High Dynamic Range Astronomical Imaging

Thor Olson; Nightscapes; Minneapolis MN

Abstract

High dynamic range imaging methods that combine multiple exposures involve a number of processing details, such as frame alignment, motion artifact removal, and pixel data weighting. When applied to astronomical scenes, these steps require additional consideration. Specifically, motion blur artifacts must be corrected by techniques that are tolerant of saturated pixels, successive image frames need to be registered with point source precision, and the weighting of shorter exposures in the HDR compositing step are limited by signal-to-noise considerations. In addition, if there is light contributed by the foreground sky (common in light polluted regions), it needs to be offset and balanced in a blackpoint normalizing procedure. This paper identifies and describes these HDR issues specific to creating astrophotographic imagery.

Introduction

High dynamic range (HDR) imaging has become a popular tool among photographers as a result of pioneering work in computational photography [2, 9], and the introduction of software tools and utilities to assemble and manage HDR images. It is particularly effective in outdoor pictures where the scene radiances cover a range that exceeds what can be captured in any single film or digital exposure.

The basic technique is to take a set of exposures, bracketing the range so that everything in the scene is properly exposed on at least one frame. The frames are then computationally blended to obtain a high dynamic range image. A tone mapping operation is performed to bring out the desired details for a given display or print. Sophisticated tone mappings take into account the spatial variations of intensity across the image, mimicking the adaptation of the human eye to a high dynamic range scene.

It would seem that HDR methods would be a natural fit for astrophotography, where the objects in a scene cover an enormous range of brightness, from the sun, moon and planets to faint nebulas and galaxies. There are a number of difficulties however, some of which are intrinsic to the HDR process but are exacerbated by astronomical subject matter (frame alignment, exposure weighting, motion artifacts), and others that are unique to it (sensor noise, sky fog).

1. Alignment and registration

Registration is a prerequisite for blending the individual exposures to form an HDR image. The frames are aligned so that the pixels are correlated to the same surface points of the same objects in the scene. Accurate registration yields good sharpness and preserves texture of course, but the objects that dominate astrophotos are stars; point sources that, if not tightly aligned from frame to frame, will cause the HDR merging process to fail.

Most terrestrial scenes can be adequately aligned by affine transformations of scaling and rotation, but image formation by lenses is not perfectly rectilinear, and this is only a first step toward the registration required here. Astrophotographic targets may not change their relative positions in the sky, but pictures of them can be acquired on different nights or even different years, making it unlikely that their positions within the frame will be identical.

This is understood in the astronomical imaging community, where it is common to add images together to increase the signal to noise ratio, and the software tools to help with this summation usually includes provisions for alignment. Not all perform the local warping necessary for the tight registration we need, but one such tool that does is a program called RegiStar [8]. It identifies the stars in a scene and performs the operations necessary to register additional frames to match the individual stars in the reference frame.

2. Motion deblur

Astrophotos are taken at the parametric edge of conventional photography. The exposures can be very long, seconds, minutes, even hours, and the optics can be operating at high magnification; focal lengths measured in meters is not uncommon. This makes the exposures vulnerable to motion errors, from vibrations induced by the shutter, (or by wind on the equipment), and from tracking errors in keeping perfect aim at the targets as the earth moves slowly beneath them.

There have been a variety of methods developed to remove motion blur from photographs [1, 3, 4, 5, 10]. One of the popular techniques in astronomy is Lucy-Richardson deconvolution, an iterative method to estimate what a source image must have been in order to produce the distorted image observed. Since each image of a star is a representation of the point spread function of the system (including the motion), it can be used in the deconvolution algorithm.

However, convolution and deconvolution are linear operators that assume the observed pixel values are the sum of the motion-disturbed signals. This is true for all of the non-saturated pixels in a calibrated radiance or luminance image, but not for those that have been clipped. The result is that artifacts occur in and near these regions. In most applications, the saturated regions are not of interest, but for our purpose, we care deeply about them. The saturated pixel locations matter because we will be identifying them to be rejected from one frame, and supplied by another. They must not artificially become

Copyright 2007 Society for Imaging Science and Technology

suddenly unsaturated, nor should they cause other nearby pixels to become incorrectly saturated.

In the deconvolution process, the saturated pixels represent limited information about the intensities in the scene. We must not ignore them, but rather use them to provide bounds on the amplitudes of the other pixels in the neighborhood. To this end, a saturation-preserving version of the Lucy-Richardson deconvolution algorithm was developed. It permitted use of frames that would otherwise have been highly misrepresentative in the HDR stack. Some examples of its behavior are shown.

