SANDIA REPORT
SAND2019-2371
Printed March 2019

Sandia National Laboratories

SAR Image Scaling, Dynamic Range, Radiometric Calibration, and Display

Armin W Doerry

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico
87185 and Livermore,
California 94550

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology & Engineering Solutions of Sandia, LLC.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@osti.gov
Online ordering: http://www.osti.gov/scitech

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Rd
Alexandria, VA 22312

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.gov
Online order: https://classic.ntis.gov/help/order-methods/

Abstract

Once Synthetic Aperture Radar (SAR) images are formed, they typically need to be stored in some file format which might restrict the dynamic range of what can be represented. Thereafter, for exploitation by human observers, the images might need to be displayed in a manner to reveal the subtle scene reflectivity characteristics the observer seeks, which generally requires further manipulation of dynamic range. Proper image scaling, for both storage and for display, to maximize the perceived dynamic range of interest to an observer depends on many factors, and an understanding of underlying data characteristics. While SAR images are typically rendered with gray-scale, or at least monochromatic intensity variations, color might also be usefully employed in some cases. We analyze these and other issues pertaining to SAR image scaling, dynamic range, radiometric calibration, and display.

Acknowledgements

This report was funded by General Atomics Aeronautical Systems, Inc. (GA-ASI) Mission Systems under Cooperative Research and Development Agreement (CRADA) SC08/01749 between Sandia National Laboratories and GA-ASI.

General Atomics Aeronautical Systems, Inc. (GA-ASI), an affiliate of privately-held General Atomics, is a leading manufacturer of Remotely Piloted Aircraft (RPA) systems, radars, and electro-optic and related mission systems, including the Predator®/Gray Eagle®-series and Lynx® Multi-mode Radar.

Contents

List of Figures

List of Tables

Acronyms and Definitions

CCD	Coherent Change Detection
CDF	Cumulative Density Function
CNR	Clutter to Noise Ratio
EO	Electro-Optic
HS	Hyper-Spectral
HVS	Human Visual System
IEEE	Institute of Electrical and Electronics Engineers
ISAR	Inverse SAR
MNR	Multiplicative Noise Ratio
NGA	National Geospatial-Intelligence Agency
NITF	National Imagery Transmission Format
PDF	Probability Density Function
RCS	Radar Cross Section
RMS	Root Mean Square
SAR	Synthetic Aperture Radar
SI	Système International (International System of Units)
SICD	Sensor Independent Complex Data
SNR	Signal to Noise Ratio

Foreword

This report details the results of an academic study. It does not presently exemplify any modes, methodologies, or techniques employed by any operational system known to the authors.

Classification

The specific mathematics and algorithms presented herein do not bear any release restrictions or distribution limitations.

This report formalizes preexisting informal notes and other documentation on the subject matter herein.

Author Contact Information

Armin Doerry awdoerr@sandia.gov 505-845-8165

1 Introduction and Background

We discuss herein Synthetic Aperture Radar (SAR) images. While other radar modes such as Inverse SAR (ISAR) also form similar images, with similar issues, we nevertheless confine ourselves in this report to SAR images.

SAR scenes may offer dynamic ranges (defines as the ratio of brightest target allowable in the scene without saturation to the noise level) of many orders of magnitude, well beyond 100 dB or so. For example, a 0.1-m resolution SAR image with a noise equivalent reflectivity of −35 $dBsm/m^2$, and allowing for a 60 dB non-saturating target requires a dynamic range of 115 dB. Finer resolutions, lower noise floor, or brighter non-saturating targets will increase this. A SAR image is simply a rendering of the scene being interrogated, with pixel values some function of Radar Cross Section (RCS) values.

Modern high-performance SAR systems will typically perform image formation using double-precision floating-point calculations. However, a final step is often to convert the resulting image to something else, often an integer representation, for storage and transmission. Several standards exist for SAR image formats, including

National Imagery Transmission Format (NITF), with latest version 2.1, and

Sensor Independent Complex Data (SICD).

Appendix B discusses some aspects of the SICD format.

