

Robert Gendler
Editor



Lessons
from the
Masters

Current Concepts in
Astronomical Image
Processing

The Patrick Moore
Practical
Astronomy
Series

The Patrick Moore Practical Astronomy Series

For further volumes:
<http://www.springer.com/series/3192>

Lessons from the Masters

Current Concepts
in Astronomical Image Processing

Robert Gendler

Editor

 Springer

Editor
Robert Gendler
Astropics
Avon, CT, USA

ISSN 1431-9756
ISBN 978-1-4614-7833-1 ISBN 978-1-4614-7834-8 (eBook)
DOI 10.1007/978-1-4614-7834-8
Springer New York Heidelberg Dordrecht London

Library of Congress Control Number: 2013941997

© Springer Science+Business Media New York 2013

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed. Exempted from this legal reservation are brief excerpts in connection with reviews or scholarly analysis or material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work. Duplication of this publication or parts thereof is permitted only under the provisions of the Copyright Law of the Publisher's location, in its current version, and permission for use must always be obtained from Springer. Permissions for use may be obtained through RightsLink at the Copyright Clearance Center. Violations are liable to prosecution under the respective Copyright Law.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

This book is dedicated to the memory of Daniel Marquardt, who during his short time with us exemplified the goodwill and camaraderie of the worldwide astroimaging community.



Preface

For more than a century and a half, astrophotography has consistently brought us vital information about our universe as well as the visual splendor of the cosmos. By way of an amazingly rich journey and through the contributions of many individuals, we have now arrived at an era where astrophotography serves as both art and science. Along the way, we have learned that art and science are not mutually exclusive objectives but together can serve the human need to learn about our place in the universe and to experience and enjoy nature on the grandest of scales.

The process of astrophotography has transitioned over time from a purely technical exercise with a defined scientific and informational purpose to an increasingly human experience where the skillful management of light and color can reconstruct powerful cosmic scenes that reach out to us and enrich our understanding and appreciation of nature on many levels. The process has become increasingly personal in a sense. A given astronomical data set can produce many different results, each one valid, and each a reflection of the eye and imagination of the individual.

Mechanically speaking, astrophotography is fundamentally a two-step process. The two processes could not be more dissimilar. The first step of data acquisition requires the precise functioning of the equipment chain. The mechanics of light collection leave little room for error. Image acquisition is dogmatic and rigid by necessity, and imposes strict rules and limits. However, the second step of image assembly and enhancement is, on the contrary, an entirely fluid process. It requires a creative and flexible mindset, and demands experimentation in order to succeed. It is truly a dynamic and creative process.

This book is more about the second step of the process. Since I began astrophotography, it is this phase of the process which I have found to be the most personally rewarding and enjoyable. People often ask me how to produce images with high visual impact. I tell them the most essential element is to view and experience as

many astronomical images as possible and emulate the ones you admire most. Only in this way can one's own photographic eye and sense of style develop and mature. Programs such as Photoshop, as wonderful as they are, can only provide the raw tools. The road to good results can only be achieved by studying the works of others and ultimately tapping into that experience to carve out your own style and direction.

The organization and objectives of this book are based on my experience as a student of astrophotography, of which I remain and will forever be. Since the advent of the CCD camera and the subsequent birth of digital astrophotography, modern astronomical imaging has become increasingly diverse and rich. In the last decade, the craft has grown exponentially, resulting in a myriad of different applications and subdisciplines. Astrophotographers today possess the technical resources to record the faintest objects from the depths of space, the finest details of the Sun, the Moon, and planets, and the sublime beauty of the local night sky here on Earth.

Certainly, technical advancements in telescope and camera technology were critical in propelling astrophotography into the modern era. Nevertheless, the final outcome of the astrophotographic process rests on the imager's ability to assemble the astronomical data into a coherent image with maximal visual and informational impact. The creative aspect of the craft is increasingly driven and defined by the successful management and enhancement of data after it is collected at the telescope. Today, the measure of success of an astronomical image is increasingly dependent on the image-processing skills and creative vision of the astrophotographer. The modern practitioner of astrophotography cannot only be proficient at taking images but must stay current with the growing array of sophisticated digital techniques used to extract the finest details, the richest colors, and the faintest signal from his or her data.