Of course the best answer is to avoid the motion artifacts to begin with. Camera mirror lockup, solid mechanical mounting, guided tracking, are methods that can reduce the blur. Although we have discussed it second, it should be noted that numerical deblurring methods if used, should be performed *before* frame registration.

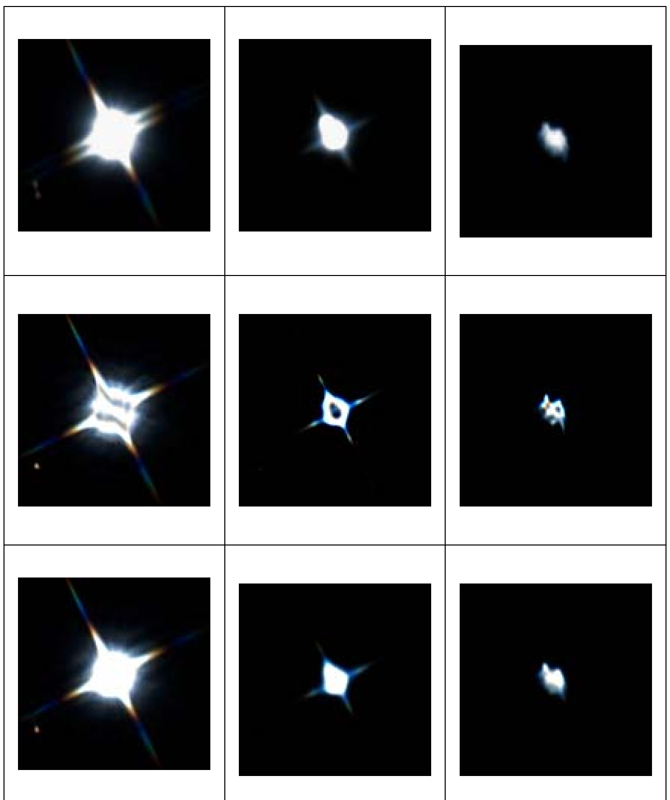

30 sec	4 sec	1/30 sec, (2x spatial scale)

Figure 1. Some examples showing the results of motion deblur algorithms. The top row are selected samples of motion blurred exposures in an HDR sequence. The middle row shows the application of Lucy-Richardson deconvolution supplied in the Matlab Imaging Toolbox. The bottom row is output from DeblurLucySat, a Matlab script developed to account for pixel saturation.

3. Noise

In astrophotography, it's all about the noise. Every effort is made to minimize it: thermally cooled sensors, dark frame removal techniques, long integration times and multiple exposures are all put to bear on the problem.

Depending on the subject and the imaging conditions, different noise sources dominate. In some situations detector noise is the problem. This is the case when the sky is very dark compared to the scene, a condition that is becoming increasingly compromised by light pollution. When the sky is *not* dark, this background level overwhelms the sensor thermal and fixed pattern noise (in a good sensor, including those in modern digital SLRs). Each case (dark sky, bright sky) requires slightly different treatment when blending frames to make an HDR image.

Conventionally, when individual frames are blended to form an HDR image, the near-black and the near-white pixels are rejected, as their information will be provided by some other frame in the set. The others are accepted, scaled by the exposure factor, and given a weighting that indicates its importance compared to corresponding pixels from other frames that might also contribute to that image location.

A variety of weighting functions have been proposed over the development of HDR [6, 9, 11], and the subject will be discussed further below, but for now we note that we must not discard the near-black pixels too easily, since they contain much if not most, of the information we seek in an astronomic scene. The importance we give the near-black pixels is a delicate balance between signal level and noise level. Because the pixel value is scaled by the exposure factor (and this factor can be very large), it is critical to avoid scaling the noise rather than the signal.

3.1. Sensor noise

There are a number of noise sources associated with image sensors, but one of the important ones is dark current, the signal generated when no light is present. It is strongly temperature dependent. There are others as well, amplifier noise, read noise, hot and cold pixels, fixed pattern noise, cosmic ray hits, and other noise sources that may be specific to a sensor's technology or its integration in an imaging system.

A section obtained from a frame taken with a Canon EOS-20Da is shown, highly amplified (figure 2). The noise in the Canon sensor is remarkably low, but we want to find its distribution so that we can estimate the likelihood that a given pixel amplitude represents signal rather than noise. If we look at the histogram, we can see that the bulk of the noise appears to be Gaussian with a small mean and sigma, but there is a residual non-gaussian "tail" that extends to much larger values.