Regardless of what an image format may represent, when displaying a SAR image for human observation, we are further encumbered by the limits of the Human Visual System (HVS). While there is much debate over what the HVS can perceive in a single image, as this is a very complex subject and a number of dependencies must be considered, commonly quoted numbers suggest the HVS can only perceive about 30-40 dB of dynamic range.[1,2] Display technology is understandably keenly aware of the HVS characteristics, or should be.

Consequently, regardless of the dynamic range of the original SAR image calculations, they need to be fit into some image format. Then, regardless of the image data, we need to contend with the limits of the HVS if those images are to be displayed for human consumption and exploitation.

We address herein two issues.

1. What should be the scaling of a SAR image in the output image data, and
2. How should a SAR image be displayed?

Image Power Measures

While the radar engineer tends to work with measures of energy and power, with respective units of Joules and Watts,[†] the Optical and Display communities tend to work with other measures and units. We address some of them here.

Irradiance – a.k.a. intensity, optical intensity – is the optical power per unit area; a power density. In the SI system, it has units of watts per square meter (W/m^2).

Luminous flux – is a measure of the perceived power of light. The SI unit of luminous flux is the lumen (lm). The luminous flux is the total radiant flux weighted by a model of the sensitivity of the human eye, called the luminosity function.

Illuminance – is the luminous flux per unit area. The SI unit of illuminance is the lux (lx). One lux equates to 1 lumen per m^2. For a wavelength of 555 nm (green), which is the peak of the luminosity function, 683.002 lux equates to 1 W/m^2.

Luminous intensity – is the luminous flux (power) per solid angle. The SI unit of luminous intensity is the candela (cd). One candela equals 1 lumen per steradian.

f-stop – a.k.a. stop, f-number, exposure value (EV) -- is the ratio of the focal length to the effective aperture. This means that an increase of one f-stop will decrease the aperture by half, and hence the intensity by half; two f-stops will decrease the intensity to a quarter, etc. The intensity as a function of f-stop is a geometric progression.

[†] Note that RCS is a ratio of re-radiated (i.e. reflected) power (W) to incident power density (W/m^2).

2 Image Representation

While it is easy enough to conjure, or even optimize, a pixel representation for a particular complex image (where each pixel is a complex-valued element that can be represented with real and imaginary components, or equivalently with magnitude and phase components),[3,4] we are often not given the luxury of our choice. We are often left with "How can we cram an image into a representation given to us, like perhaps 16-bit integer real and imaginary data. We further stipulate that most often, such requirements include a linear magnitude representation, where RCS is proportional to the square of the pixel magnitude, forgoing any companding.[3,4]

Given that a SAR scene's dynamic range is bounded by the noise floor on the low end, and the brightest reflector on the high end, there is little sense in the image rendering (allocating pixel values to) levels to below the noise floor or above the brightest expected target. We allow that other characteristics of the data might modulate this somewhat.

How noise levels relate to RCS, and hence pixel values, is a function of the radar equation, which can be found in many texts and papers, including a previous report.[5] We stipulate here that desired maximum noise-equivalent reflectivity, σ_N, for various target scenes at Ku-band might be as follows:

General image analysis	–25 dBsm/m^2
Coherent Change Detection (CCD)[6]	–35 dBsm/m^2
Oil-spill mapping[7]	–50 dBsm/m^2, or better

These values should be considered indicative, and not hard rules. We recognize that the noise level needs to be far enough below the scene's clutter or target level so that we can properly exploit the image, which often includes assessing characteristics of shadow areas (areas of no return). This, of course, is dependent on the actual scene being imaged. Average clutter levels for various scene types has been cataloged by Long,[8] and Ulaby and Dobson,[9] among others.

Furthermore, the average clutter level often even sets the effective noise floor via multiplicative noise. Multiplicative noise is manifested by signal channel nonlinearities, which impact sidelobe responses and cause energy to spill away from the desired pixel location. This limits how "dark" small shadow regions can be relative to the surrounding clutter. Since this offending noise energy is proportional to signal level, it is termed "multiplicative." Even high-quality SAR systems may exhibit a Multiplicative Noise Ratio (MNR) of perhaps –20 dB or so, meaning an average noise floor due to this phenomenon of perhaps no lower than 20 dB below the average clutter level.

