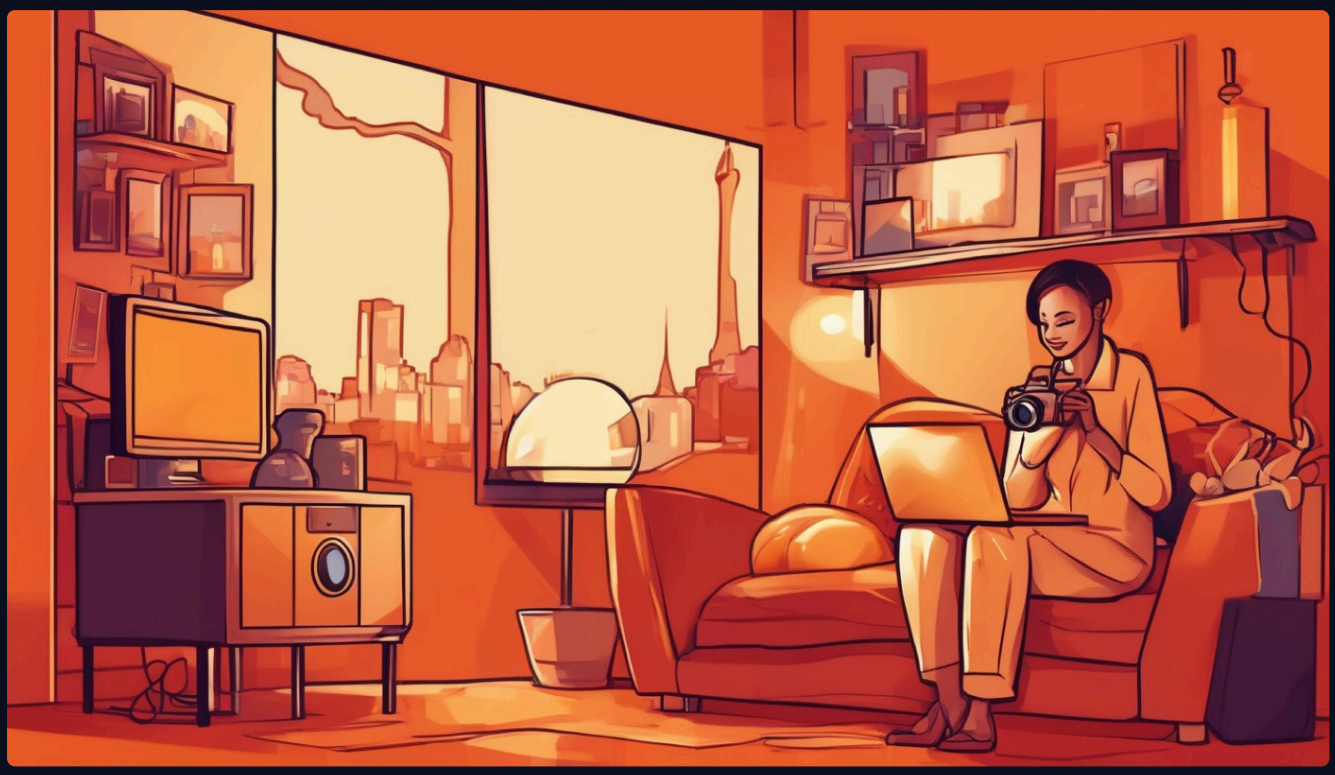
Then the electronics add their own. **Read noise** is the small random error added every time a photosite is measured, the same whether the site is full or nearly empty. It matters most in the darkest parts of the picture, where the real signal is tiny. **Pattern noise** is different: a fixed structure, such as faint bands or a grid, that repeats from frame to frame. It is less common now, but it is the kind that shows when you lift shadows hard.

Dynamic range is the span between the brightest thing a sensor can record before a well fills, and the darkest thing it can record above its own noise. It is often quoted in stops. It is worth saying that dynamic range is **measured** under defined conditions, and figures from different tests are not always comparable, so treat a single number with care.

Many modern sensors read each photosite in more than one way at once, at a low gain for the highlights and a higher gain for the shadows, and combine the results. That is **dual gain**, and it is why some cameras have what looks like a second base ISO, a higher setting that is unusually clean.

And a popular myth. **More megapixels does not, by itself, mean more noise.** Each smaller photosite is noisier, but when the picture is viewed at the same size, many small photosites together gather roughly the same light as fewer large ones. Compare images at the same output size, not pixel by pixel.

