Pushing the Limits

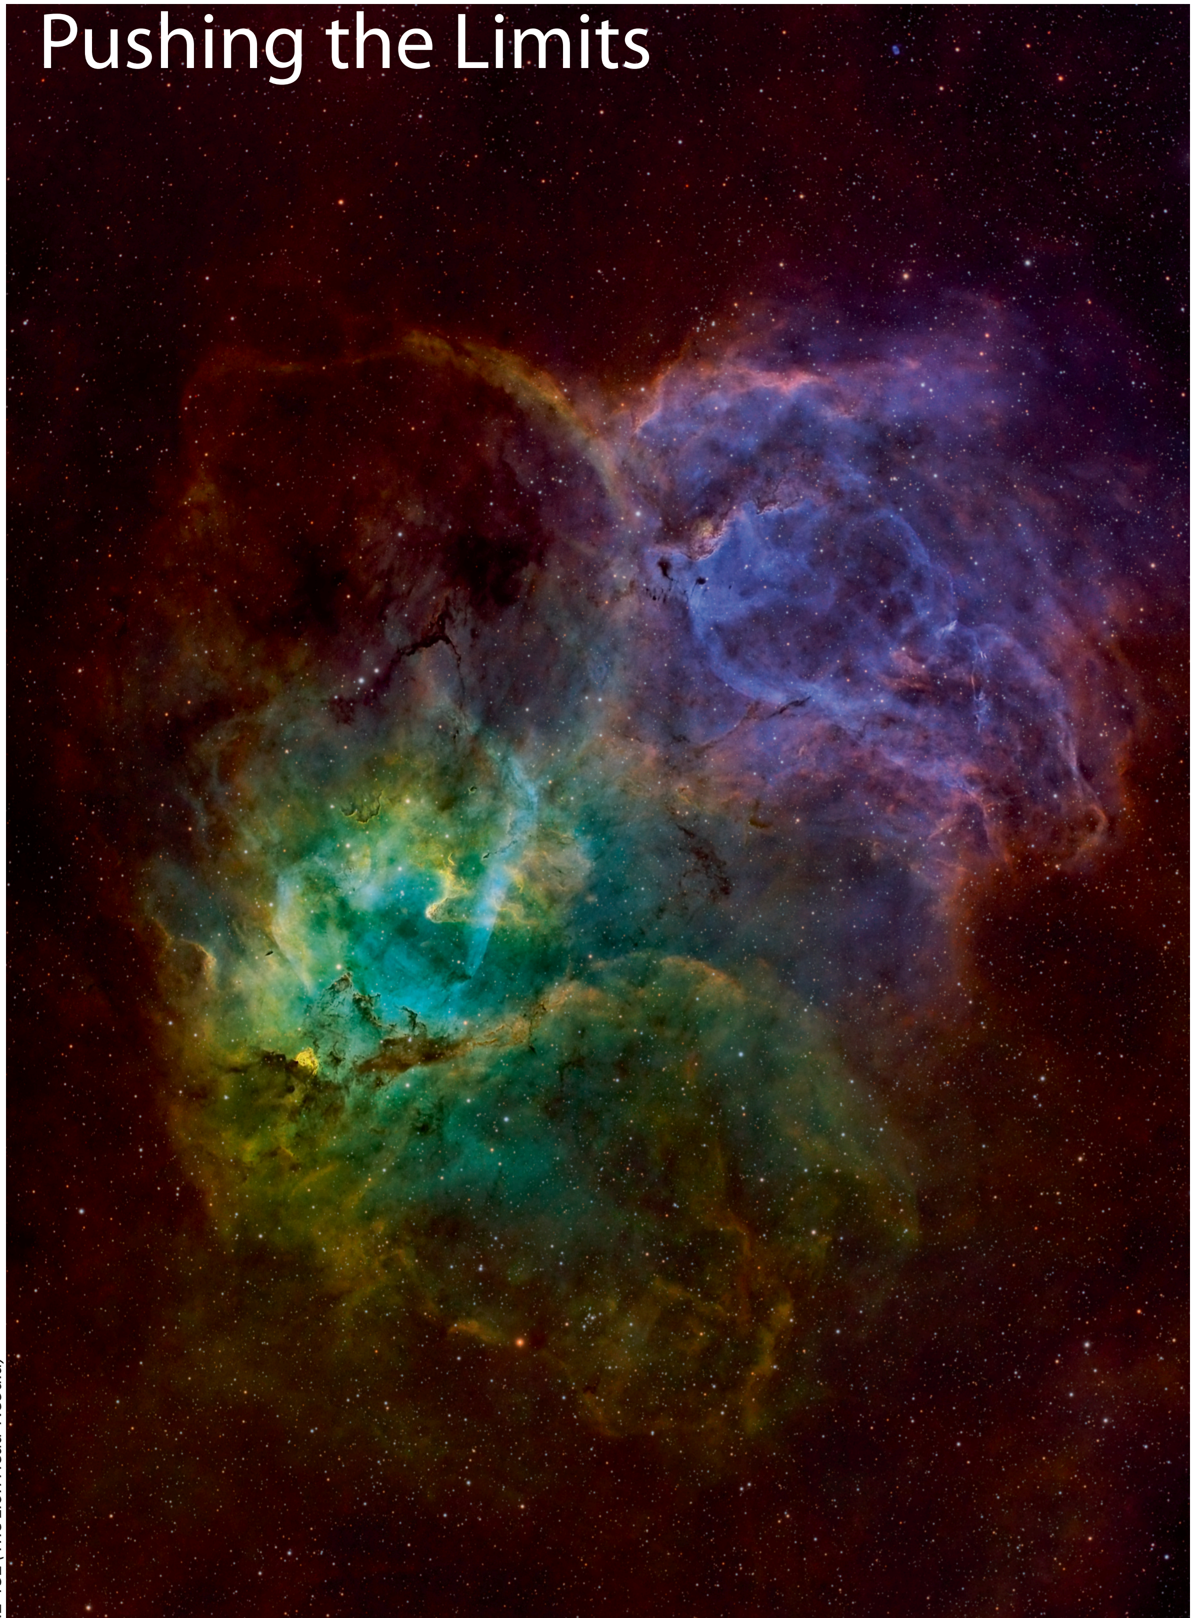

SH2-132 (The Lion Head Nebula)

Down to Earth

A reminder of the significant challenges that astrophotographers face and how they might be met.