From Raynal, et al.,[10] we are given that at Ku-band a reasonable upper bound for cultural clutter discretes is perhaps 45 dBsm, although we can certainly conjure brighter targets. We note that even a 60 dBsm target is not a particularly unimaginably large reflector, achieved with a 2.5-m diameter circular flat plate, or a 3-m trihedral.[11] Neither of these dimensions is unreasonable for man-made structures and objects.

Equally important is the distribution of RCS values between these two limits. This entails an understanding of the nature of clutter in the scene. We note that a scene of uniform distributed clutter, by virtue of the Central Limit Theorem, processed with linear calculations into a SAR image, will manifest as complex Gaussian distributed output values. The magnitude of these will be Rayleigh-distributed (see Appendix A), with mean-squared value proportional to mean RCS of the pixel, which itself is proportional to the normalized clutter reflectivity, σ_0.

Looking ahead a little, the HVS needs about 100 intermediate gray levels over its dynamic range to interpret the image as having smoothly varying shades.[12] We might interpret this as suggesting that the bulk of the distributed clutter span at least 100 magnitude pixel values (counts). This would place the RMS clutter level no lower than a magnitude pixel value of perhaps 30 counts or so.

We summarize this discussion with some mathematics. We accordingly define

σ_0 = normalized clutter reflectivity with units m^2/m^2,
σ_N = noise-equivalent reflectivity with units m^2/m^2,
ρ_r = slant-range resolution of the radar with units m,
ρ_a = azimuth (cross-range) resolution of the radar with units m, and
ψ = local grazing angle. (1)

Normalized reflectivity values are commonly expressed in units of $dBsm/m^2$ which is $10\log_{10}(\,)$ of the reflectivity values with the units cited in Eq. (1).

A pixel dominated by clutter will have a mean RCS value of

$$\sigma_{clutter} = \frac{\rho_r \rho_a \sigma_0}{\cos\psi}. \tag{2}$$

A pixel dominated by noise will exhibit a mean RCS value of

$$\sigma_{noise} = \frac{\rho_r \rho_a \sigma_N}{\cos\psi}. \tag{3}$$

Note that both $\sigma_{clutter}$ and σ_{noise} are mean-squared values of a Rayleigh-distributed random variables. We define the ratio of $\sigma_{clutter}$ to σ_{noise} as

$$CNR = \frac{\sigma_{clutter}}{\sigma_{noise}} = \frac{\sigma_0}{\sigma_N} = \text{Clutter to Noise Ratio (CNR).} \tag{4}$$

This is analogous to Signal-to-Noise Ratio (SNR), except that the "signal" in this case is the clutter field. This, of course, ignores the contribution of MNR.

For the clutter to be observable, and exploitable, in the SAR image, we need a substantially positive CNR. While any specific number might be subject to some debate, we opine that 10 dB is not an unreasonable minimum CNR for shadows to be exploited, and for reasonably acceptable CCD. That is, we would like the noise level to be at least 10 dB below the lowest clutter reflectivity we wish to exploit. We note that two integer bits is about 12 dB. We are also mindful that an upper limit on CNR is likely due to multiplicative noise, and might be perhaps 20-25 dB for typical high-quality SAR images of largely uniform clutter scenes.

Of course, a SAR scene is often composed of more than simply uniform clutter, quite often containing patches of locally uniform clutter with varying reflectivity characteristics, and can be expected to also contain clutter discretes, as previously discussed. Nevertheless, a Rayleigh characteristic to the pixel magnitude remains a reasonably good approximation, albeit perhaps with some extra energy in the tail of the distribution due to clutter discretes.

W now examine a representative SAR image, illustrated in Figure 1.

Figure 1. Example SAR image.

Table 1. Image parameters for Figure 1.

Center frequency	16.7 GHz
Slant-range resolution	0.1016 m
Cross-range resolution	0.1016 m
Noise Floor σ_N	–30 dBsm/m^2
Nominal Slant-Range	6.858 km
Grazing angle at aperture center	30.3545 deg

Figure 1 renders a SAR image of a Sandia National Laboratories test facility in Albuquerque, New Mexico, USA. The data was collected on 30 March 2001 by a Sandia Ku-band testbed system. Some image parameters are given in Table 1.