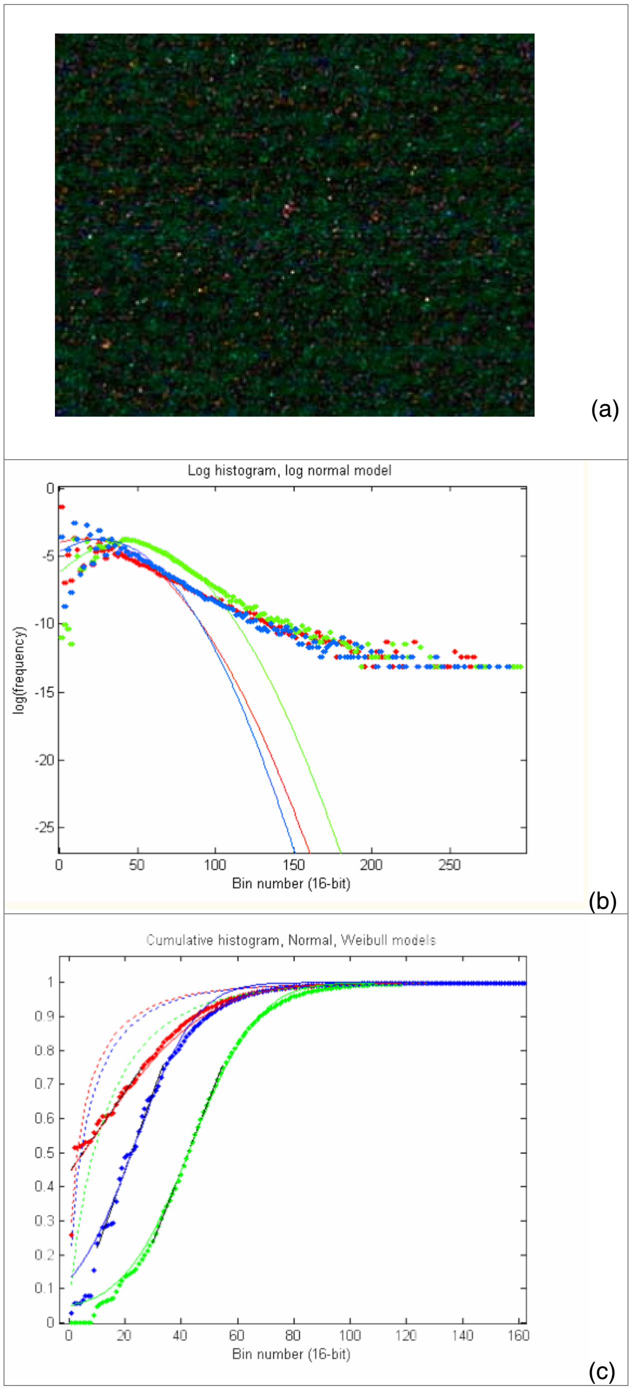

Figure 2. a) A 320x256 section from the background of a 1-sec exposure made by a Canon EOS-20Da at f/9.9, ISO-800. The values have been multiplied by 255 from their 16-bit linear representation to show the visual character of the noise.
b) To see the behavior in detail, a log of the histogram is taken. It is obvious that the fitted normal distribution falls off too rapidly compared with the observed noise.
c) A normal Gaussian (solid) can be fit to the early part of the distribution function, but the tail is best fit by a Weibull model (dashed). A blend of the two noise processes is needed to match the data.

The exponential falloff can be modeled as a Weibull distribution [7], and blended with the Gaussian that matches the dominant component. The result is a probability distribution that can be used to estimate the likelihood that a pixel of a given amplitude is the result of noise. It will be used to provide a weighting factor later when the frames are merged.

The modeled cumulative noise distribution transitions from a normal distribution, F_1, to a Weibull distribution, F_2. The parameters can be obtained by examining characteristics of the cumulative histogram of noise data.

$$\begin{aligned} F_1 &= \Phi(\mu_1, \sigma_1), \\ F_2 &= 1 - \exp\left(-(a_2 x)^{\gamma_2}\right), \\ F_n &= \alpha(x) F_1 + (1 - \alpha(x)) F_2 \end{aligned} \qquad (1)$$

Where $\alpha(x)$ is a sigmoidal blending function whose 50% point is set where the average of the two components matches the observed distribution.

3.2. Sky background (er, foreground) noise

When the sky is not dark, the sensor noise becomes largely irrelevant, being swamped by the signal from the sky itself. This is noise in that even though it is coming from the layer of air we are looking through, we wish it wasn't there. Technically, this is a foreground object, but the appearance and perception is that it is part of the "background".