It is incredible to think that amateurs are routinely imaging deep-sky objects at a quality level to rival some professional observatories of twenty years ago. Even so, astrophotography is not a "walk in the park", and for many of us, it is a constant crusade against forces mainly beyond our control. We operate at the very limit of our equipment, and yet making beautiful images is possible. These challenges form a significant part of the final reward and recognition. Your mission, should you choose to accept it, is to boldly go where no sane human has gone before.

This chapter reminds us of the real-world challenges and the high-level techniques we use to overcome them. It touches upon many technical topics without too much detail, which would obscure the high-level message. These essential topics have a more rigorous treatment in later chapters.

Location

The location where we image from greatly impacts what we can achieve. As well as the limits imposed by our particular imaging horizon, the weather, light pollution and geography influence what we image, how we image and how long it will take. It is important to recognize this limitation at an early stage, for though it is exciting to engage in a frenzy of retail therapy, it may not necessarily result in better images. In my case, while improved technique, equipment, and considerable patience have dramatically improved my image quality, my location is now the limiting factor and fights further quality improvements.

Imaging the Invisible

Astrophotography is mostly about taking images of exceedingly dim subjects sprinkled with surprisingly bright points of light and usually in the presence of considerable light pollution. To give you an idea of the challenges, it is illuminating (pun intended) to compare the imaging conditions of a typical studio portrait with those of an emission nebula. Using a sensor's gain, quantum efficiency specifications, and the resulting image file's pixel value allows one to estimate how many photons land on a pixel (technically, a photosite) in each scene. In the case of the studio-flash portrait, about 25,000 photons hit a pixel during the brief flash. In the case of the nebula example, the image values of a faint glowing hydrogen cloud indicate just 40 photons (on average) emitted from the nebula land on a pixel during a 20-minute narrowband exposure. That makes the hydrogen cloud about 10 *billion* times darker, or, in photographic terms, over 33 stops. One would have to image for 1,000,000 seconds to get an equivalent pixel value. At the same time, the brightest stars saturate/clip sensor pixel photosites in a fraction of a second. Challenging times are ahead!

Object Brightness

The apparent magnitude of an object infers the object brightness. This measure is one of the various standard parameters in the stellar and deep-sky catalogs. However, this measure has to be used with caution; there are several terms that we use to imply brightness, namely luminosity, flux and magnitude. Luminosity relates to the total light energy output from a star, flux is a measure of energy over a unit area, and the intensity reduces with viewing distance, obeying the inverse square law. It is possible to measure flux with a telescope and a sensor, but it is not a convenient measure for comparing object brightness.

The apparent magnitude of an object provides a convenient unitless measure of its relative intensities from Earth. Magnitude is a back-to-front logarithmic unit; a one-unit increase is 2.5x less bright, and a 5 unit increase is 2.5^5x (100x) less. For instance, Sirius has a magnitude of -1.47, and the faintest object observable from the Hubble Space Telescope is about +31, or 2.4×10^{13}x dimmer. Fig.1 shows a range of objects at different apparent magnitudes and what is typically required to detect them.

The note of caution referred to earlier relates to the apparent magnitude of an object. This is based on the amount of light emitted from the whole object and for large objects, such as some galaxies or nebulae, the "surface" appears fainter, compared to a smaller object of the same magnitude. As such, one cannot convert the apparent magnitude of an object into a simple exposure calculator.

Dynamic Range

We all appreciate that deep-sky objects are demanding to image and, on the whole, exceedingly dim and come in all shapes and sizes. We are also familiar with the general concept of reducing or increasing exposure in response to object intensity. In astrophotography, however, it is often the case that the various objects in an image will have very different intensities, i.e., the subject has a high dynamic range. With a fixed telescope aperture, exposure is controlled by sensor gain (or camera ISO) and exposure duration. The best setting is elusive, and for scenes that include objects with a large intensity range, an exposure may have clipped highlights, lose faint details or both.

This is not a new problem; in the days of film, if one over-exposed transparency film, the highlights washed out, whereas, with color and mono negative film, underexposing shadows was of more concern. Transparency film can distinguish tones over about a five stop range (32:1), whereas modern sensors have a dynamic range of about 12 stops (4096:1) and are ~16x more efficient at capturing photons. Even so, capturing the faintest and brightest details of most deep-sky scenes in a single exposure is impossible. Sensor models have different dynamic range capabilities and may change with operational settings. Understanding these is useful when selecting and using a camera for deep-sky imaging. Dynamic range is discussed with other sensor specifications in a later chapter.

Sensors, Warts and All

No sensor is perfect, and their imperfections become apparent when we operate them at their limit. For example, they typically waste 10–30% of the photons that fall onto the sensor surface and the electronics, which convert the electrical charge into a voltage and convert to a digital value, introduce their own errors. These include consistent conversion errors as well as random variations. Worse still, some sensors generate unwanted artifacts, the most obvious of which is amp glow.