We note that the image of Figure 1 shows a fenced compound of buildings surrounded by natural clutter typical of the high Chihuahuan desert common in New Mexico. We observe parked vehicles as well as sparse vegetation. We further note that the bulk of the pixels are of distributed clutter rather than specular reflectors (clutter discretes).

While pixel values in Figure 1 have been manipulated to display the subtle clutter features, the histogram of the underlying 16-bit magnitude data is illustrated in the plots of Figure 2. We observe in the top plot that pixel values are preponderantly at the very lower end of the magnitudes. The middle plot shows a tail of higher pixel magnitudes corresponding to clutter discretes, including two pixels that saturate at the 65535 maximum value. The bottom plot details the Rayleigh characteristic of the natural clutter pixel magnitudes. Noise for this image has an RMS magnitude pixel value of about 12.7. Figure 3 renders the histogram in terms of RCS values with units of dBsm. This calibration can be calculated from system parameters.[13]

In the histogram rendered in Figure 3, the shape has been distorted due to pixel RCS in units of dBsm being a logarithmic function of pixel magnitude. With pixel statistics available via the histograms of Figure 2 and Figure 3, we make the further observations for this specific image as follows,

- The noise level is sufficiently represented in the magnitude image pixel data, with an RMS value of about 12.7, corresponding to a noise equivalent reflectivity of –30 dBsm/m^2, and to an average RCS of about –50 dBsm. Noise contribution is overwhelmed by the clutter, as is expected in the image by noting that collective shadow area is considerably less than collective natural clutter area.