4. 4 Raw and JPEG



A **raw file** is, roughly, the measurement itself: the number from each photosite, still behind its single colour filter, before demosaicing, before white balance, before any tone curve. Alongside it the camera stores the settings it would have used. That is why a raw file opened without adjustment looks flat and dull. **Nobody has made any decisions about it yet.**

A **JPEG** is a finished picture, and to make it the camera's processor makes every one of those decisions for you, in a fraction of a second. It demosaics. It fixes the white balance. It applies a tone curve and a colour style. It sharpens, and it reduces noise. Then it throws information away twice. First it reduces the depth of each value: the JPEG standard's baseline process, the one cameras use, stores **eight bits for each colour channel**. Then it compresses, discarding fine detail the eye is least likely to miss.

Bit depth is easy to put into numbers. Eight bits gives **two hundred and fifty-six** levels per channel. Raw files commonly hold twelve or fourteen bits: twelve bits is **four thousand and ninety-six** levels, and fourteen bits is **sixteen thousand, three hundred and eighty-four**. Where does that matter? Not in a picture viewed as it came out of the camera. It matters when you **push** a file, lifting shadows by several stops or pulling back a bright sky. With few levels to spread, smooth gradients break into visible steps, called banding.

Raw files themselves come in versions. **Uncompressed** raw keeps every value as it was recorded. **Lossless compressed** raw packs the same values into a smaller file, the way a zip file does, and gives back exactly the same data when opened. **Lossy** compressed raw is smaller still, and discards

information permanently, usually information you would struggle to see. Know which one your camera is set to.

So is raw always better? No. **JPEG is the right answer** when the picture needs to go somewhere now, when you are shooting thousands of frames of an event, or when the camera's own colour style is exactly what you want. Many photographers record both at once: the JPEG to use today, and the raw in case a picture needs more later.

5. 5 Getting the pictures off



The pictures live first on a **memory card**. Most cameras use one of three families. **SD cards**, the familiar stamp-sized card. **MicroSD**, the fingernail-sized version used by phones, action cameras and drones. And **CFexpress**, a faster, larger card used by many high-end cameras for high-resolution bursts and video.

The numbers printed on a card are easy to misread, so it is worth knowing what they promise. The SD Association defines **speed classes**, and each one guarantees a **minimum sustained writing speed**. A card marked **Class ten** will write at no less than ten megabytes a second. **U three** guarantees thirty. The **video speed classes** say it most plainly: **V thirty is thirty megabytes a second, V ninety is ninety**. That is a floor for recording without dropped frames, not a top speed. The bigger number on the label, the one advertised in large type, is usually a peak reading speed, which tells you how fast the card might copy to your computer, not how reliably it records.

To copy the files, you have two main choices. A **card reader** takes the card out of the camera and plugs into the computer. A **cable** connects the camera itself. A good reader is usually faster and saves the camera's battery. A cable avoids handling the card at all.

A cable also allows **tethered capture**: the camera sends each frame straight to the computer as it is taken, so it appears on a large screen immediately. In a studio, that means a client can see the picture, and focus can be checked properly, before the model has changed position.

Then there is **wireless**. Almost every modern camera can send pictures to a phone. It is convenient for sharing a handful of images, and it is the right tool for that job. For a full card, it is still the slow

option: moving hundreds of large raw files over the air is far slower than moving them through a reader. Use wireless for the few, a reader or a cable for the many.

6. 6 Not losing them



Film negatives sit in a drawer for decades. Digital files disappear in an instant, and almost always for the same few reasons: a card formatted before it was copied, a single drive that fails, a laptop that is lost. This chapter is about making that impossible, and it starts with a habit.

Ingest is a deliberate step. Copying pictures off a card should never be an afterthought done while the kettle boils. Copy the card to your computer, then check that the files are there and open. Only then is the card safe to reuse. Many photographers go further and copy to two places during ingest, so that the pictures exist twice before the card is ever cleared.

Next, **naming and folders**. The test is simple: will you find this picture in ten years? The structure that survives is usually the most boring one. A folder for each year, and inside it a folder for each shoot, named with the date first so they sort in order, followed by a few words about what it was. Rename files at ingest if you like, but decide on one pattern and keep it. **A clever system you stop using is worse than a plain one you always use.**

Then backups. The rule most often quoted is **three, two, one**. The United States National Cybersecurity and Communications Integration Center set it out in two thousand and twelve: **keep three complete copies of your information, two of them on different media, and one of them stored off-site.** In practice for a photographer, that can be the working copy on your computer, a second copy on an external drive, and a third somewhere else entirely: another building, or a cloud service. A copy in the same room as the computer does not survive the same fire, flood or burglary.

Finally, the cards themselves. **Format a card in the camera**, not on the computer, so the camera sets it up the way it expects. Format it only after the pictures are safely copied. If a card starts giving errors, retire it rather than trusting it with one more shoot. And if pictures are ever lost, stop using that card immediately; recovery software can often find files that have been deleted, but only until new pictures are written over them.

7. 7 The catalogue



There are two kinds of software for looking at your pictures, and the difference shapes everything after ingest.