Consequently, the mastering of astronomical image processing has become the essential task for the modern astrophotographer. As techniques and methods expand and evolve at an ever-increasing pace, staying current with the latest information has become the primary mission of the dedicated imager. Yet, the delivery of that information has become challenging. There are no formal university courses, apprenticeships, or degrees in astrophotography. As a result, astronomical imagers have become a self-taught breed relying predominantly on a combination of field experience, web-based resources, and informal instruction to learn the vast nuances of this extraordinarily complex and challenging craft. There exists today a paucity of books on astronomical image processing primarily because it is impossible for one author to cover the full range of subject areas with the expert precision each area deserves.

This book – by its very nature, a collection of works by individual contributors, each distinguished in their particular areas – will attempt to accomplish this task by covering in systematic detail each of the major subdisciplines of astrophotography. This approach offers the reader the greatest opportunity to learn the most current information and the latest techniques directly from the world's foremost innovators in the field today. Each chapter covers the specific processing techniques, methods, and strategies unique to each of the major subdisciplines of astrophotography. A large portion of the book is devoted to “deep sky imaging” since it represents the


foundation of astrophotography. Individual chapters cover specific challenges within the realm of “deep sky imaging,” such as bringing out faint structure, enhancing small- and large-scale detail, noise reduction, and narrowband imaging. To add some practical elements to the instruction and ensure that all pertinent subject matter is covered, I asked three world-renowned imagers to write chapters in which they describe their preferred “deep sky” workflows in step-by-step detail.

The information in the pages ahead may not answer every processing question that arises, but it will provide the reader with the necessary skills and strategies for successful high-level imaging. The collective input from multiple contributors, all with different perspectives and experience, is extraordinarily valuable but also increases the chance that there may be some overlap in the information delivered. In this sense, overlap is a good thing. Some redundancy reassures and confirms that certain methods are more time proven than others. The field of astronomical image processing is not a static one, but a markedly fluid and ever-evolving craft. The challenge of trying to cover the most current material in a rapidly changing field is similar to chasing a moving target. Nevertheless, the pages ahead should offer the reader a strong foundation of skills and knowledge to build on for years to come.

For many, astrophotography is an adventure of the mind, heart, and spirit. It is a difficult and demanding journey, but one that pays back immensely and offers an unparalleled sense of discovery. The objective of this book is to provide guidance to those individuals who want to take this amazing journey themselves. They want to learn the powerful tools and techniques of image processing so they, too, can experience the joy and wonder of astrophotography. It is my hope that the pages ahead will provide motivated imagers with the necessary tools, techniques, and strategies to help them find their own path and direction in astronomical imaging.

Avon, CT, USA

Robert Gendler



The Theory of Astronomical Imaging..... 1
Stan Moore

High Dynamic Range Processing..... 29
Ken Crawford

Intensifying Color 51
R. Jay GaBany

Revealing Small-Scale Details..... 73
Ken Crawford

Bringing Out Faint Large-Scale Structure 97
Rogelio Bernal Andreo

Narrowband Imaging 115
Don S. Goldman

**Widefield Imaging: Selected Strategies for Processing
Light-Contaminated Data** 131
Stephen A. Cannistra

Noise Reduction Techniques 149
Tony Hallas

Deep Sky Imaging: Workflow 1 159
Adam Block

Deep Sky Imaging: Workflow 2	193
Johannes Schedler	
Deep Sky Imaging: Workflow 3	217
Steve Mazlin	
High Resolution Lunar and Planetary Imaging	233
Damian Peach	
Secrets to Successful Earth and Sky Photography	261
Babak A. Tafreshi	
Imaging and Processing Images of the Solar Corona	293
Fred Espenak	
Catching Sunlight	319
Alan Friedman	
Aesthetics and Composition in Deep Sky Imaging	343
Robert Gendler	
Hybrid Images: A Strategy for Optimizing Impact in Astronomical Images	363
Robert Gendler	
Author Biographies	377
Index	383



The Theory of Astronomical Imaging

Stan Moore

This chapter presents a conceptual framework for understanding the basic physics and dynamics of astronomical imaging. The core concept is that an image captures information about astronomical objects (stars, galaxies, and such) and the quantity and quality of that information places strict limits on the image.

An understanding of these concepts can inform many decisions and expectations regarding astronomical equipment and imaging techniques. Although this information theory approach can be understood without solving mathematical equations, such an understanding is incomplete. The basic equations presented below are invaluable for understanding the sometimes unexpected behavior of signal-to-noise ratios and such equations are necessary for the analytics and predictions that can be put to practical uses.

The Photonic Nature of Light and Noise

Astronomical images present information from distant objects. That information is conveyed by light photons that are electromagnetic waves with a particle (quantum) nature. Photons are emitted and detected on an individual quantum basis, and this has profound effects on information carried by the light. The quantum photonic nature of light limits the quality of information about an imaged object due to inherent uncertainty associated with quantum information. This uncertainty is called Poisson noise.