Amp Glow

Amp glow is most noticeable in CMOS sensors and varies between models. The unwanted signal is generated within the sensor chip, on account of the integrated conversion electronics and while successive sensor designs generally reduce the issue, it is still apparent in long exposures. For example, in fig.2 the amp glow intensity in this long exposure exceeds the narrowband image intensity. Removing the amp glow and the pixel-to-pixel inconsistencies requires precise exposure calibration, using perfectly matched dark frames with the same sensor temperature and exposure time (fig.3). In practice, the necessary precision favors temperature-regulated, refrigerated sensors and relegates the role of DSLR or mirrorless cameras to light duties. Once, the high price of CCD-based astro cameras made this a significant investment, but today's modern CMOS designs are more affordable and, once you have used one, there is no going back.

visibility	apparent magnitude	# objects brighter	example / notes
human eye urban sky	-1	1	Sirius (-1.5)
	0	4	Vega
	1	15	Saturn (1.5)
	2	50	Jupiter (-2.9 to -1.6)
	3	<200	Andromeda Galaxy (3.4)
human eye dark sky	4	500	Orion Nebula (M42)
	5	1,600	Uranus (5.5-6.0)
	6	4,800	Eagle Nebula (M16)
binoculars with 50-mm aperture	7	14,000	Bode's Nebula (M81)
	8	42,000	Crab Nebula (M1)
	9	121,000	M43 Nebula in Orion
typical visual 8-cm aperture	10	340,000	NGC 4244 Galaxy
	11	-	Little dumbbell (M76)
typical visual 15-cm aperture	12	-	beyond Messier Catalog
	13	-	Quasar 3C 273
typical visual 30-cm aperture	14	-	Galaxy PGC 21789 nr. Pollux
	15	20,000,000	IC 4617 Galaxy nr. M13
10-cm refractor, sensor,10x30s suburban sky	16	-	faint star in image with simple stacking, about 20,000 times more sensitive than by eye alone
10-cm refractor, sensor10x 300s suburban sky	18	-	faint star in image with simple stacking, about 1,000,000 times more sensitive than by eye alone
rural sky	20	-	typical *background magnitude* in rural area
Hubble Space Telescope	31	-	galaxies 13.3 billion light-years distant

fig.1 This table highlights the limits of perception for the aided and unaided eye over a range of conditions and indicates the number of objects within that range. The advantage of CCD/CMOS imaging over an exposure of 5–50 minutes is overwhelming. For Earth-bound imaging, the general sky background and noise, indicated by the shading, will eventually obscure faint signals, from about magnitude 18 in suburban areas. The Hubble Space Telescope operates outside our atmosphere and air pollution and, at its limit, it detects magnitude 31 objects. Its sensitivity is approximately 150,000x better than an amateur setup. The James Webb Telescope's mirror collects 6x more photons than the HST!

fig.2 Some modern CMOS sensors have amp glow, which becomes increasingly obvious and may exceed Ha emissions (as here) with long exposures. With care, however, calibration using temperature- and time-matched files removes this glow, leaving behind a little random noise.
Precise calibration is difficult to achieve with consumer digital cameras as it is impossible to regulate the sensor temperature without modification and even in a stable ambient, the sensor temperature increases with use.

fig.3 This is the same view as in fig.2. This image is a combination of 40 calibrated and registered exposures. The calibration process removes the amp glow and constant pixel-to-pixel conversion errors (and at the same time, things like vignetting and dust spots too). Combining 40 exposures reduces the random noise by about 6x, allowing one to stretch the image to reveal the faint Ha details. The amp glow area does, however, have a slightly higher noise level, on account of its accompanying shot noise.

Sensor Noise

It is easy to become fixated on sensor noise and specifications. Unfortunately, these specifications lead to misleading model comparisons, particularly when their pixel sizes are different. The term "noise" is often used as a general term to refer to any unwanted (error) signal. Under normal circumstances, it would be invisible, but in astrophotography, it requires extreme image stretching to reveal faint details that increase local contrast 100-fold or more and reveal hidden issues. In a nutshell, there are several error mechanisms:

- pixel gain differences (constant)
- pixel offsets (constant)
- pixel variation (random)
- pixel variation (temporary patterns)

Image calibration is the key weapon in our arsenal to combat sensor pixel to pixel inconsistency. This is something that conventional photographers seldom do knowingly. The image calibration process evaluates your sensor, operating under the same conditions as the image exposure. It uses exposures taken in the dark (dark frames) and of an even illuminated target (flat frames) to correct the error in each pixel in each image exposure. However, this only fixes the constant errors; things like hot pixels, pixel shifts, gain variations in the conversion electronics, and the mean level of thermally-generated electrons. Image calibration does not reduce random pixel variation. This random noise also reduces the ability of a sensor to record very bright and dim subjects simultaneously (I like to think of noise as fog – nothing in fog is high-contrast). The sensor's dynamic range is a computed specification from a photosite's electron capacity (full-well depth) and its read-noise level. Each time the read noise doubles, the effective dynamic range halves.

Random noise mechanisms include sensor read noise and the randomness of any thermally generated electrons during the exposure, which increases with temperature and duration. The noise from these thermally generated electrons is most significant in DSLRs and mirrorless cameras as sensors heat up with operation and, without cooling, reach 30 °C or more with extended use. For best results, it is best to use a dedicated astro camera with a built-in cooler. I operate mine at -15 °C, achievable in most ambient conditions. By chance, my Fuji camera and astro camera use the same Sony sensor, and refrigeration reduces the thermally generated random noise by about 15x. Without cooling, the thermally generated noise from a 5-minute exposure is obvious, at about 4x the sensor read noise. The two equivalent images in fig.4 compare the thermal noise levels at two sensor temperatures.