The noise associated with this illumination will be shot noise, and the Poisson statistics are approximately Gaussian at the levels encountered in light polluted skies. A typical image histogram of an astrophoto (figure 3) is dominated by the sky color, dark, but not black, and also not neutral, not much different from taking a picture of a uniform gray card under unbalanced lighting. Most of the spread in the histogram is the shot noise, proportional to the square root of the signal level. As one takes frames at varying exposures, the offset and the spread will change accordingly. At half the exposure, the offset will be half, and the spread (sigma) will be 0.7x. This behavior will affect how we blend successive exposures and how we perform tone mapping on the final HDR image.

Copyright 2007 Society for Imaging Science and Technology

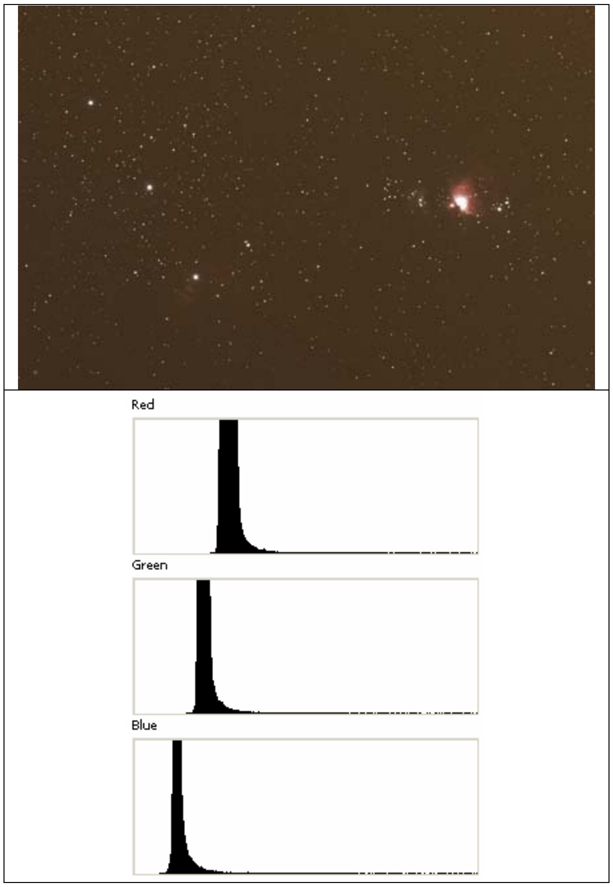

Figure 3. A typical raw astrophoto is dominated by the background sky color, whose histograms are shown here. The sky is not neutral; usually it is the color of local light pollution. The scaling of the noise with signal can be seen here, broadening the distribution for channels with higher amplitudes.

4. Weighting functions

When frames are merged into an HDR image, their individual pixel values are scaled by the frame exposure to bring them into a common numerical coordinate system, what has been referred to here as the "base exposure" (the longest exposure). The scaled values are then averaged by adding them up and dividing by the number of frames. Of course this doesn't work when pixels are saturated, so we apply a weighting in order to reject them. In general, we would like to weight the pixel values according to how much we "trust" them to represent the desired signal, instead of representing noise or other artifacts. The merge operation to obtain pixel intensity (level) I, can be written as:

$$I = \frac{\sum_q W(I_q) \cdot (1/e_q) \cdot I_q}{\sum_q W(I_q)} \quad (2)$$

where I_q represents a pixel from frame q, having an exposure e_q. The weighting function $W(I_q)$, depends on the numeric values from the frame. Mitsunaga and Nayar identified a weighting function derived from the camera sensitivity function for raw camera values as they are converted to radiometric intensities, assuming a constant noise variance [6]. We will be expanding on that result, taking into account our understandings of the noise distribution of the camera, and the desire to not increase the noise levels of the sky background or of signal quantization.

Specifically, we assume the image frames have already been converted to intensities, and absorb their weighting function as a component in a cascade of such weightings:

$$W(i) = W_n(e, i) * W_{sky}(e, i) * W_{sat}(i) \quad (3)$$

The individual terms are weightings for sensor noise W_n, sky background W_{sky}, and saturation avoidance W_{sat}. We will focus on the forms for the first two terms in the chain.

4.1 W_n

Consider first the noise weighting. As successive frames are amplified by their exposure ($1/e_q$ is always greater than 1), the noise found at the lowest levels will sneak into the visible portions of our image. To avoid this, we ask the question, "Is this (original) pixel level more likely to be noise, or signal?" The likelihood of it being due to dark noise can be derived from the distribution function found earlier (equation 1). A useful ratio can be formed with the probabilities of a pixel being due to signal P_s, or to noise P_n:

$$R(i) = \frac{P_s(i)}{P_s(i) + kP_n(i)} \quad (4)$$

When k=1, and the probabilities of noise and signal are equal, R is 0.5. When the probability of noise is dominant, R tends toward zero; with strong signal, R approaches 1. Parameter k can be used to tune the desired signal to noise result. Not knowing the signal statistics, P_s is set to a uniform distribution (all signal levels are considered equally likely). Our desired weighting can now be expressed:

$$W_n(i) = \frac{1}{1 + k \cdot 2^N \cdot f_n(i)} \quad (5)$$

Where N is the number of bits used to represent the pixel level in a frame, and $f_n(i)$ is the noise mass density, the derivative of the distribution F_n, integrated over a digital increment at level i.