Poisson noise can be demonstrated by measuring rainfall over an area by collecting rainwater in many buckets within a limited area. Water in each bucket is measured at the end of the storm to reveal that each bucket collected slightly different amounts of water due to the randomness of the falling drops. Averaging all buckets will produce a more accurate estimate of the rainfall than any single bucket.

Poisson noise is caused by random variations of quantum accumulations and exhibits a mathematical predictability. The uncertainty of an accumulation of random discrete events is described by:

$$\text{Poisson noise} = \sqrt{\text{number_of_events}}$$

Noise denotes the range of values for which there is 68 % likelihood that the true value is contained within that range, akin to the “margin of error” stated in opinion polls.

The inherent Poisson noise of light signals means that it is not possible to increase information by indefinitely amplifying or distilling a signal because the noise/uncertainty is intrinsic to the signal itself. So there are strict physical limits to what any camera can do, and it is useful to understand those limits and how they impact imaging decisions. It is interesting to note that this dynamic does not exist for pure wave energy, and imaging would have a very different dynamic if it were not for the quantum nature of light. Non-quantum energy does not embody the inherent uncertainty of Poisson noise. Noise in a non-quantum system is caused by extraneous sources that interfere with the pure signal so the noise level is unaffected by signal strength. A stronger non-quantum signal quickly overpowers noise, and S/N becomes nearly linear. In such a system, doubling a strong signal would also double S/N, whereas in a quantum system, it is necessary to quadruple the signal to double the S/N. CCD images offer proof that light has a profound quantum nature.

This image demonstrates limitations imposed by the quantum nature of light. This is a photon counting image of M57, the Ring Nebula. Each dot is a photon. For the short exposure time (30 ms) either a photon arrived at a location or not. The image was taken with a photon counting ICCD having zero read noise (ZeroCam) (Fig. 1).

The Virtual Image

The virtual image consists of information conveyed to the focal plane by photons from astronomical objects, regardless of camera (“before” the camera). Characteristics of the virtual image set limits on information that can be obtained by a camera. So it is important to understand the physics and dynamics of the virtual image, then examine how that interacts with the camera.

Light from astronomical objects is collected by a telescope aperture for a certain length of time. The collected photons are conveyed by the optics to the focal plane, where the virtual image is formed. Information about each object is contained in the virtual image. At root, that information consists of the location of every photon collected.




Fig. 1 M57 (Ring Nebula) 100 ms exposure from a photon counting camera

An object's information is based on the number of photons collected from the object, which is proportional to the area of the telescope's aperture and collection time. Focal length has no effect on the number of photons collected from an object. Thus focal length and f-ratio have no impact on the virtual image other than determining linear scale, which is irrelevant for angular properties.

Signal to Noise (S/N) Ratio

Information about the object is contained in the object's photon count, or "signal." Noise is the uncertainty of that information. The proportion of signal-to-noise is the "Signal to Noise Ratio," S/N or SNR. S/N is a widely used measure of information quality. S/N determines ability to discern contrasts in an imaged object.

An "object" can be anything, even an arbitrary patch of sky. But most meaningfully, an object is a star, galaxy, or nebula that occupies an angular area in the sky. A star object is the simplest. All stars have the same angular extent for a given system/exposure. Stellar S/N is a straightforward measure with many applications; for example $S/N=3$ is the lowest S/N for star astrometric detection. Star S/N determines astrometric and photometric accuracy. The simple object S/N of large objects is affected by object size, so for large objects it may be preferable to use an angular measure such as square arcseconds. An angular measure is similarly used for galaxies – magnitude per arcsecond squared. For a large object, higher S/N reveals

finer contrasts, so a fuller description of S/N would include MTF terms, but that is beyond the scope of this chapter.

It is worth noting that this object information theory approach to imaging is not commonly employed (at least not explicitly), and that the usual definition of S/N is restricted to single pixel dynamics, oblivious to object/informational S/N. Pixel-based S/N is very easy to measure (at least for undifferentiated, non-dimensional illumination such as sky glow) and that ease and seeming simplicity can be seductive. This chapter primarily concerns the fidelity of representing an actual astronomical object and is less concerned with internal camera electronics or non-specific visual densities at particular image scales, as is a primary concern of “pixel S/N.” Simple pixel S/N is a superficial concept that can too easily be misapplied, failing to account for the larger context of object fidelity, as illustrated by “the f-ratio myth” below. Pixel S/N can “fail to see the forest for the trees.” However, most of this chapter and equations can easily be simplified to “pixel S/N” as a sub-domain of object S/N.