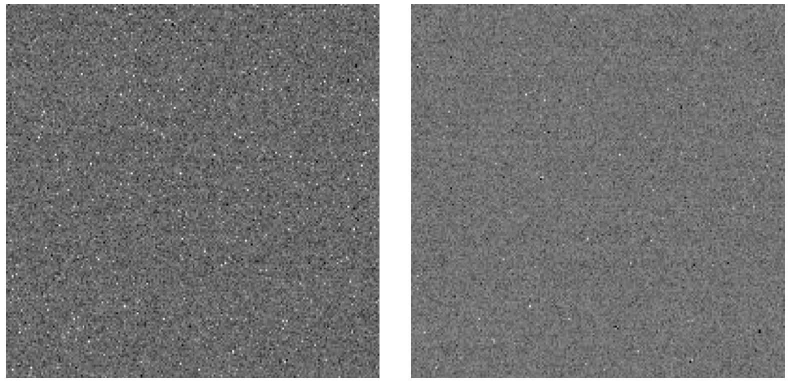

fig.4 The magnified image on the left is a 10-minute exposure with the sensor at 5 °C. On the right, another at -15 °C. In both cases, the mean level was subtracted from the image and a similar stretch applied to each, to reveal the dark noise. Some DSLR and mirrorless cameras report their sensor temperature over the USB connection and in practice, can rise above 30 °C.
This is thermal noise and increases with exposure and doubles with every ~5 °C rise in sensor temperature. Most cooled astro cameras are capable of reducing the sensor temperature 35–40 °C below ambient, which lowers the dark noise level by 128–256x.

Without diverting ourselves into too many technical details at this time, the key takeouts are to use a cooled camera to reduce the thermally generated current (and hence noise) and calibrate every image exposure with matched dark and flat frame exposures. But, unfortunately, the elephant is still in the room.

The Problem with Light Pollution

Light pollution is a growing issue for many of us. Urban living conspires with personal safety requirements and encourages households and local authorities to illuminate our surroundings at night. Even those in uninhabited deserts suffer each month as the Moon illuminates the sky. Each of its forms affects our imaging differently, as do the coping strategies. Even with the same amount of illumination, light pollution may vary from night to night, depending on the amount of dust, water vapor and aerosols in the atmosphere. These scatter and reflect the incident light in all directions. A few weather forecast applications predict atmospheric transparency and infer light pollution to some extent. I have found some of the best imaging periods occur after it has been raining, as the air is usually cleansed and has less backscatter.

Many believe the objection to light pollution is due to the general obscuring glow in an image. This glow is typically brighter near the horizon, making it even more objectionable and, practically, may limit an exposure duration. This, however, is not the real issue since, with care, image processing removes the glow and gradient, leaving behind a uniformly neutral dark background with faint details of the deep-sky object. The real demon is still there, however, lurking in the shadows (literally). This is where we come face to face with the immovable force of physics, and this crucial and inescapable property requires a little explanation.

Shot Noise

Photons behave as discrete particles and are emitted randomly because of the laws of physics (they are also waves, but we are not going into that just now). We normally do not perceive that randomness but know the average emission rate sets the subject brightness. Over a suitable duration, a photosensor converts enough photons into electrons and converts them into a digital value at the end of the exposure. It is useful to think of photons like raindrops; we understand the difference between a shower and a downpour, but in each case, the raindrops land randomly over the back yard, and you cannot predict where the next one will land, or when. Photons behave similarly, and even in an image of an evenly-illuminated uniform target, each sensor pixel captures a slightly different number of photons (just as the number of raindrops hitting two precisely equivalent targets is different). The unwanted randomness between pixel values is noise (or shot noise, to be precise). These combine with other sources of noise caused by interference and sensor electronics.

The shot noise level increases with the number of captured photons (or raindrops). There is an important distinction here; we are not talking about light level (brightness) but the total exposure. If 10,000 photons land on a sensor pixel over 0.001 or 100 seconds, it makes no difference to the shot noise. Even more interesting, it is the same deal if 10,000 photons are captured over 100 x 1-second exposures and then added together. The difference is noticeable and the example in fig.5 compares two short but slightly different exposures of the same bright scene. (These images were carefully processed to isolate the random noise from the mean exposure.)

Thankfully, the shot noise does not increase at the same rate as the exposure; if the captured photons increase by 4x, the noise level only increases by 2x. Conversely, quartering the exposure halves the shot noise. In fact, a simple equation predicts the shot noise for a given exposure level. The same equation holds true for other signals; for example, the photons from the object and the thermally generated electrons in the sensor.

fig.5 Shot noise increases with exposure on account of the nature of light. These magnified images have had their mean levels adjusted and equally stretched to reveal the noise. The left image is a 0.01-second flat exposure and the right is 0.04 seconds. Increasing exposure increases shot noise, which can dominate the other noise sources in each exposure.

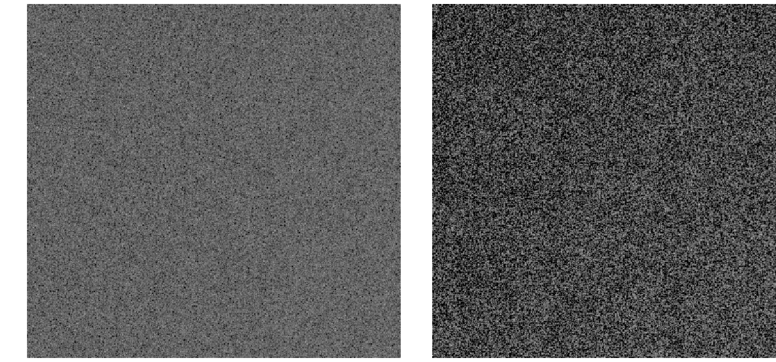

$$\textit{shot noise level} = \sqrt{\textit{signal level}}$$

So how does this relate to light pollution? If we consider two identical exposures of a dim object, taken from areas of low and high light pollution, we can assess the shot noise associated with the sky background by noting the pixel value increase between the dark and light frame for an area of blank sky in each case. At the dark site, the faint signal mostly competes with sensor noise and its shot noise. For the light-polluted site, while the general average glow is removed during image processing, the random shot noise is left behind and the faint signal mostly competes with sky shot noise.

