

UV/Optical/IR Astronomy Notes

1. *The Atmosphere.*

The atmosphere is a major problem in the UV/Optical/IR. It has the following negative consequences:

- It blurs images, leading to spatial resolution typically an order of magnitude or more worse than the diffraction limit.
- It glows, adding a strong signal to all observations. In many cases, the Poisson noise in this signal is the dominant noise source, decreasing your sensitivity by an order of magnitude or more.
- It absorbs. At certain wavelengths it is opaque. At other wavelengths, the amount of absorption is time dependent (as clouds come and go).

1.1 **Blurring: Astronomical seeing.**

The blurring caused by the atmosphere is called “seeing”. A good description of this is on the following web page:

http://en.wikipedia.org/wiki/Astronomical_seeing

(You won’t be tested on the mathematical details)

A few points this article didn’t mention.

- The refractive index of air is less as you go to longer wavelengths. Thus the seeing tends to improve at longer wavelengths. Indeed, by the time you get to the mid-IR, you are often diffraction limited rather than seeing limited.
- For many older telescopes, it turned out that much of the seeing was being generated very close to the telescope itself – in thermals coming off the mirror, off the dome, or heat plumes from electronic racks etc. For example, the median seeing at Siding Spring is thought to be ~1.2 arcsec, but the median observed with the 2.3m is more like 1.5 – 2.0 arcsec, presumably due to these effects. Modern telescopes are designed to minimise this “dome seeing”
- Typical seeing at the best sites is ~0.7 arcsec in the visible and ~0.5” in the near-IR.

What can be done to overcome this resolution limit? The best solution is to put your telescope in space, but that is ridiculously expensive. Another solution that is very popular right now is adaptive optics. The idea is to measure the distortion caused by the atmosphere and fix it in real time by introducing a deformable mirror into the light path, and bending it so as to introduce equal and opposite distortions.

Notes on adaptive optics can be found here:

<http://www.ctio.noao.edu/~atokovin/tutorial/intro.html>

These are rather too complex for the purposes of this course – you should read them, but I won’t test you on all the more abstruse details. Here are a few of the key points you will, however, need to understand:

- Adaptive optics, as currently practiced, is far from perfect. It often gathers ~20% of the light into a very sharp core, but can smear the remainder over a very large area.
- At present, adaptive optics only works in the IR. On Gemini, for example, it does at great job at 2 microns (K-band), a reasonable job at 1.6 microns, but a marginal job at shorter wavelengths.
- Adaptive optics only corrects a tiny field of view – typically about 10'' (arcsec) across (known as the isoplanatic patch). Experiments are underway to dramatically enlarge this field of view (so called multi-conjugate adaptive optics).
- Adaptive optics requires you to have a bright “guide star” close to your scientific target. The light from this guide star is analysed to work out what correction to apply. Until recently, this had to be a real bright star - typically 12th mag or brighter. Which had to be within ~10'' of your science target. Which meant that most targets could not be observed. Recently, lasers have been used to create artificial guide stars wherever you want them. You still need a so-called “tip-tilt guide star”, but it can be fainter and further from your target, opening up much more of the sky.


1.2 Sky Glow

The sky glows at UV through IR wavelengths. This glow comes from a number of different sources:

- Emission lines. A number of emission lines, mostly molecular (O-H bonds are the worst culprits). These are really bad at some wavelengths, and absent at others. They vary rather rapidly with time. Generally they are not a problem in the UV and blue, are a bothersome in the red optical, and are a complete curse in the near-IR. Military night vision equipment (erroneously called “starlight scopes”) often amplifies and uses this light. Blocked by clouds.
- Thermal (black body radiation). The upper atmosphere glows like a 250K black body. Irrelevant in the optical, but becomes the dominant glow beyond about 2 microns wavelength. Becomes much worse if even thin cloud is present, because you detect the thermal glow from the Earth’s surface scattered off these clouds. Typically less of a problem on cold nights and at cold observatories.
- Moonlight. When the moon is up, moonlight is scattered off molecules in the atmosphere. Typically has quite a blue spectrum (the same spectrum as the blue daytime sky, in fact). Depends on the phase of the moon.
- Zodiacal light. Sunlight reflecting off dust grains in the inner solar system. Concentrated towards the ecliptic. Typically a little weaker than other sources of sky brightness from the ground, though it still can make ~0.4 mag of difference on dark nights in the V-band, if you are observing a target on the ecliptic. This is the fundamental limit for space telescopes.
- Light pollution. If your observatory is near a city, this will probably be your dominant source of noise: artificial light scattered off the atmosphere into your line of sight.