Object S/N of Pure Signal

Considering only the photons collected from the object and the associated Poisson noise:

$$\frac{S}{N} = \frac{nPhotons}{\sqrt{nPhotons}} = \sqrt{nPhotons}$$

The non-linear nature of simple S/N occurs repeatedly in astronomical imaging situations and issues (Fig. 2).

This simple equation contains advice for imaging. To double the S/N it is necessary to quadruple the signal, which can be accomplished by exposing 4× longer or using an aperture with twice the diameter. That relationship embodies a “hitting the wall” dynamic whereby reasonably long exposures cannot be significantly improved without expending inordinate amounts of time. For example, it is necessary to shoot an additional 12 h to double the S/N of a 4 h exp.

Object S/N with Sky

In addition to signals from astronomical objects the virtual image also accumulates photons from the atmospheric sky (high-altitude emissions and light pollution). Sky photons create a diffuse signal over the entire field. Because it is evenly diffuse the overall intensity of sky light can be subtracted but the sky’s Poisson noise cannot be removed. This affects object S/N.

$$\frac{S}{N} = S / \sqrt{S + S_{ky}}$$




Fig. 2 Signal to Noise as a function of signal in the absence of other noises

where

S=Number of photons from object.

Sky=Number of photons from sky glow that occupies the same area as the object.

This equation contains further advice for imaging. Objects that are significantly brighter than the sky produce S/N only slightly less than would be the case for darker skies. But S/N for objects significantly dimmer than the sky is strongly constrained by sky brightness. For example, if the sky is 3× brighter than the object then S/N is about half as much as the same object in a much darker sky and so it would take 4× more exp time to compensate for the sky noise. Here is a graph of the effect on object S/N due to moderate sky glow (10 photons) at various signal intensities (1 =same S/N as if no sky glow) (Fig. 3).

Point-Spread Function (PSF) and FWHM

The virtual image information basically consists of photon arrivals. Every photon arrives at a different location (even if the distances are subatomic) and this location information constitutes the image resolution. At less than a particular distance,




Fig. 3 Differential effect of sky glow (10 photons) on S/N by signal intensities (1 = same S/N as if no sky glow)

location precision has little or no effect on information quality because the resolution of the image is constrained by physical limits of diffraction, optical quality, atmospheric stability (seeing), tracking, and so on. This limit can be characterized as a point-spread function (PSF).

A point-spread function (PSF) is the result of blurring a point-source. Stars have incredibly small angular diameters that are essentially point-sources. The atmosphere, 'scope, and mount diffuse the starlight into a fuzz-ball that when graphed shows a bell-shaped curve (PSF). Although telescopes produce a complicated diffraction ring PSF the convolution from seeing over time intervals blur the PSF to closely approximate the Gaussian function.

The Gaussian function contains a parameter that defines the width of the PSF. This function width is a reliable and constant measure of resolving power regardless of intensity. Function width is a critical factor in the Nyquist sampling theorem discussed below. The commonly used FWHM (Full Width at Half Max) characterizes the PSF resolution by measuring the width (diameter) of the star at half of peak intensity. FWHM is directly related to the Gaussian function width by: $\text{FWHM} = 2.355 * \text{width}$.




Fig. 4 Full Width at Half Maximum (FWHM) is constant regardless of star brightness

Star images will exhibit different “sizes” – brighter stars seem to be larger than dimmer stars. That is due to the bell curve of PSF, where the apparent diameter of a star is determined by particular intensity level. FWHM is a constant measure of PSF and resolution regardless of intensity. Both dim and bright stars have the same width at their half maximum intensity (of course that maximum intensity is different for different stars) (Fig. 4).

Sampling

To capture the virtual image as a digital representation it is necessary to sample the image information. Imagine overlaying the virtual image with a grid (graph paper) that combines/collapses photon locations into a limited number of grid squares. Each square embodies two types of information: the location/area of the square and the number of photons encompassed within. Each grid square is a pixel.

Nyquist Sampling Theorem

The virtual image itself is unaffected by different grid overlays or pixels sizes, but the sampled image is degraded if the sampling scale (pixel size) is insufficient to

preserve resolution. The Nyquist theorem postulates that to preserve information the sample size (or rate) must be no larger than the function width of the signal. The STD of Gaussian PSF is the function's width, so the Nyquist sample size for a Gaussian function measured as $\text{FWHM} = 2.355$ pixels (see PSF above). But because pixels are square with a long diagonal and because astro-images are routinely resampled (e.g., alignment registration), critical sample size should be adjusted upwards to $\text{FWHM} = 3.0\text{--}3.4$ pixels.