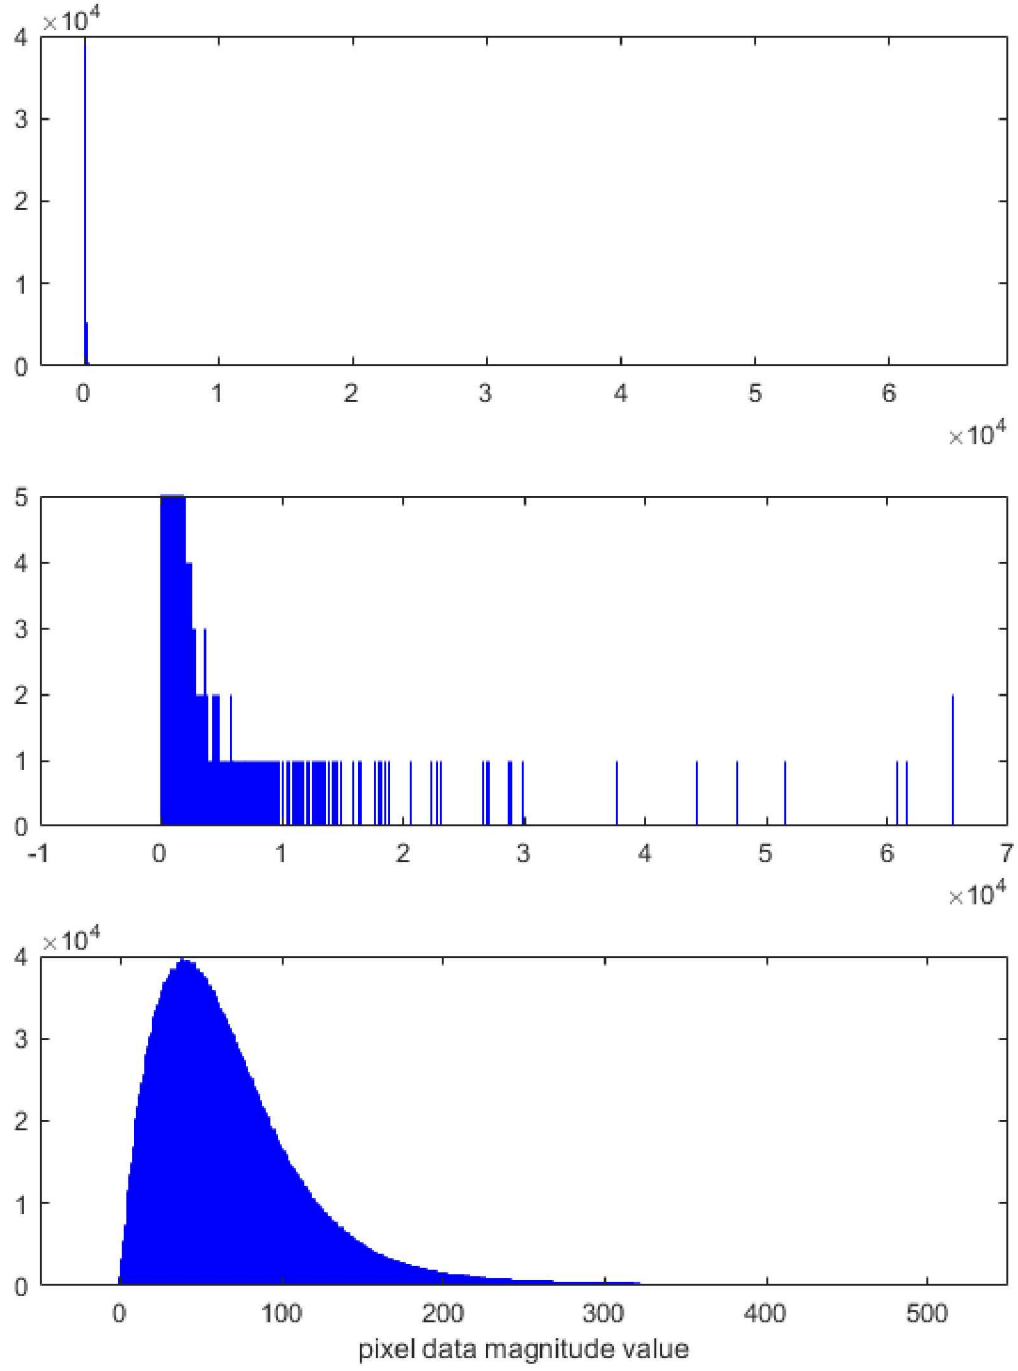

Figure 2. Histograms of magnitude pixel data. Pixel magnitude values range from 0 to 65535.

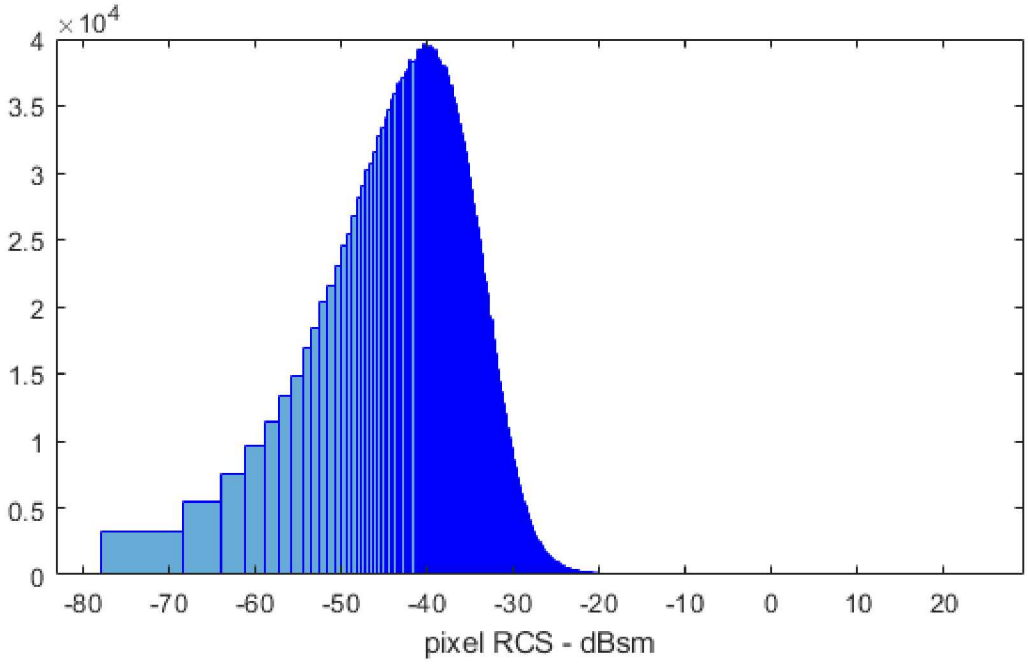

Figure 3. Histogram of magnitude pixel data converted to RCS values.