Competing with Light Pollution

There are two strategies to lessen the effect of light pollution: reduce the amount of light pollution hitting the sensor and increase the exposure. As you have guessed already, increasing the exposure is an all-round good thing to do and we will consider that last, in a broader context. Let us look at the various forms of light pollution and the potential specific workarounds. There are several approaches to reduce the impact of light pollution on our astrophotography:

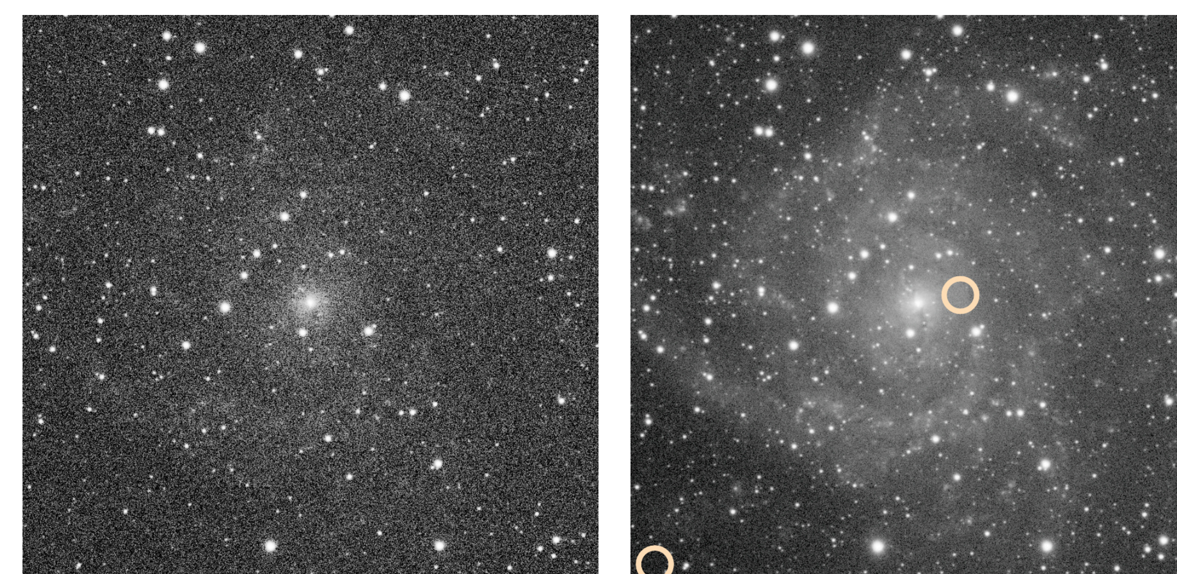

figs.6,7 These closely-cropped stretched images of this target indicate the sample areas used in the analysis in fig.8. The image on the left (fig.6) is a single, calibrated 120-second exposure through a red filter. It is hard to see the target's structures in a single frame and the image on the right (fig.7) is a combination of 53 calibrated and aligned images. (Fig.7 also demonstrates the benefit of integrating multiple exposures to improve dynamic range and noise levels.)

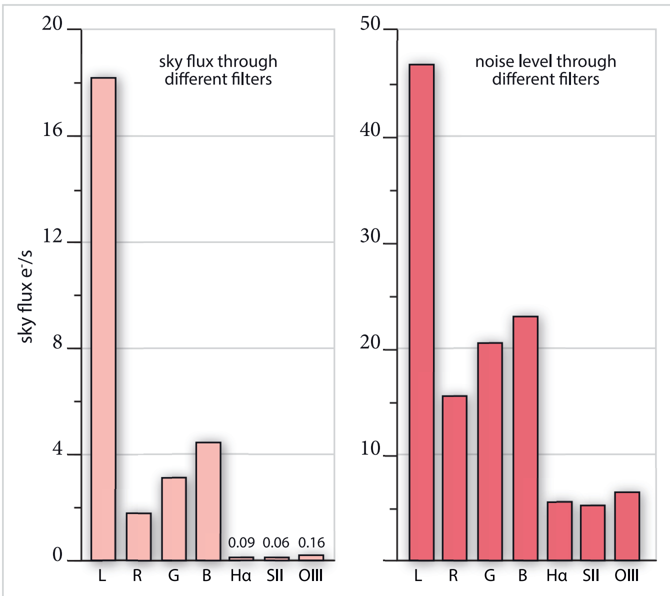

fig.8 Using the sample areas in fig.7 for differently-filtered exposures, together with the calibration file values, it is possible to break a single exposure's pixel values into their constituent components. On the left-hand graph are the image pixel values before calibration. One can see how different filters reduce the sky flux value for the same exposure duration.
The magnified graph on the right correspondingly shows the effective noise level in a 120 second exposure (SQRT(sky flux x 120 secs) for each filtering strategy. One can see that the RGB filtering is already halving the residual noise and the narrowband filters reduce the background noise level by about 8x,

- light-pollution filters
- separate RGB filters on a monochrome sensor
- narrowband filters (on color or mono sensors)
- image from a darker site
- (capture more photons)

These all improve the signal-to-noise ratio but in different ways. In the first four, we reduce the amount of light pollution hitting the sensor and, more significantly, reduce the substantial shot noise that comes with it.

Filtering Strategies

If your local street lighting uses traditional sodium lamps, a color image of what appears to be a dark sky reveals a muddy orange color as the light bounces back to Earth from all the airborne contaminants. These emission lamps emit distinct colors, and although there are several different types, their primary outputs are at 590 and 600 nm. The less popular mercury vapor lamps emit over a broader spectrum but with strong components at specific wavelengths in the violet, blue, green, and yellow-orange colors. The coping strategies in these conditions take advantage of the specificity of the light-pollution colors. We use unique thin-film (dichroic) optical filters to block the unwanted artificial colors and pass the others.