Here are some spectra of the night sky, showing these components:

Optical night sky spectrum at a dark site.



Note the very faint sky at blue wavelengths (when the moon isn't up). A strong oxygen line at 557.7nm is the first problem, but then when you get beyond ~750nm, all sorts of OH sky lines give you trouble.

Near-IR night sky



At these near-IR wavelengths, OH lines are almost everywhere. With a high resolution spectrograph you may be able to work between them, but this is not possible with imaging or low resolution spectroscopy. Beyond about 2.3 microns, you see a sharp climb due to thermal black-body radiation.

The dips seen at ~1.4 and 1.7 microns are due to absorption – the atmosphere is almost opaque at these wavelengths (as discussed in the telescope notes).

If you are doing spectroscopy, you will need to consult graphs like these to check whether a particular feature you are interested in lies on top of a sky emission line or not. For imaging, or spectroscopy where you don't know which wavelength is crucial (such as measuring redshifts), average tables of sky brightness are useful. Here are typical sky brightnesses at Siding Spring observatory, in magnitudes (Vega) per square arcsecond.

	U	B	V	R	I	J	H	K
Dark sky	22.8	22.5	21.5	20.8	19.3	15.0	13.7	12.5
6-day moon	21.2	21.3	20.8	20.4	19.2	15.0	13.7	12.5
Full moon	18.5	18.8	18.5	18.9	18.2	15.0	13.7	12.5

These numbers should be regarded with some caution. They vary from observatory to observatory, and from night to night. The exact filter you choose makes a difference, particularly in the K-band (special K_n and K' filters are more widely used than the standard K-filter, because they cut out much of the sky glow – for them the sky brightness is more like 13.5). Looking in the ecliptic makes a difference, as does even the faintest trace of cirrus cloud.

1.3 Sky Absorption

Sky absorption was discussed in the telescope notes.

2. Detectors

At wavelengths shortward of 1.1 microns, the detector of choice is typically the Charge Coupled Device (CCD). CCDs are made of silicon, and rely on incoming photons producing electron-hole pairs. The band-gap energy of silicon corresponds to a wavelength of 1.1 microns: photons to the red of this have insufficient energy to produce electron-hole pairs. So for infra-red cameras, a more exotic semiconductor with a smaller band-gap is needed. Typical choices are InSb (indium antimonide) and MgCdTe (Mercury cadmium telluride).

2.1 CCDs

See the notes on the following web page:

http://www.ing.iac.es/~smt/CCD_Primer/CCD_Primer.htm

Check out the first three powerpoint presentations on this page (Section 1, 2 and 3) – they are an excellent introduction to CCD detectors. There is a lot of information here, but CCDs are so crucial to so much of astronomy that these are important things to know about.

A common CCD problem to watch out for:

- **Fringing.** As mentioned in one of the powerpoint presentations, silicon is almost transparent to red photons. This means that they sometimes bounce backwards and forwards several times between the front and back of the CCD before being detected. If the spectrum is dominated by some narrow emission lines, as is usually the case in the red, this can lead to interference fringes: at some locations on the chip, these bounces lead to constructive interference of the light in particular sky lines, while at other wavelengths the interference is destructive.

Note – when you get data from a CCD or IR detector, you get a number of counts in each pixel. One count is not, in general, one detected photon (which produced one free electron). The counts are Analogue-Digital Units (ADU). The ratio of free electrons to ADU is called the *gain*. You can look up the gain in the manual for a given detector, or work it out by looking at the noise statistics in the data.

2.2 Infrared Detectors

Infrared detectors currently work a little differently from CCDs. Like CCDs, photons produce electron-hole pairs, and the charges pile up in the different pixels. Unlike CCDs, however, you can measure the amount of change in any given pixel non-destructively, and without reading out the whole chip. Typically, you will read the values in each pixel over and over again, watching as the charge accumulates, and beating down your read-out noise by repeat measurements.

3. Data Reduction (Imaging)

In this section I'll discuss the stages you go through to reduce optical imaging data. Most of this is described in the CCD notes referenced above, and will not make sense until you've read them.

The basic idea of data reduction is to get rid of instrumental artefacts in your data, and restore it to being a true image of the sky. If you had a perfect detector, data reduction would not be necessary. But detectors are never perfect, though some modern CCDs can come quite close.