4.2 W_{sky}

The other weighting factor depends on the intensity level associated with the background sky. If there is no significant component from the sky, this term can be set to unity, but when there is a "sky fog" level, we want to ensure that adding subsequent frames does not decrease the signal to noise of the sum. It is well known that we can add additional frames at that same exposure level, and the signal to noise ratio will increase by the square root of the number of such frames added. This does not hold true when the frames are scaled from shorter exposures.

If we make a weighted sum of two variables $ax_1 + bx_q$, the mean and variance of the result will be

$$m_w = am_1 + bm_q$$

$$\sigma_w^2 = a^2\sigma_1^2 + b^2\sigma_q^2 \quad (6)$$

Set a=1 to represent the base exposure. What is the limit on weight b when exposure frame data x_q is added such that the signal to noise ratio of the weighted sum m_w/σ_w does not exceed the signal to noise ratio of the base frame m_1/σ_1? The mean and variance of x_q is (with base exposure e_1 set to 1):

$$m_q = e_q m_1$$

$$\sigma_q^2 = e_q s_1^2 \quad (7)$$

The constraint for weight b is:

$$b \le \frac{2}{1 - e_q} \quad (8)$$

As expected, we may add additional *equal* exposures (at any weighting!) and not decrease the signal to noise, but once e_2 is smaller, we cannot add it with such impunity. We would like to add it with weight ($1/e_q$), to match the data scale of the base exposure, and we can, for a while. When e_q is ½, we can scale the frame pixel data by 2, and still come out ahead; the signal to noise ratio of the sum is improved, but by the time the exposure drops to 1/3, we have reached the limit. This is the shortest exposure that can be added with its full exposure multiplier (3) and not worsen the signal to noise ratio.

This helps identify our weighting term W_{sky}, which operates on $(1/e_q)I_q$ (in equation 2):

$$W_{sky}(e_q) = 1 \qquad e_q > \frac{1}{3}$$

$$W_{sky}(e_q) = \frac{b_{\max}(e_q)}{(1/e_q)} = \frac{2e_q}{1 - e_q} \qquad e_q \le \frac{1}{3} \quad (9)$$

We could conceivably apply this exposure-specific weight to each subsequent frame, but for the large signals, all this accomplishes is to make both numerator and denominator equally smaller. When the pixels in the base exposure have saturated, this weighting is moot.

Alternatively, we could recognize that we are mostly interested in protecting the noise amplitude when near the background sky level. We also know that in the base exposure, there is no information below this. We can therefore make a weighting for those smaller pixel levels that corresponds to where the sky is found in subsequent exposures. This results in a function that is independent of exposure, but is now indexed by pixel value. A plot of this weighting is shown in figure 4 for a specific background sky level.

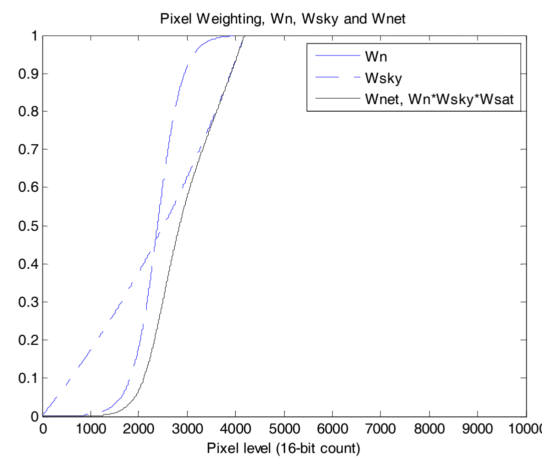

Figure 4. W_{sky} is the maximum weighting we wish to apply to images whose exposure is less than the base exposure in order to maintain signal to noise ratio in the HDR sum. It is combined with W_n, our confidence that the pixel represents signal rather than noise of the sensor. W_{sky} is independent of exposure, whereas W_n depends on the exposure of the frame being added.

5. Blackpoint Normalization

As described earlier, the light polluted sky is an unwanted component in an astrophoto. Whereas much color processing of conventional images concerns the whitepoint and white balance, astrophotographs require proper blackpoint processing for their best presentation. When an HDR image has been assembled, "black normalization" is an important part of the tone scale operator that is applied to make a rendering of the image.