A **browser** shows you what is in a folder. Move the folder and it simply shows you the new place. Nothing is remembered about the pictures except what is stored in the files.

A **catalogue** keeps a database about your pictures: where each file is, the ratings you gave it, the keywords you added, the edits you made. Most catalogue software edits without touching the original file at all. It stores your adjustments as instructions and applies them when you view or export. The original raw file stays exactly as the camera wrote it, which means **every edit can be undone, years later**. The price is discipline: move files outside the catalogue and it loses track of them.

A catalogue is only as useful as the information in it, and the honest advice is to add only what you will actually keep doing. A **star rating** for quality. A **flag** for picks and rejects. And a few **keywords** for the things you will one day search for: people, places, subjects. A handful of keywords applied to every shoot beats an elaborate scheme abandoned after a month.

Then comes the part that most improves your photography, and it is not editing. It is **selecting**. Suppose a day's shooting produced four hundred frames, and the story needs twelve. The way through is in passes, not in one heroic sitting.

On the first pass, move quickly and reject only the obvious failures: out of focus, eyes shut, a missed moment. Do not agonise.

On the second pass, flag the pictures that genuinely work. Where several frames show the same moment, choose one and let the others go. **Near-duplicates are where most of the four hundred go.**

On the last pass, look at the survivors together, as a set. Which twelve tell the story, with variety of distance, of pace and of mood? Here a strong picture sometimes goes, because it repeats another.

What you end with is small and good. And you have learned something that no setting teaches: **which of your pictures actually work, and why.**

8. 8 A sane editing order



Editing software offers dozens of controls, and the order you use them in matters, because each step changes what the next one is working on. Here is an order that avoids undoing your own work.

First, the profile and the white balance. The profile is the starting interpretation of the raw data: how colours and contrast are rendered before you touch anything. Choose it first, because every later adjustment sits on top of it. Then white balance, which sets what counts as neutral. Get neutral right and every colour decision afterwards is made against the truth, not against a colour cast.

Then exposure, then contrast, then the tone curve. Exposure sets the overall brightness. Contrast sets how far apart the tones sit. The tone curve is the fine instrument: it lets you shape exactly which tones get brighter or darker. Work from coarse to fine, and you avoid fighting a curve that was set before the exposure changed underneath it.

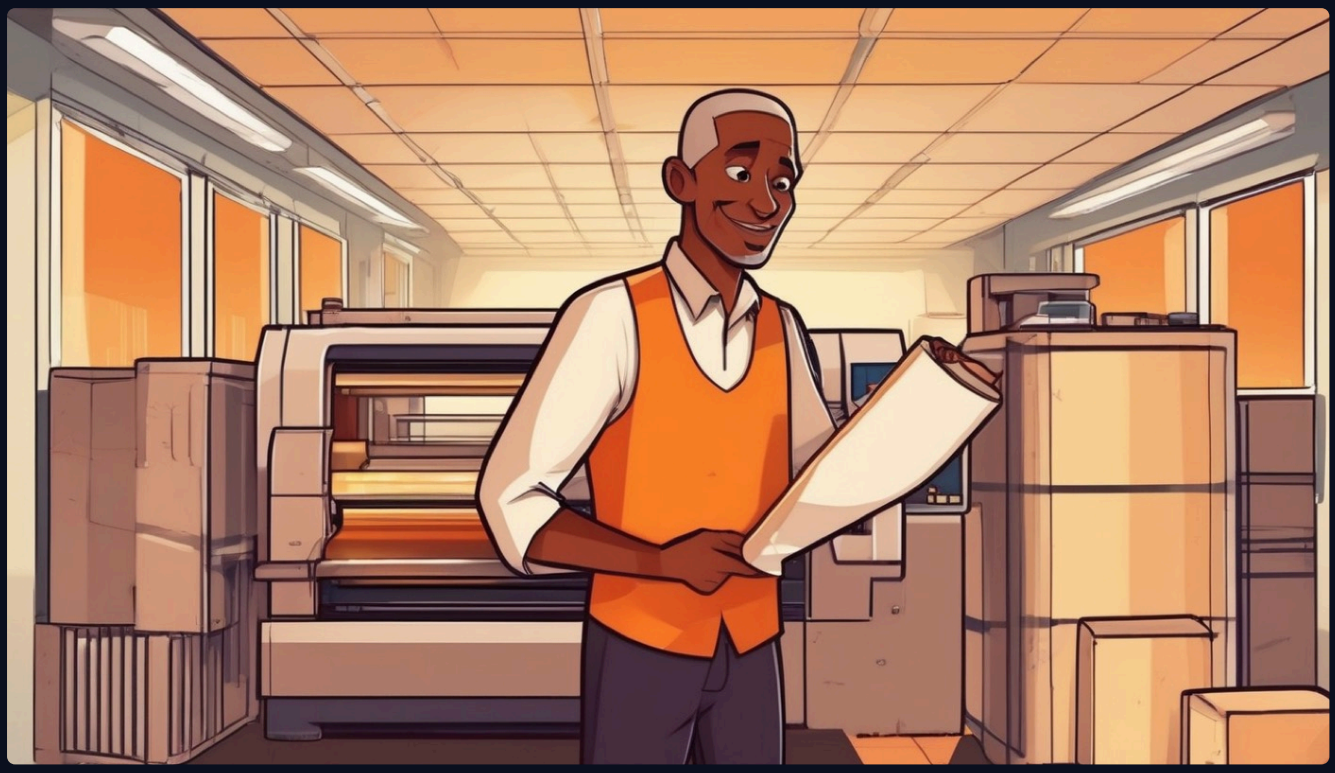
Then local adjustments. Everything so far applied to the whole picture. Now you work on parts of it: brightening a face, darkening a sky, drawing attention away from a distracting corner. Modern software builds **masks** for this, often automatically finding a subject or a sky. This is the digital descendant of dodging and burning in module nine, done with a mouse instead of a card.