1. **Bias subtraction.** Typically during the day before or after your observing run, you take a series of “bias frames” – zero second exposures with the shutter closed. You combine a bunch of these to make a “master bias” with reduced noise. This is then subtracted off all other data frames.
2. **Overscan Subtraction.** Often, when reading out a CCD, the read-out electronics is run a few times without clocking the CCD. Thus you are reading out, but there should be no charge coming off the CCD to measure. This results in some apparent extra rows of the CCD, which should contain pure bias. This is a check

on the bias, which can be interpolated from these extra rows and subtracted off the data.

3. **Dark subtraction.** Most detectors will slowly accumulate charge even in complete darkness with the shutter closed. This is due to thermal excitation of electron-hole pairs. For modern astronomical cooled CCDs, this is typically negligible. But for some IR detectors, and amateur CCDs (which are less cooled) it can be significant. To get rid of this, you should take some long exposures (same length as the exposures you plan to take during the night) with the shutter closed – and often with the dome in darkness to prevent light leaks. You should take several and average/median them to reduced noise. These are then subtracted off the science data.
4. **Bad Pixel Removal.** Most CCDs have some pixels that just do not work. You may have to interpolate over these, or mask them so that they do not contribute to the final reduced image.
5. **Linearity correction.** Generally, you hope that if twice as much light falls on a given pixel, you'd detect twice as much charge (i.e. the detector is linear). Unfortunately, this isn't always the case. Modern CCDs are pretty linear until you reach "saturation", when charge starts leaking into adjacent bins. As long as you keep the number of counts per bin comfortably below this limit, you should be safe. But IR detectors are often significantly non-linear. The way you correct for this is to turn on lights in the dome, pointed typically at blank white screen. Take a series of exposures with the telescope pointed at this screen, varying the exposure time. If non-linearity is a problem, you will find that the recorded signal is not proportional to the exposure time when the number of counts becomes large. You can fit a polynomial to this curve and use it to correct your science data.
6. **Flat fielding.** Not all pixels have the same sensitivity. This is partially due to irregularities in the CCD – for example the surface coatings may not be entirely uniform, or the CCD thickness may vary slightly. In modern CCDs these variations may be pretty small, but not always. You also get variations in sensitivity due to things blocking the light in front of the CCD, such as dust grains on the dewar window, and features in the telescope optics (often, for example, some light is blocked from reaching the extreme edge of the detector – this is known as vignetting). You measure this by taking lots of images of something completely blank and uniform. Typical choices are "dome flats", where you take images of a blank illuminated screen inside the dome, and "twilight flats" where you take pictures of the bright twilight sky, before it gets dark enough to see stars. You combine several of these flats to beat down the noise, then divide your science images through by them. Flat fielding accurately is often very difficult, especially for wide-field images, and will often limit the accuracy of your data. The main reason is that the flat field pattern is typically a function of wavelength – and the thing you are using to flat field (e.g. the twilight sky) has a different spectrum from your science data.
7. **Sky Subtraction.** You may still find, after all this, that the sky background in your image does not look uniform. This can be fixed by taking lots of science exposures during the night, each with the telescope pointed at a slightly different

location. You can combine these frames, using some statistic (such as removing the top pixel, or taking a median) that removes stars. This “sky” image can then be subtracted from your data. This is the best way to get rid of fringing (a mottled pattern caused by sky line radiation interfering with itself as it bounces back and forth within the CCD).

8. **Cosmic Ray Removal.** Your image probably still contains some spuriously high pixels, due to cosmic rays (actually often terrestrial background radiation) hitting the CCD during an exposure. The best way to remove these is to split your exposure into several shorter sub-exposures, and remove all pixels that are dramatically larger in one sub-exposure than in the others (e.g. by sigma-clipping or taking a median). If you don’t have multiple exposures, there are programs available which try to get rid of cosmic rays by filtering out any really sharp peaks. This works because the light from real stars and galaxies is blurred by the seeing, whereas cosmic rays are not.

In infrared astronomy, things are sometimes slightly different. One problem is that the sky is typically very bright – you may only be able to expose for a few seconds before the sky brightness alone saturates your detector. And typically the sky is much much brighter than the science objects. So you typically take lots and lots of very short images, and combine them in software to get your final image.