Referring back to the histogram in figure 3, the differences in peak locations represents the distinct color of the sky through which the photograph was taken. We want to remove the light added to the exposure by the sky, but it is important to not simply offset the peaks to zero. This would result in visual artifacts by hard clipping to black. Even though the spread in the histogram mostly represents noise, there is also signal from faint objects in the scene (the reason for the asymmetry in the skirts of the distribution). We need to preserve the noise in order to have the image show pleasing shadow detail, but remove the offset introduced by the light of the sky.

To do this, identify the minimum signal levels found in the image. This will not be a constant, especially in the presence of

Copyright 2007 Society for Imaging Science and Technology

light pollution, which shows as an illumination gradient across the frame. Several image processing operators can obtain an estimate for the sky, but the basic concept is find the minimum pixel levels over a local neighborhood. A min filter with a range of around 1/10th the image size works well. An image that represents an approximation to the background sky is obtained.

We now subtract the sky offset image from the HDR data. This effectively removes all of the (empty) bins in front of the image data. The resulting histogram now reaches to black, but is still not balanced; the peaks of the histogram retain the color of the sky, if not its brightness.

The next step is to identify those peaks and apply a transfer function that aligns them. The maximum is used as the target value, avoiding loss of data from bin merging (we can always compress this region later if we wish). It is important to preserve the unity slope of the mapping once we have achieved the peak balance. Our goal is to remove the color cast of the sky, without impacting the color relationships of the remaining objects in the scene. Plots and resulting histograms are shown in figure 5.

6. Tone mapping astronomical images

The assembled and blackpoint-normalized HDR image is still in need of proper rendering to a display or print. Since there is no "true appearance" of these scenes, there are no rules for technically correct tone mapping. The goals of the image author must guide the selection of possible renderings. Spatially varying tone operators can be highly successful in this task, but what will be described here is a pair of simple, global operations that also yield acceptable results.

A common goal in astrophotography is to show the detail in faint nebulas without excessive grain noise, and without clipping the bright areas. Once the black level is established, a simple gamma function serves this purpose. Empirically, one can find the exponent that brings the bulk of the histogram into the lower half of the tone scale, leaving the residual tail (usually comprising stars) to reside in the top half. This is then followed by an application of an S-curve, which accomplishes several things: it amplifies the detail in the faint objects, compresses the noise near black (without clipping it), and compresses the brightest stars in the scene, without clipping them.