Images with $\text{FWHM} < 2.3$ pix are "under-sampled". Under-sampled images involve information from the virtual image. For example, close double stars may merge and small details may wane in an under-sampled image. The degradation may be acceptable for modest under-sampling ($\text{FWHM} = 2$ pix) but can be significant for severe under-sampling ($\text{FWHM} < 1.5$ pix). Regardless, some images are deliberately under-sampled out of necessity to display a wide field on a small screen. Also there is an aesthetic that values the "pinpoint" stars of under-sampled images.

Images with $\text{FWHM} > 3.5$ pix are "over-sampled." Over-sampled images contain all of the useful resolution information of the virtual image, more than is strictly necessary. A modest amount of over-sampling can be useful for stacking and other processes such as deconvolution. But significant over-sampling ($\text{FWHM} > 5$ pix) can have undesirable consequences for real-world cameras (discussed below).

The Nyquist criterion can be used to quantify the spatial coverage of an image by dividing image size by the PSF width ($3 \times \text{FWHM}$). The result is the number of pixels necessary to capture and display the information. For example, if the corrected field of a scope is 30 arcmin and actual resolution (including seeing) is $\text{FWHM} = 3$ arc-sec then the critically-sampled image size would be $1,800 \times 1,800$ pixels.

Image Scale and Pixel Size

Distances in the virtual image can be measured three ways: angular, linear, and pixel. For example the angular diameter of Jupiter might be:

- Angular: 40 arc-seconds
- Linear: 200 microns for focal length = 1 m
- Pixel: 40 pixels for focal length = 1 m and pixel = 5 microns

Angular distances are constant for all focal lengths and pixel sizes. Linear distance depends on focal length. Pixel measurements depend on both focal length and pixel size. For purposes of analysis and comparison it is often most useful to use angular measurements.

Object S/N Versus Pixel S/N

The virtual image of an astronomical object can be characterized by "object S/N," which is unaffected by sampling (unless severely under-sampled). In the sampled

image, the object's signal still consists of all photons from the object, regardless of which grid squares (pixels) they occupy. It is a simple matter to sum the relevant pixels' signals to derive the object's signal.

A single pixel in the sampled image may contain some object signal (photons from the object) but not the entire signal, which is distributed over several pixels. It is easy to measure the signal of a pixel and calculate pixel noise and S/N statistics. This "Pixel S/N" is a simplification that can be useful for understanding S/N dynamics, but it must be understood in the context of object information or it can result in misunderstandings and incorrect applications.

Pixels Are the Trees that Can Distract from Seeing the Forest!

Pixel S/N is most misleading when applied to different angular samplings. Consider a virtual image of a square object 2×2 arcsec and total intensity of 100 photons with no sky glow. Sample that image with 0.5" and 1" pixels:

- Object S/N = $\sqrt{100} = 10$
- 1" pixel: object occupies 4 pixels with mean signal = 25; pixel S/N = 5
- 0.5" pixel: object occupies 16 pixels with mean signal = 6.25; pixel S/N = 2.5

The object's image is not somehow superior with the larger pixel because the pixel S/N is twice the S/N of the smaller pixel. It is the very same virtual image and only the number of pixels used to represent it has changed (i.e. the overlay grid has a different mesh size). This misunderstanding is the basis of the "f-ratio myth," and one way to unwind that myth is to use a "level playing field" by normalizing pixel sizes.

The F-Ratio Myth

Experienced photographers know that faster f-ratios mean shorter exposure times. Advertisements for telescope focal reducers often claim the faster f-ratio will take the same image in much less time. Similarly, it is often claimed that CCD binning will double S/N or shorten the exposure by $4\times$.

From an object information perspective, f-ratio per say is irrelevant (other than possible secondary effects of camera noise, discussed below). An object's photons are collected by aperture and time and that virtual image information is unaffected by focal length or sample size. It is not possible to magically create more information by simply changing sample size. Varying angular pixel size changes the pixel S/N but does not intrinsically affect an object's S/N (Fig. 5).