Then colour. Hue, saturation and luminance controls let you change individual colours: calm a green that is too loud, deepen a blue. Calibration adjusts how the primaries themselves are interpreted. Grading adds an overall mood, warmer shadows or cooler highlights. Colour comes late because it is the most taste-driven step, and it depends on everything before it.

Last, noise reduction, and then sharpening, in that order. Sharpening makes edges crisper, and it cannot tell an edge from noise. Sharpen first and you sharpen the noise, making it more visible for the noise reduction to fight. Reduce noise first, then sharpen what is left.

None of this is a law. But **when an edit starts to feel like a fight, it is often a sign that steps were done out of order.**

9. 9 Colour management and export



A picture's numbers only become colours once you say which **colour space** they belong to. The same three numbers mean a different colour in a different space.

The one to know first is **sRGB**, defined in nineteen ninety-six by engineers at Hewlett-Packard and Microsoft, and later made an international standard. Its specification defines exactly where its red, green and blue primaries sit, a white point called D sixty-five, and a tone response that works out at roughly gamma two point two. It is the default for the web, and **if you do not know where a picture is going, sRGB is the safe answer**.

Adobe RGB, nineteen ninety-eight, is wider, and its own specification shows exactly where. **Its red and blue primaries are identical to sRGB's. The difference is almost all in the greens and cyans**, which matters for some print work. **Display P3**, used by many phones and laptops, is also wider than sRGB. **ProPhoto RGB** is very wide indeed, and is used as a working space inside editing software rather than for delivering pictures.

Then, why the print did not match the screen. Usually the screen was too bright. The sRGB specification assumes a reference display of **eighty candelas per square metre**; Adobe RGB's assumes **a hundred and sixty**. Many screens are set far brighter, so pictures edited on them print dark. **Calibrating** the monitor with a measuring device sets its brightness and colour to a known standard, and it is the single biggest step towards prints that match.

Sharpening has one more stage, at the end. **Output sharpening** depends on where the picture is going. A small image for a screen needs a little. A print needs more, because ink spreads on paper

and softens edges.

Before publishing, **resize** to the size actually needed, and check the **metadata**. A photograph can carry the camera's settings, your name, and sometimes the exact location where it was taken. Decide what should leave with it, and remove what should not, particularly the location of your home.

And the last check: **soft proofing**. Good software can show on screen how the picture will look in a particular printer, ink and paper, so you can adjust before the print is made rather than after.

That is the whole journey, from a photon to a finished picture. **The light was measured once; everything after that was a decision.**

Sources and credits

Primary sources

- US Patent 3,971,065, 'Color imaging array', inventor Bryce E. Bayer, assignee Eastman Kodak Co, filed 5 March 1975, granted 20 July 1976 Green at every other position (half the elements); the eye's limited ability to discern blue detail.
- Michael Stokes (HP), Matthew Anderson (Microsoft), Srinivasan Chandrasekar (Microsoft), Ricardo Motta (HP), 'A Standard Default Color Space for the Internet - sRGB', Version 1.10, 5 November 1996 (w3.org) — primaries, D65 white, transfer function, 80 cd/m² reference. Later standardised as IEC 61966-2-1, which was not read.
- Adobe Systems, 'Adobe RGB (1998) Color Image Encoding', Version 2005-05, May 2005 — primaries (green 0.21, 0.71), D65 white, gamma 2.19921875, 160 cd/m² reference white.
- SD Association, 'Speed Class' (sdcard.org), last modified 31 October 2023 — C10 10 MB/s, U3 30 MB/s, V30 30 MB/s, V90 90 MB/s as minimum sustained sequential write.
- ISO/CIE 11664-2:2022, CIE standard illuminants.

Not regulated financial advice.