In this instance, light-pollution filters precisely exclude the main street lamp colors. There are many on the market; most exclude the two sodium and four mercury lamp emissions yet pass the common nebula emission colors with remarkable efficiency. Common sizes are 1.25- and 2-inch sizes and are increasingly made in smaller tailored packages, to insert inside a consumer digital camera. Many telescopes have an internal 48-mm thread in the focus tube or field flattener to accept a 2-inch filter. The established manufacturers are Hutech (IDAS), Baader, and Astronomik, with new entries from SkyTech and STC. If you hold one of these up to a street lamp, it all but disappears, and while these filters significantly reduce specific wavelengths they do not eliminate light pollution as a whole. At the same time, they often improve the color balance (and sky gradient) of an image taken on a color camera. As a result, they are most commonly used with conventional (color) cameras, but not exclusively.

Monochrome sensors present more flexible filtering options. These sensors do not have Bayer filter arrays and have a broader and higher overall sensitivity. However, they are rarely used without filtration since they are sensitive to UV and IR light, and it is typical to use them with discrete filters, the choice of which affects the transmitted light pollution. A natural color image requires one to image through Red, Green, and Blue filters (RGB) or perhaps Sloan (g',r',i') filters, for a slightly different effect. Some sets, as well as the filter mosaic in front of a color sensor, have overlapping filter

responses (similar to human vision). In contrast, others intentionally omit the yellow light wavelengths associated with low-pressure sodium street lighting and yellow light from the deep-sky object.

The big guns in the monochrome arsenal are narrowband filters. Their incredibly selective transmission passbands exclude all light other than a specific emission nebula color. The principal nebulae emissions are fortunately in the blue-green and deep red regions and do not correspond to those from urban lighting. This removes most of the light pollution and the associated shot noise. As a result, many inner-city astrophotographers almost exclusively do narrowband imaging of emission nebula.

New Lighting Developments

The efficiency benefits of LED lighting are hard to ignore. Many countries are steadily replacing older vapor lamp systems with white LEDs. This is a worrying problem for astrophotography as, for all its benefits, the backscatter from the ground appears worse. Even more significant is the spectrum of the light itself. An LED is monochromatic, with the wavelength set by the manufacturing doping process. A white LED is a combination of separate red, green, and blue LEDs, and consequently, its output covers a broad spectrum and is more difficult to filter out. LED color is defined by its wafer doping; there are no color standards and, unlike a sodium vapor lamp, whose color is determined by quantum physics, there is considerable variation between manufacturers.

Our traditional coping strategies are less successful with white LED urban lighting; while some recent light-pollution filters, for example, the IDAS D2 LPS filter, have been designed to remove the intense blue spectrum associated with LEDs, they still pass the lower intensity green and red wavelengths. At the same time, they are reducing useful light from the deep-sky object and on color sensors, may affect color balance. Unfortunately, this intense blue light pollution passes through the blue filter of an RGB filter set, but again, the discriminating nature of narrowband filters will continue to work well.

Time and Place

Sometimes the only recourse is to image from a darker site, with less light pollution, or when conditions are at their best. For instance, the Moon is a regular nuisance and, for about 10 days each month, the reflected light from its surface floodlights the sky and degrades image backgrounds (especially when the atmospheric transparency is poor). In addition, moonlight has a broad spectrum, peaking in the orange-red wavelengths. The obvious coping strategy is to avoid imaging during this period, or when the Moon is below the horizon. One may also consider imaging with narrowband filters when the Moon is up or choose targets away from the Moon in the opposite direction or the two northern quadrants.

Exposure, Exposure, Exposure

Finally, our universal solution to reduce random noise (from light and sensor) also improves the effective dynamic range and allows us to capture faint fuzzies. This golden rule is to capture more photons. Each time you quadruple the captured photon count, the shot noise level only doubles. At the same time, the sensor read noise in any exposure stays the same. Any combination of the following will help:

- use a bigger aperture (diameter, not f/ratio)
- take longer exposures
- take and combine multiple dithered exposures

I chose my words carefully; these three are about capturing more light by increasing light intensity or effective exposure duration. It is not the same as increasing the gain (camera ISO). This does not affect the amount of captured light and potentially reduces the photon count. The temptation is to shorten the exposure duration at high gain (ISO) to avoid over-exposure. Sensor gain may affect image noise for other reasons; in some models, especially those with 12- or 14-bit ADCs, increasing internal sensor gain reduces the quantization error, which adds to the read noise. This is a marginal improvement, compared to the overwhelming shot noise from typical light pollution.

Aperture

Each time you double the aperture of an instrument, its photon capture capability quadruples. In astrophotography, the aperture of a telescope is the diameter of the primary mirror or glass lens. This is initially confusing to photographers who often incorrectly use the same term to refer to f/ratio. A telescope is ultimately a light-gathering device. If you compare the pupil in your eye with a telescope's aperture, it is easy to appreciate the benefit. Big is better. Over a given period, a 2,000-mm focal length f/10 with a 200-mm aperture captures 8x more photons than a 200-mm focal length f/2.8 lens with a 71-mm aperture. (Incidentally, an iPhone has an aperture of about 2.5 mm, capturing

6,400x less light.) Telescope prices increase rapidly with aperture size, and there is an economic and physical limit to how big you can go.

Exposure

The second option is to take a longer exposure. Doubling the exposure time doubles the photon count and fills up the sensor photosites. At the same time, this makes full use of the sensor's dynamic range capability and registers faint details out of the noise. Exposure duration and aperture size are two sides of the same coin, each trying to use the full range of the sensor. In one case, it does so by increasing light intensity and, in the other, recording for longer. However, both have the same practical limit, which is the onset of unwanted highlight clipping. Exposure is, therefore, almost always a compromise between the extremes of registering faint details and highlight clipping, and the optimum setting depends on an individual's artistic goals. For instance, I like my images to have plenty of color variation and, for that reason, my exposure durations use the full range of the sensor and just clip on a few pixels corresponding to the brightest star cores. However, as a colored object brightens, it loses its apparent color saturation and when bright image pixels are stretched, the result is usually an unsightly white blob. Therefore, I double-check that my exposures are not over-exposing a bright galaxy core or nebulae before committing an exposure sequence.