However, like most myths, there is some truth in the f-ratio myth. Certainly it is true when varying f-ratio by changing aperture while keeping focal length constant, as is the case for normal photographic exposure control. And the myth's f-ratio




Fig. 5 Testing the “f-ratio myth”: same aperture, same exposure times but very different FL and f-ratios

dynamic applies to non-normalized pixel S/N analyses, which can provide assessments of superficial visual quality at particular image scales. But perhaps most interestingly, the f-ratio myth’s predictions can be approximately correct for short exposures, though for reasons entirely different than the myth’s supposition, and this is due to imperfections of real cameras (read noise in particular).

Normalizing Pixel Size

To compare pixel based characteristics of different-sized pixels it is necessary to mathematically normalize the pixel sizes so that the angular sizes are equal (equivalent to defining an “object” with constant flux over a fixed extent). This is especially important when evaluating different cameras for use on a particular telescope. Camera specifications are per pixel. At first glance it might seem that a detector with 10 μm pixels having full-well depth of 40,000e⁻ is deeper than a 5 μm detector

with full-well depth of 10,000e- until you consider that the depth per actual area is identical. A star that occupies 4 pixels in the first camera will occupy 16 pixels in the second and so each smaller pixel need not be as deep as the larger pixel.

The same dynamic applies to dark current; e.g., a 10 μm detector with dark current of 0.1e-/pix/s has the same effect as a 5 μm detector with dark current = 0.025e-/pix/s. Noise must be normalized quadratically (sqrt of the sum of squares), so it is a bit more complicated to normalize read noise; e.g., a 10 μm detector with read noise = 10e- produces half as much noise as does a 5 μm detector with read noise = 10e-.

Pixel sizes can be normalized from one detector to another or to a standard reference. Pixel sizes can be angularly normalized, such as 1 arcsec (to compare different ‘scopes’); or linearly normalized, such as 1 mm (for detectors used on the same ‘scope’). For example, to normalize detector characteristics to 1 mm:

$$\text{Well capacity per square micron} = \frac{\text{full_well}}{ps^2}$$

$$\text{dark current per square micron} = \frac{\text{pixel_dark}}{ps^2}$$

$$\text{Read noise per square micron} = \sqrt{\frac{rn^2}{ps^2}}$$

Where ps = pixel size in microns.

Real Cameras

A real camera is unable to perfectly capture a virtual image. Some information from the virtual image is lost (QE) and the uncertainty associated with the information is increased (camera noise). Both of those effects decrease object S/N.

Photon/Electron Accumulation, ADU, DN, and Gain

Imaging photon detectors such as CCDs and CMOS’s accumulate photon information by converting incoming photons to stored electrons. For normal detectors each stored electron represents one photon (intensified detectors such EMCCD and ICCD employ many electrons per photon). At the end of an exposure the stored charges are read out to produce numeric digital values for each pixel. The pixel values are called ADU (Analog to Digital Units or Arbitrary Digital Units) or DN (Digital Number). The ADU values are related to electron/photon counts by Gain:

$$\text{Photons detected} = \text{Electrons measured} = \text{Gain} * \text{ADU}$$

Gain is usually documented for each camera and can be measured via various methods, most of which employ evenly illuminated frames to link measured noise (STD) with photonic Poisson noise. S/N and noise equations are quadratic (squares and square-roots) and require measurement in real quanta, so it is necessary to convert ADU to electrons:

$$electrons = gain * ADU$$

Quantum Efficiency (QE)

Quantum efficiency (QE) is a measure of the percentage of incoming photons that are detected by the imaging sensor. For example, if QE=0.66 (or 66 %) then the sensor fails to detect 1 out of every 3 photons. QE is determined by material characteristics of the sensor and is also affected by the obstruction of overlaying circuit elements. Micro-lenses on each pixel are commonly used to route light around the overlaying circuit elements. “Back illuminated” sensors are fabricated to allow light to enter the sensor’s back side, avoiding all circuit elements.

Camera Noises

As shown above, sampling (pixel size) has no intrinsic effect on object information (S/N) in the virtual image. But real cameras contribute noise at the pixel level and so the effects of that noise on objects’ S/N must vary with pixel size. It requires more small pixels than large to encompass the same object, which imparts more camera noise into the object’s S/N. There are two primary sources of pixel-based camera noise: read-out and dark current.

Read Noise

Read noise (or read-out noise) is uncertainty resulting from physically measuring each pixel’s signal. CCD binning combines signals from multiple pixels prior to read-out; thus CCD read noise is associated with logical (binned) pixels, not physical pixels. CMOS detectors cannot be binned because read-out is direct access, so CMOS noise is always associated with physical pixels.