- The natural clutter has a median reflectivity of about −17 dBsm/m^2, resulting in a median RCS of about −37 dBsm in the image, suggesting an average CNR of about 13 dB. The median pixel magnitude value is 58 counts. The median statistic was used to minimize the influence of outliers; especially bright discrete reflectors. Its relationship to other Rayleigh statistics is discussed in Appendix A.

- The maximum discrete clutter RCS in this image saturates the magnitude image pixel data at a value of 65535 counts corresponding to 24.22 dBsm. We do note that the imaging geometry for this image is an oblique angle with respect to the cultural features, which would minimize any high-RCS flashes from man-made objects.[‡]

- Although 16 bits of magnitude would allow up to 96 dB of dynamic range, the noise level in this image limits the rendered effective dynamic range to about 74 dB. This suggests that scaling the floating-point data prior to integer conversion to push the noise level to lower pixel values could increase the rendered dynamic range by up to perhaps 20 dB.

All things considered, this image has been fit rather nicely within the 16 bits of dynamic range afforded to it. We might have benefitted slightly by scaling the magnitudes by ½, thereby yielding an extra 6 dB of dynamic range, but it would have mattered to no more than two pixels in the entire image.

[‡] A common 'trick' in setting up SAR data collections is to not use imaging geometries corresponding to cardinal directions. Man-made structures are often oriented in North-South and East-West directions, so imaging at non-cardinal angle geometries will minimize bright flashes from dihedrals formed by the sides of these structures.

2.1 Design Example – Rural Scene

Consider a Ku-band VV polarization SAR image of a rural scene with the following parameters.

Resolution (range and azimuth)	0.3 m
Grazing angle	30 deg.
Mean noise-equivalent reflectivity	–35 dBsm/m^2
Mean clutter reflectivity	–17 dBsm/m^2
Maximum desired clutter discrete	+45 dBsm
Image format	16 bits unsigned magnitude, 16 bits phase
Companding	none, assume linear magnitude

The question is "How do we best represent this image within the allotted bits?" which we interpret to mean "What scale factor should we use to relate pixel value to RCS?"

From these parameters, we calculate (rounded to the nearest dB)

Mean noise RCS	–45 dBsm
Mean clutter RCS	–27 dBsm

Although the difference between the noise level and the maximum clutter discrete is 90 dB, well within the available 96 dB afforded by 16 bits of linear magnitude, if the RMS noise level were placed at 1 count, the mean clutter level would be only 18 dB above the quantization level, which might make the clutter excessively quantized. A better clutter level might be 30 dB above the quantization level, which gives better rendering of small clutter variations. This would place the quantization level (the RCS represented by a magnitude of "1" count) at –57 dBsm. Corresponding relationships would be

Quantization level	1 count	0 dBq[§]	–57 dBsm
Mean noise RCS	4 counts	12 dBq	–45 dBsm
Mean clutter RCS	32 counts	30 dBq	–27 dBsm
Maximum clutter discrete RCS	65535 counts	96 dBq	+39 dBsm

We note that with this scaling, we are only 6 dB shy of our desired maximum clutter discrete value, suggesting slightly more frequent saturation than desired. However, because the maximum clutter discrete RCS is so far out on the tails of the clutter PDF, this is not expected to be particularly troublesome.

[§] We define "dBq" as dB with repect to the quantization level.

We also note that with this clutter level, the MNR will likely keep the actual noise floor greater than 0 dBq, even if the front-end receiver noise drops.

With these relationships, we may accordingly calculate a calibration scale factor to convert any pixel magnitude value to RCS as

$$Sf = \frac{\sigma_{RCS}}{p^2} \approx 2\times 10^{-6} \text{ = RCS calibration scale factor,} \tag{5}$$

where

σ_{RCS} = reference RCS in units m^2, and
p = pixel magnitude value in counts corresponding to σ_{RCS}. (6)

In units of dBsm, the calibration calculation for an arbitrary pixel magnitude value becomes

$$\sigma_{RCS,dBsm} = 10\log_{10}\left(p^2 Sf\right). \tag{7}$$

Of course, this may also be written as

$$\sigma_{RCS,dBsm} = 20\log_{10}\left(p\right) + 10\log_{10}\left(Sf\right). \tag{8}$$