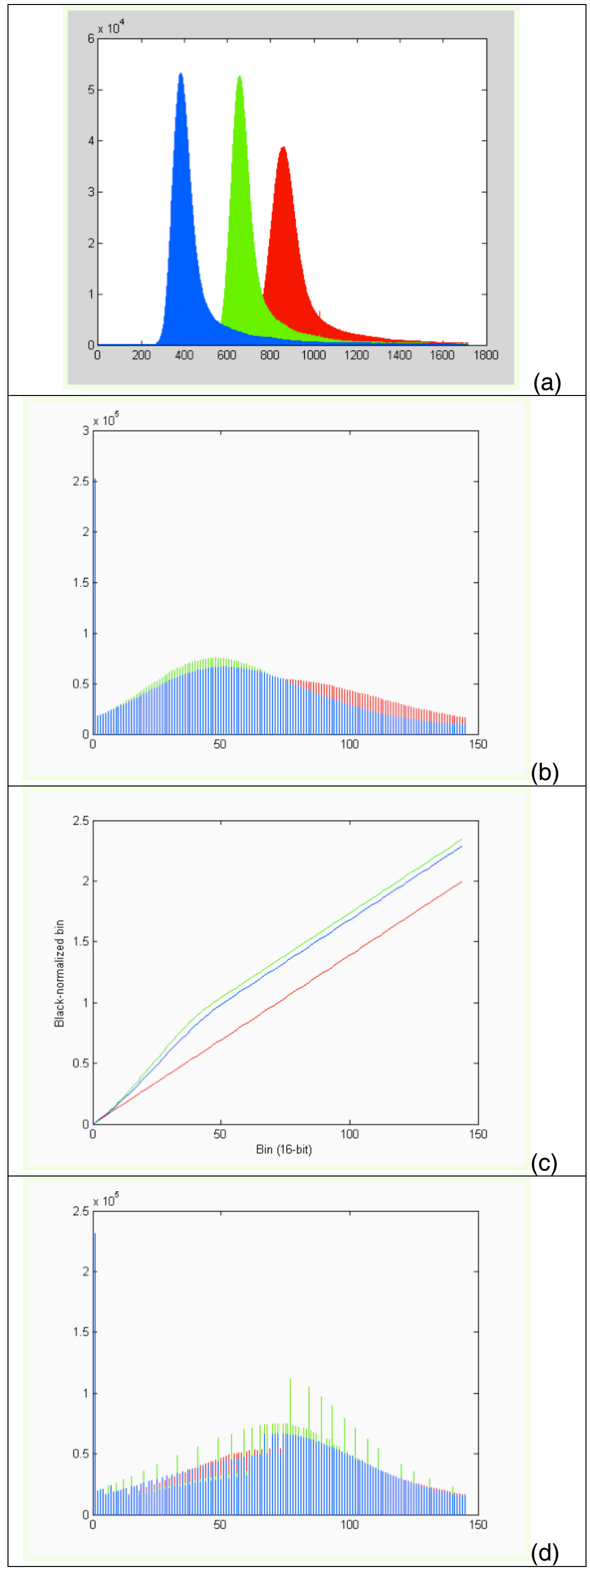

Figure 5. a) *These are the low bins in a 16-bit linear scale for an HDR image (M42, example 2).* ***b)*** *The min image is subtracted from the original. All of the original data is still present (apart from some minor clipping at zero), we have removed the bins with no counts.* ***c)*** *A transfer function maps the peaks of the offset sky to the same (highest) position, preserving relative bin positions at the peak and continuing at unit slope thereafter.* ***d)*** *The result shows the channel histograms nearly superposed (albeit with bin counts re-distributed accordingly).*

Summary

A number of issues have been discussed that require special care when creating high dynamic range images of astronomical scenes. Some of these are already part of the assembly of any HDR image:

- Image registration requires more than just offset, rotation and scaling because stars are point sources and optical systems are not rectilinear.

- Motion artifacts include deblurring long exposures. It must be done in a way that preserves the state of saturated pixels.

In addition, there are new issues that are specific to astronomical images:

- Low amplitude pixels are not rejected and must be carefully weighted according to the noise characteristics of the imaging sensor.

- The background sky level sets another type of weighting that needs to be applied in order to avoid unwanted noise that would result from the scaling of pixel values.

- When present, unwanted background sky must be offset and neutralized to obtain a pleasing rendering of the scene.

Some examples of applying these methods are shown.

References

[1] J. M. Coggins, L. K. Fullton, B. W. Carney, "Iterative/Recursive Deconvolution with Application to HST Data," The Restoration of HST Images and Spectra II, Space Telescope Science Institute, 1994, R. J. Hanisch and R. L. White, eds.

[2] P. E. Debevec and J. Malik. "Recovering High Dynamic Range Radiance Maps from Photographs," SIGGRAPH 97 Conference Proceedings, pp. 369-378, ACM SIGGRAPH, August 1997.

[3] P. Favaro, M. Burger, and S. Soatto. *Scene and motion reconstruction from defocused and motion-blurred images via anisotropic diffusion*. In Proc. 8th European Conference on Computer Vision (ECCV'04), pages 257--269, Prague, May 2004..

[4] R. Fergus, B. Singh, A. Hertzmann, S. T. Roweis, W. T. Freeman, "Removing Camera Shake from a Single Photograph," SIGGRAPH 2006 Conference Proceedings, ACM SIGGRAPH, August 2006

[5] L. B. Lucy, "An iterative technique for the rectification of observed distributions," Astronomical Journal, v 79, n 6, June 1974.

[6] T. Mitsunaga and S. K. Nayar. "Radiometric Self Calibration," in Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, Fort Collins, CO, IEEE, June 1999.

[7] NIST/SEMATECH e-Handbook of Statistical Methods, http://www.itl.nist.gov/div898/handbook/ ,7/18/2006, "A Gallery of Distributions," http://www.itl.nist.gov/div898/handbook/eda/section3/eda366.htm.

[8] RegiStar, Auriga Imaging, http://www.aurigaimaging.com

[9] E. Reinhard, G. Ward, S Pattanaik, P. Debevec. *Hugh Dynamic Range Imaging*, Morgan Kaufman Publishers, San Francisco, CA, 2006.

[10] R. L. White, "Image Restoration Using the Damped Richardson-Lucy Method," *The Restoration of HST Images and Spectra II*, *Space Telescope Science Institute, 1994*, *R. J. Hanisch and R. L. White, eds.*

[11] F. Xiao, J. E. Farrell and B. A. Wandell, "Psychophysical thresholds and digital camera sensitivity: the thousand photon limit"

Author Biography

Thor Olson's day job is at Electronics for Imaging, where he applies his background in physics and signal processing to problems in imaging, printing, color measurement, and color management.