Multiple Exposures

In the quest to increase the number of captured photons through intensity and duration, the sensor design sets an upper limit, typically in the range of 15,000–80,000 electrons. Unfortunately, this is insufficient for photographing faint fuzzies in the presence of light pollution. The solution is to combine multiple exposures. Conceptually, we can either add successive exposures, creating a "super sensor" with unlimited electron capacity or, more usefully, average multiple exposures, usually in a 32-or 64-bit file format.

There are subtle differences between the strategies; quadrupling the intensity allows one to capture the same number of photons in a quarter of the time and halves the random dark noise in the exposure. Doubling the exposure count and doubling the exposure duration also have slightly different results; in the former, we have multiple read-noise contributions and in the latter, just one. For that reason, the optimum exposure strategy is to use the largest telescope aperture that you can afford/lift, with an exposure duration that makes full use of the sensor capacity but keeps clipped pixels to a minimum and then I combine lots of them.

Combining multiple exposures increases dynamic range, allowing us to capture faint fuzzies and bright stars in the same final image. We have only touched upon this earlier, but as sensor noise levels increase, it reduces our ability to distinguish different tones. If the sensor uncertainty is 2 electrons, it effectively halves the ability to discern discrete electron levels. Now the fun part; if we consider four identical deep-sky exposures, they are not truly identical, because of random noise. This is due to the photons and electrons occurring randomly in time. Averaging these four exposures also averages (smooths) out the randomness, reducing it by 2x. If our sensor uncertainty was 2 electrons, in the averaged version, it is now 1. The simple act of averaging four exposures doubles the dynamic range (increased by 1 stop). If we average 64 exposures, we improve dynamic range by 8x or three stops (bits). With sufficient exposures, it is possible to have an effective dynamic range that exceeds the number of levels in a 16-bit file, hence the recommendation to combine exposures to 32- or 64-bit files.

Optical Resolution and the Environment

Advertising and consumer pressure tempt us to over-indulge in telescope purchases for astrophotography. Many optical and physical properties distinguish a good telescope from a bad one. Just like with any other pursuit, knowing what is important is the key to making the correct purchasing decision. In the case of resolution, the certainty of optical performance, backed up by physical equations, is a beguiling one for an engineer. I have to frequently remind myself that these are only reached under perfect atmospheric conditions (which I have yet to encounter). The final image quality in all forms of photography has many factors, and the overall performance is a consequence of all the degradations in the imaging chain. It is easy to misinterpret the image and blame the optics for any defect. The simple objective truth is that many amateur telescopes resolve finer detail than an astro camera can record and the atmosphere allows. Assuming well-corrected optics, the resolution of a telescope improves linearly with the size of the aperture diameter. Unlike conventional photography, astronomers are more interested in angular resolution, and we conveniently express telescope resolution in arc seconds. As apertures increase, the shimmering of our atmosphere, what we call astronomical seeing, sets a ceiling on resolution performance.

Astronomical Seeing

Astronomical seeing is an empirical measure of the optical stability of our atmosphere and, conveniently, is also measured in arc seconds as it affects angular resolution. Air refracts light like glass and its index is affected by pressure, temperature, and humidity. Turbulence causes rapid localized changes in air density and parallel beams deviate through refraction, causing stars to shimmer or blur when viewed through a telescope. At any one time, the light beams pass through adjacent small air pockets with different refractive indices. Astronomers look through about 20 miles of the atmosphere (looking straight up) and double that, closer to the horizon. Turbulence is most significant in the denser air near the ground or from tiny convection currents within the telescope tube. A series of brief exposures or videos show how a star jumps about, with some badly blurred and others remarkably sharp. During a long exposure, the photons from this twinkling image accumulate onto the sensor, creating a smeared star image. How long is long? I ran an experiment and measured the star sizes for thousands of exposures taken at different exposure durations (fig.9). With 0.1-second exposures, the stars were consistently small, increasing rapidly and in variation with 0.5- and 1-second exposures. The trend changed with longer exposures, with less variation and diminishing size increases. This suggests that, in this experiment at least, image shifts occurred around a sub-1-second time-frame, rather than fractional or double-digit exposure values.

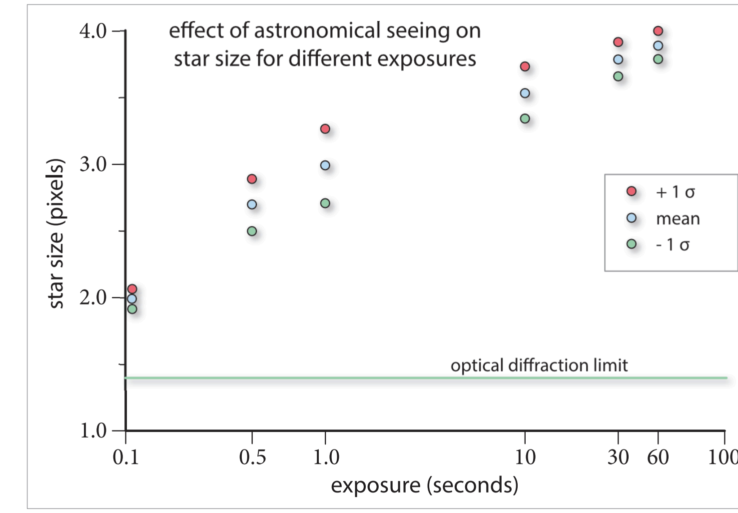

fig.9 This shows the star size and size variance for exposure times from 0.1 to 60 seconds. In this case, astronomical seeing is a short-term abrupt anomaly in the atmosphere that occurs every second or so, rather than more frequently.