Read noise is commonly specified for astronomical CCDs and can be measured empirically. A simple measure of read noise can be done using two “bias” exposures (short-exp dark frames): subtract one bias frame from the other then measure the Standard Deviation (STD) of the resulting frame; convert STD to electrons (ADU * gain=electrons) and divide by sqrt(2).

Dark Current

Dark current consists of electrons generated by the molecular and atomic effects of heat. Dark current continually accumulates during the exposure and is normalized by dark subtraction calibration. Dark accumulation is subject to the uncertainty of Poisson noise that cannot be removed via calibration. So it can be important to minimize dark current by cooling the detector.

Object S/N with Sky and Real Camera

This equation captures the primary significant terms and dynamics for most imaging situations and contains many useful insights:

$$\frac{S}{N} = S * QE / \sqrt{(S + Sky) * QE + f * p * rn^2}$$

where:

S=Number of photons from object

Sky=Number of photons from sky-glow that occupies the same area as the object

QE=Quantum Efficiency (electrons/photon)

f=number of frames used in stack (f=1 for single frame)

p=number of pixels covered by object (p=1 for pixel S/N)

rn=read noise (electrons)

These noise and S/N equations are “quadratic.” Terms are squared before being added, then the square-root of that sum produces the result. The squaring and square-root operations produce a dynamic that may be non-intuitive or even counter-intuitive, and many aspects of imaging may not follow common sense. Quadratic summation is non-linear, so increasing or decreasing the signal or noise usually does not produce proportional changes in S/N; for example, it takes four times more photons to double S/N.

Due to the squaring, high valued terms can potentially render low-value terms insignificant, meaning the resulting S/N is nearly the same with or without the low-value terms. This dynamic can be exploited by the imager; for example, the above equation reveals how exposure times and sampling can be set to produce “sky limited” exposures.

Example Application of S/N Equation: Sky Limited Exposures

An important dynamic of the S/N quadratic equation is the concept of “sky limited,” whereby it is possible to relegate read-noise to near insignificance by manipulating exposure times and/or sampling (pixel size). Total noise is often dominated by the highest noise, whereby other noises are not great enough to significantly

affect the quadratic sum. For bright objects, the Poisson noise of the signal itself is the limiting factor. But for dim signals the S/N may be dominated by sky and/or camera noises, which limits the S/N. An exposure is “sky limited” when the background noise is dominated by sky rather than camera noises.

Using “pixel S/N” ($p=1$) to simplify the basic concept:

$$\text{Background noise per pixel} = \sqrt{\text{skyRate} * t * f + f * rn^2}$$

where:

skyRate = sky electrons per second per pixel

t = number of seconds per frame

Total exp time = $f * t$

Thus fewer frames require longer exposures per frame to maintain the same total exp time. Conversely, shorter exposures per frame require more frames.

$$\text{Background noise per pixel} = \sqrt{\text{skyRate} * \text{totalTime} + f * rn^2}$$

Stacking more frames per total exposure time increases the effect of read-noise while the sky term remains constant. Thus it would seem necessary to limit the number of frames by taking the longest possible exposure per frame. However, the quadratic equation contains a non-linear loophole: the exposure per frame only need be long enough for the sky noise to limit the total noise, regardless of number of frames. To see how that dynamic might work in a single frame ($f=1$), it is necessary to plug in a few typical values:

For:

$$\text{skyRate} = 10e - /s$$

$$rn = 12 e -$$

If totalTime = 1 s then background noise = $\sqrt{10 + 144} = 12.4$. The background noise is limited by readout; if sky was zero then $N=12$, which is hardly any different. But if $t=100$ s then background noise = $\sqrt{1,000 + 144} = 33.3$, which is limited by sky; i.e. if read-out was zero then $N=31.6$.

Effect of Sampling on Sky Limit

Object S/N provides further insight. The background contains no object but is subject to angular sky flux:

$$\text{Background noise per area} = \sqrt{\text{skyAreaRate} * t * f + f * p * rn^2}$$

where:

skyRate = sky electrons per second over the area

t = number of seconds per frame

Examine the dynamic of a single frame ($f=1$):

$$\text{Background noise per area} = \sqrt{\text{skyArea} + p * rn^2}$$

The number of pixels per area:

$$p = \frac{\text{area}}{\text{pixelSize}^2}$$

Smaller pixels result in more pixels per area and thus increase the impact of read-noise, while leaving angular/object signal and sky noise unaffected. This dynamic increases sky limited exposure times for smaller pixels.