Another difference is with the flat fielding. At IR wavelengths, much of the signal in your supposedly flat field is actually thermal emission from the telescope, rather than coming from the supposedly flat dome screen. If you divide through by this, you will get the wrong brightnesses. What you have to do is take dome flats with the illumination of the screen on and off. The difference between these two is the true flat field – the remainder is thermal emission from the telescope.

The sky subtraction step in the analysis is then crucial – as it gets rid of this thermal glow.

4. *How bright is that star? Photometry.*


The normal approach to measuring the brightness of some object is as follows:

- Measure the number of ADU in your image of that object. Divide through by the exposure time to get the flux in ADU/sec.
- Observe a standard star of known brightness, using the same telescope/detector/filter. Measure the number of ADU/sec from this standard star.
- Use the ratio between the two fluxes to calculate the brightness of your object.

This only works in “photometric” weather conditions – i.e. no cloud. Otherwise the amount of light lost in the cloud will probably be different for the standard star and your object.


To get higher precision, you may need to allow for colour differences between the standard star and your target (as described in the telescope notes). Also, if they were observed at different elevations in the sky, the amount of atmospheric absorption might be different.

One of the biggest complications, however, is the question of how you measure the number of ADU coming from a given object. This is not as straightforward as you might think... In principle, this sounds pretty easy – you'd just go into the image and see how many counts there were in the relevant pixel. In practice, however, it's a bit complicated, as in almost all practical situations, the light will be spread over many pixels.



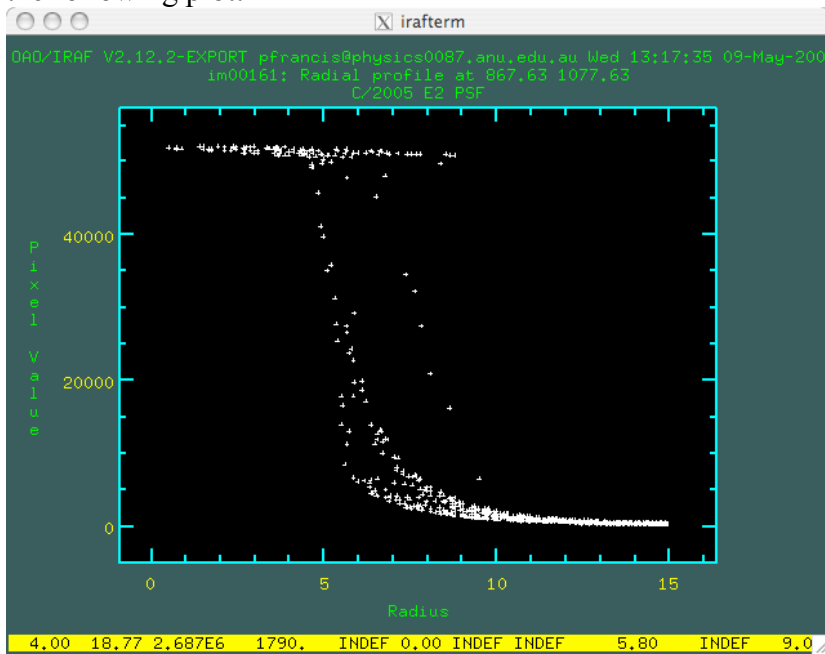
Consider, for example, the above image (taken with the 2.3m imager). Everything in this image is a star in our own galaxy. The true size of any star disk, as viewed from the Earth, is a tiny fraction of a pixel. But as you can see in the image, the detected light from each star is spread out. This is mostly due to seeing, but you can also faintly see diamond spikes sticking out from the brightest stars (A and B) – these are caused by diffraction around the struts holding up the secondary mirror.

In addition, star A has a nasty stripe extending up and down from it. This is due to charge bleeding – this star is so bright that the number of electrons generated in the central pixels of the CCD was more than the potential well could hold – some of them leaked into adjoining pixels. Here is a surface plot showing what the pixel values look like around A:



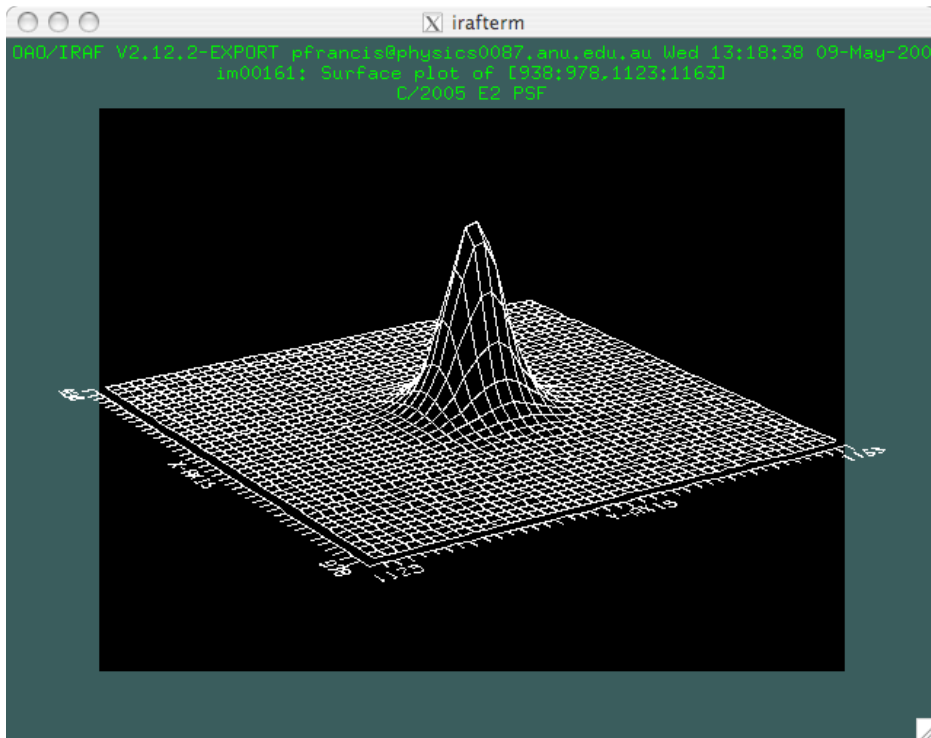
There is nothing you can do to measure the brightness of a star like this – you need to go back and observe with a shorter exposure time or with a smaller telescope.

Star B looks better, but if you plot pixel value against distance from the centre, you get the following plot:




This star, though not bleeding so dramatically, has clearly saturated the CCD (hence the flat topped profile), and is also useless for measuring brightnesses.

Star C, however, is not saturated:




This is a nice normal “point spread function” or PSF. The PSF shows how the light from something that is really a point (such as a star) is smeared by seeing and the telescope optics. Here is a radial plot:



You can see that most of the light falls within ~4 pixels of the centre, but there is still detectable light falling beyond 10 pixels out.

So how can you measure the total amount of light received from a star like this? In principle, you could just add up all the counts detected within (say) 15 pixels of the centre. There is probably a bit of light further out than this, but not a lot. This approach (a big photometric aperture) is fine for nice bright stars like this. But what about a faint star like the following one taken from the same image:



Now you can see all the noise in the sky pixels. If you add in all the flux within 15 pixels, you will be adding in lots of noise, which may well swamp your signal. But if you only include the flux from where the starlight clearly exceeds the noise (say the central 4 pixels), you'll miss a lot of light!

There are LOTS of different ways of addressing this problem. One is to measure the amount of ADU for stars in quite a small aperture – i.e. just a circle of radius 4 pixels around the centre. But you look at a bright star and work out the correction factor: how much to multiply the number of ADU in the central region to get a decent estimate of the total flux. A second technique is to use optimal weighting: you assume that the star has the same PSF as some bright star in the field, and estimate the size of it from all the pixels, weighting each pixel inversely by the variance (as in the stats notes).

What is important is that you use a self-consistent method when measuring the brightness of your target and of the standard star(s). Bear in mind that the seeing and hence the PSF may well have varied between the two observations.

This is how you measure brightnesses for stars and other point-like objects. Extended fuzzy objects, such as galaxies and comets, are much more complicated still. The methods described here are not usable, and there is no *a priori* way of knowing what the true radial profile is. Different ways of measuring the brightness of such objects can give DRASTICALLY different answers.

As usual, the best thing to do is simply to be very explicit about what approach you've used. For example, you could quote an isophotal magnitude. This means that you took all the parts of the galaxy above a certain surface brightness (a certain number of ADU/pixel) and added up their brightness. This will not be the total galaxy brightness, because there will be faint outer regions you did not include. But it is measurable and reproducible.

5. *Calculating the sensitivity of a Telescope: Imaging*

Finally, I'll show you how you use your knowledge of statistics, and of how telescopes work, to calculate the signal-to-noise ratio you might expect from a given imaging observation.