Astronomical seeing changes with location, season, time, and weather conditions and can be forecast to some extent. Some of the weather apps targeted at astronomers include seeing forecasts. These apps are constantly changing and evolving, and a search will establish the current ones that apply to your location. For a prime site, a seeing condition of 0.5 arc seconds is possible, but values in the range of 1.5–3.5 are more typical. More often than not, the prevailing seeing conditions will limit any telescope's resolution. The table in fig.10 shows the theoretical limits of visible light resolution for several popular amateur telescope sizes compared to typical seeing conditions. It is sobering to realize the limitation imposed by typical seeing conditions through the atmosphere is equivalent to a telescope with an aperture of ~3 inches (~75 mm).

The overall imaging system resolution is a combination of the separate resolutions for the optics, sensor and atmosphere (and tracking errors during the exposure). As a result, it is always worse than the weakest link in the chain. In a typical system, these might be 1, 3, and 2 arc seconds, respectively, making the system resolution of 3.7 arc seconds almost 4x worse than the telescope's.

Coping with Seeing

Our options are limited. Seeing conditions are sensitive to inconsistencies in the dense atmosphere closest to Earth and generally improve with altitude and proximity to large expanses of water (due to the moderating effect on thermal generation). The mountain observatories in Hawaii and the Canary Islands are good examples of prime locations. Seeing conditions also change with the season and the amount of daytime heating. The local site also has an immediate bearing; it is better to image in a cool open field than over an expanse of concrete that has received a day's sunshine. Likewise, telescopes are best acclimatized to their surroundings for several hours before use to reduce convection currents and stabilize the lens and mirror dimensions. Astronomers choose remote sites not to be anti-social; they need to find high altitudes, clear skies, low light pollution, and low air turbulence.

Some also resort to "lucky imaging", a technique borrowed from planetary and solar imagers. This takes thousands of very short exposures, ruthlessly rejects the blurred images, and combines the rest. The result improves the system resolution closer to the optical limit. However, it goes against some of our golden rules for reducing noise and for it to work effectively, requires a bright target, small pixels (for resolution), large apertures (for resolution and photon capture), sensors with very low read noise, and a capacious disk drive.

Tracking, Focus, and Resolution

It is worth remembering that seeing conditions are not the only resolution-robbing effect. Tracking errors from mechanical flexure, alignment, and manufacturing tolerances cause a star to drift or shift during the exposure, blurring the result. Autoguiding and complex multi-parameter mechanical and atmospheric refraction modeling can help, if done well, or make matters worse. These are complicated subjects and they have their own chapters. As far as the fundamentals are concerned, best practice includes carefully polar-aligning a telescope mount, making sure everything is as rigid as possible, and using autoguiding to correct any tracking errors that the mount cannot correct for itself. In practice, I consistently achieve 0.5 arc second accuracy with autoguiding, better than typical prevailing seeing conditions.

Not surprisingly, poor focus also ruins image resolution. Unlike conventional photography, it is not a one-time thing; the mechanicals are less stable; tube length and optics change with temperature, and heavy mirrors move around. The focus position also changes with the filter in the optical path. Accurate focus, repeated checks, and adjustments are necessary; you guessed it, focusing has its own chapter.

Summary

Image noise is the major hurdle in the path of achieving quality results. The mind-bogglingly low light levels from the object are competing with larger forces. In many ways, image noise has more impact than focus and tracking errors, which are more process and mechanical in nature. Unfortunately, image noise is often not entirely understood, and many practitioners concentrate on noise sources within the sensor and forget about the random nature of light itself.

There are no shortcuts; the signal-to-noise ratio (and hence quality) of an image is improved by:

- using a bigger aperture (captures more photons)
- taking longer exposures (without clipping)
- taking and combining more exposures
- cooling the sensor
- imaging from a darker site
- using filters to exclude light pollution

In my semi-rural location, a typical image has 20–50 hours of exposure with a cooled astro camera over several weeks or months on account of the weather. A year's effort may result in half a dozen images. The prevailing weather conditions dictate the terms of this hobby and are an unavoidable consideration. Keen amateurs are relocating their system to a remote dark site in increasing numbers, in search of more frequent clear nights and the chance to operate at an alternative latitude.

I have started using dual-camera systems to make the most of each rare opportunity. One of the exciting developments in acquisition software overcomes a limitation of existing interface standards and allows dual-imaging systems to operate more efficiently in tandem on a single mount. This enables doubling-up exposures or, with careful alignment, side-by-side images for a wide-field mosaic image.

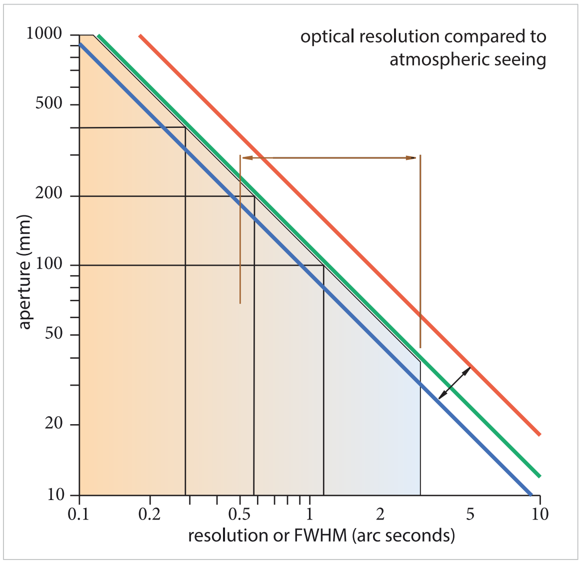

fig.10 The chart above indicates the diffraction-limited resolution for visible light, in arc seconds, for any given aperture in relation to the practical limits imposed by typical seeing conditions (0.5–3 arc seconds).