Practical Considerations for Sky Limited Imaging

There are several sky limited exposure calculators on the web, but it is easy enough to estimate sky limited exposure times from the above equations.

$$\text{Background noise per pixel (including read noise)} = \sqrt{\text{skyRate} * t + rn^2}$$

$$\text{Background noise per pixel (excluding read noise)} = \sqrt{\text{skyRate} * t}$$

Use a threshold factor h to equalize the two (e.g., $h=1.10$ for 10 % tolerance):

$$h * \sqrt{\text{skyRate} * t} = \sqrt{\text{skyRate} * t + rn^2}$$

Solve for sky limited exposure time t :

$$t = \frac{rn^2}{\text{skyRate} * (h^2 - 1)}$$

Approximately 10 % tolerance:

$$t = \frac{rn^2}{\text{skyRate} * 0.2}$$

Measure and Calculate Sky Limited Exposure Time

Use a dark subtracted image to measure the mean or median background ADU to calculate skyRate:

$$\text{skyRate} = (\text{background} - \text{pedestal}) * \text{gain} / \text{exposure_time}$$

$$\text{pedestal} = \text{ADU added in dark subtraction}$$

$$(\text{determined by software, usually } 100 \text{ or } 0)$$

Example

Calculate a sky limited exposure with 10 % tolerance using a camera with [gain=2, pedestal=0, read noise = 10e-]

Reference a 100 s image with [100 ADU average background]:

$$\text{skyrate} = 100 * 2 / 100 = 2 \text{ e- / sec / pix}$$

$$t = 100 / (2 * 0.2) = 250 \text{ sec}$$

Impossible Limits

There are situations where it may be impossible to take sky limited exposures. The sky flux per pixel (skyRate) may be so low that the above equations produce untenable exposure times or long exposure times that result in excessive saturation of brighter objects. Dark skies, slow f-ratio, filters and smaller pixel size all decrease skyRate. Increasing pixel size increases skyRate, which can become a trade-off between S/N and resolution.

The sky limited exp times necessary for filters may be prohibitive, especially for narrowband filters. Even broad-band color filters can result in long exposure times. The LRGB technique of color imaging employs binned color data, which makes it easier to take sky-limited filtered exposures due to the increased skyRate of the larger pixels. Another way to understand that dynamic is that by reducing the number of pixels in the object(s) so the total read-noise is also reduced.

Image Construction and Processing

The principles of object information can also be applied to image construction and processing. Basic astronomical image construction consists of frame calibration to remove most of the pseudo noise due to predictable variations from the camera. Most deep space images are constructed from multiple frames and the principles of S/N can optimize “stack” operations. Additionally, S/N statistical processing can identify and modify bad data within or outside of the stacking operation. An image processing such as blurring or sharpening modifies object information in predictable ways.

Image Construction: Frame Calibration

The individual pixels of digital cameras vary from each other in two important ways: dark current and quantum efficiency (QE). These variations can produce undesirable artifacts such as a “salt and pepper” look, bright/hot or dark/cold pixels, anomalous columns, subtle grids, rectangles, and so on. Additionally the ‘scope and optical elements produce varying obstructions such as vignette and dust spots. These effects may obscure or impair images. But these effects can be known, predicted and removed via calibration.

Calibration: Dark Subtraction

During an exposure, thermal molecular motions knock electrons into pixel wells and these electrons are indistinguishable from photon-generated electrons. Thermally generated electrons are called “dark current.” Due to electronic structures and material/fabrication imperfections, different pixels accumulate dark electrons at different rates. Fortunately, each pixel’s dark current is constant for stable temperature. This makes it possible to create a calibration “dark frame” to characterize the dark current. Each image pixel is calibrated by subtracting the corresponding dark frame pixel, thus removing the dark current.

According to information theory it is not possible to exactly remove the dark current from each pixel because that information is quantum (electrons) and thus subject to the uncertainty of Poisson noise. So each pixel contains dark noise from the image and dark subtraction. The total noise from dark current after dark subtraction is:

$$dark\ noise\ per\ pixel = \sqrt{darkCount + darkSubtractorNoise^2}$$

where:

$$darkCount = total\ dark\ current\ electrons/pixel$$

$$darkSubtractorNoise = dark\ frame\ noise$$

darkSubtractorNoise can be decreased by combining many dark exposures into a “Master Dark Frame.” For the most part, dark exposures contain noise from dark current and read-out.

$$\begin{aligned} darkSubtractorNoise\ per\ pixel &= \frac{\sqrt{f * (darkCount + rn^2)}}{f} \\ &= \frac{1}{\sqrt{f}} * \sqrt{darkCount + rn^2} \end{aligned}$$