Why would you need to do this? Often there is no integration time calculator available. You could well find yourself in a job where you have to produce an exposure time calculator: they are all based on these equations. Or you might be trying to come up with a concept for a new telescope: the equations will give you an estimate of what such a telescope could do.

Step 1: your target. You need to know how bright your proposed target is. Less obviously, you also need to know how big it will appear to be. If it is an extended target (e.g. a galaxy) this will simply be the apparent area of this source (i.e. the number of square arcsec it subtends). If it is a point source, the area will be determined by the seeing you expect. If you expect (say) 1.2" seeing, that means the PSF will drop to half at a radius of 1.2". So to get most of the flux from your object, you'd need a circular aperture of radius perhaps $\sim 2''$. Which means that the subtended area is $\sim \pi r^2 = 4.52$ square arcsec.

Step 2: Photons per second. You can convert the magnitude of your proposed target into a flux, in f_λ , $\text{W m}^{-2} \text{nm}^{-1}$ above the atmosphere. Use your filter bandpass to convert this into power per unit energy (i.e. for the V-band with its bandwidth of 90nm, you'd multiply by 90). Now use $E=h\nu$ to convert the energy flux into a photon flux (number of photons per second per square metre).

Step 3: Losses. You can now estimate the number of photons per second reaching your detector. You'll typically lose $\sim 20\%$ of the light passing through the atmosphere, and $\sim 10\%$ per mirror reflection in the telescope. Some more will be lost in the instrument and filters: the manual should tell you how much. Multiply all these losses together.

Step 4: Detector Quantum Efficiency. You now know how many photons are hitting your detector. But not all of these will be detected. Look up the "Detector Quantum Efficiency" (DQE, or often just QE) at your wavelength of interest, and multiply by it to work out how many electrons will actually be produced by your object, per unit time

Step 5: Object noise. If your object gives n_o electrons per second, you will record a total of n_{ot} electrons in an exposure of length t seconds. Using Poisson statistics, your object noise $\sigma_o = \sqrt{n_{ot}}$. Often the web page of an instrument will allow you to skip steps 1-4: it will quote the number of photons received from a star of a given magnitude, and you can scale things from there.

Step 6: Sky level. Repeat all the above for the sky emission. Look up the expected sky brightness in magnitudes per square arcsec. Multiply by the number of square arcsec

subtended by your source, to get the amount of sky light that will be falling onto the same pixels in your detector. Convert into the number of electrons you will actually detect per second n_s . Your sky noise is then $\sigma_s = \sqrt{n_s t}$.

Step 7: Dark Current. If the dark current is significant (not often these days), you will need to know how many electrons per pixel it gives you, in your chosen exposure time, $n_d t$. The dark noise is then $\sigma_d = \sqrt{n_d t}$.

Step 8: Read-noise. Whenever you read out a pixel, noise is added to it. You can look up this noise $\sigma_{r \text{ pix}}$, or measure it from a bias frame. This does not depend on how long you exposure for, as it is only introduced when you read out the CCD. But you get this noise from each pixel on which your target's light falls. You can work out how many pixels this is by dividing the subtended area of the object by the area (in square arcsec) of each pixel. If the object covers n pixels, the total read-noise will be $\sigma_r = \sigma_{r \text{ pix}} \sqrt{n}$.

Step 9: Putting it all together. Now you know all the sources of noise, you can put them together. Your final signal will be $n_o t$. Your final noise will be the quadrature sum of the object noise, sky noise, dark noise and read-noise. Thus the final signal-to-noise ratio s/n will be:

$$\frac{s}{n} = \frac{n_o t}{\sqrt{n_o t + n_s t + n_d t + \sigma_r^2}}$$

Special Cases:

It is worth looking at some of the more common special cases of this equation.

- Broadband imaging of faint objects. In this case, the sky noise is usually vastly bigger than everything else, so the equation boils down to: $\frac{s}{n} = \frac{n_o}{\sqrt{n_s}} \sqrt{t}$.
- High resolution spectroscopy of bright objects: in this case the object noise dominates, and $\frac{s}{n} = \sqrt{n_o t}$
- High resolution of somewhat fainter objects: in this case the read-out noise dominates and $\frac{s}{n} = \frac{n_o t}{\sigma_r}$. This is the only case where the signal-to-noise ratio varies linearly with the exposure time.

To be continued – the second part of these notes will cover spectroscopy.