At night however, he attempts astrophotography and can sometimes be found outside, in a field, in the dark, studying the view through his telescope, wondering about how his camera will record the faint objects of the night sky. Some of his work can be found at his website www.nightscapes.net.

Copyright 2007 Society for Imaging Science and Technology

Example 1: Vega

A set of exposures ranging from 30 seconds to 1/2000 second were deblurred, (manually) registered, and combined using weightings that prevented low level sensor noise from being amplified.

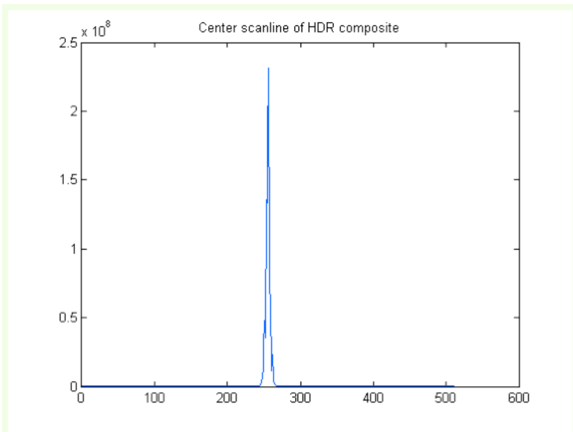

a) A plot of the central scanline in the HDR composite of the 0.3 magnitude star Vega (red channel). This is essentially a plot of the point spread function of the optical system. Note the scale is in digital counts x 10^8.

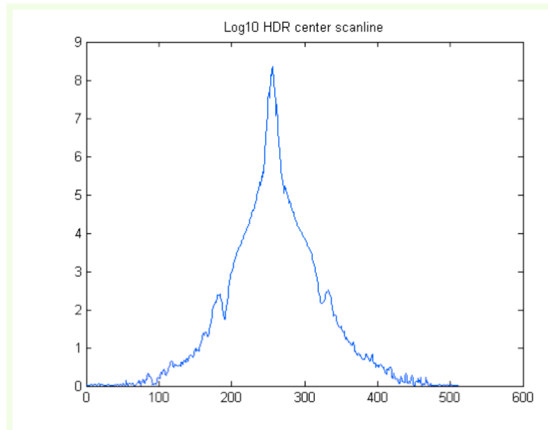

b) A log plot of the same scanline revealing flare and diffraction spikes.

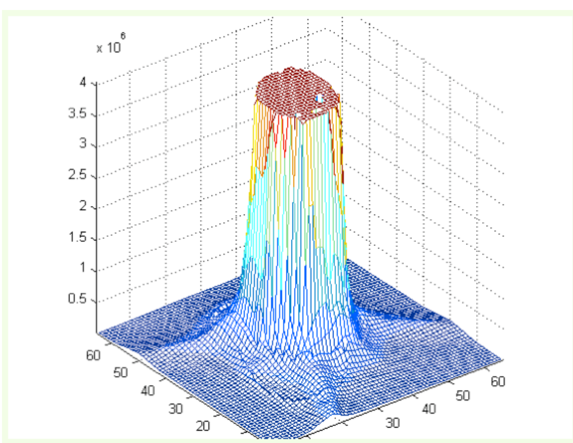

c) The bottom 1-percent of the point spread function.

d) A rendering of Vega. The colored diffraction spikes are an artifact of the sensor's color sampling mosaic. The loops and contours at the core are a result of imperfect de-blurring and alignment. The 9th and 10th magnitude stars are visible in the same view as Vega, a 0-magnitude beacon. This scene shows a dynamic range of objects whose brightness differs by 100,000:1.

Examples 2 and 3: M42, the Orion Nebula

Here are some exposures of M42. The equipment was a Canon EOS 20Da on a Televue 85 operating at f/5.6. Camera settings were at ISO 1600, fixed white balance at 6500K. The exposures range from 30 to 300 seconds, with the longest ones clearly showing the "fog level" of the background sky. Even though these were taken on a dark night in remote Arizona, the sky is not black, and it shows up in these exposures.

			
30 sec	60 sec	120 sec	300 sec

Example 2: Composited HDR of M42, the Orion Nebula.
Motion deblurred, star-registered, sky background removed, gamma and S-curves applied. Nebula detail is visible while avoiding highlight clipping

Example 3: More Nebulas in Orion.
A wider view showing the belt and sword stars of this constellation utilizing the methods described in the paper. In addition to M42, the Flame, and Horsehead nebulas are revealed.

Copyright 2007 Society for Imaging Science and Technology