Comments

We note that the mean noise RCS in our example is 12 dB above the desired minimum noise level of 0 dBq. This suggests that if we had brighter clutter, say from an urban scene, then we might increase the scale factor by up to 12 dB, to a maximum of perhaps $Sf \approx 3.16\times 10^{-6}$, and still maintain the desired minimum noise level. Evidence suggests that even a high urban clutter scene might have a mean clutter reflectivity of perhaps –9 dBsm/m^2 at Ku-band for VV polarization.[10] This represents only an 8 dB increase over our design example. Were we to know that we were imaging such an urban scene, all other parameters unchanged, we might select a scale factor of perhaps $Sf \approx 1.26\times 10^{-6}$ to better accommodate it. Note that this is about scaling the "image" data, and not the RF gain for raw data collection.

Of course, all of this is predicated on that we are assuming an image format with pixels limited to 16 bits of unsigned magnitude with no companding. Were we to employ companding, or even a floating-point number representation, then the problem of squeezing our image with its dynamic range constraints into a finite number of bits is considerably eased. We do note also that the IEEE does maintain a specification for "half-precision" floating-point numbers using only 16 bits.[14]

2.2 Design Example – Ocean Oil Slick Detection

Consider a Ku-band VV polarization SAR image of an ocean scene with the following parameters.

Resolution (range and azimuth)	3 m
Grazing angle	10 deg.
Mean noise-equivalent reflectivity	–50 dBsm/m^2
Mean sea clutter reflectivity	–40 dBsm/m^2 (sea-state 1)[15]
Maximum desired clutter discrete	+45 dBsm
Image format	16 bits unsigned magnitude, 16 bits phase
Companding	none, assume linear magnitude

The amount of backscatter reduction due to an oil slick depends on a number of factors, including operating frequency, grazing angle, and polarization. The contrast appears best for VV polarization at about 30 degree grazing angle, decreasing somewhat for shallower grazing angles.[7] A survey of literature suggests that contrast might be from 5 to 15 dB, depending on the factors cited.[16,17] Clutter discretes are essentially maritime vessels in the scene.

The question is still "How do we best represent this image within the allotted bits?" which we interpret to mean "What scale factor should we use to relate pixel value to RCS?"

From these parameters, we calculate (rounded to the nearest dB)

Mean noise RCS	–40 dBsm
Mean sea clutter RCS	–30 dBsm

The difference between the noise level and the maximum clutter discrete is 85 dB, well within the available 96 dB afforded by 16 bits of linear magnitude, if the RMS noise level were placed at 1 count. Let us again place the desired/expected clutter level at 30 dB above the quantization level, which gives better rendering of small clutter variations we seek to map oil slicks. Corresponding relationships would be

Quantization level	1 count	0 dBq	–60 dBsm
Mean noise RCS	4 counts	20 dBq	–40 dBsm
Mean clutter RCS	32 counts	30 dBq	–30 dBsm
Maximum clutter discrete RCS	65535 counts	96 dBq	+36 dBsm

We accept that many maritime vessels are likely to saturate. With these relationships, we may accordingly calculate the calibration scale factor to convert any pixel magnitude value to RCS as

$$Sf = \sigma_{RCS} \Big/ p^2 \approx 1.0 \times 10^{-6}. \tag{9}$$

“It is better to be looked over than overlooked.”
-- Mae West

3 Image Display – Gray Scale

While optimizing the display of 'any' image could consume a lifetime of scientific study, we beg the reader's indulgence to express herein some basic observations and to propose "something that works."

The intent now is to accept a SAR image with typical characteristics as noted in the previous section, and to render it in a manner to convey information to the human observer. We opine that there is probably a balance between maximizing the information transfer, with information defined as a rigorous mathematical calculation, and maintaining some sort of "natural" look to the image. This "natural" aspect is important because it plays to the expectations of an observer, and brings to bear his toolkit of experiences and synergism with other sensors.

Recall that commonly quoted numbers suggest the HVS can only perceive about 30-40 dB of dynamic range. This amounts to 5-7 bits of dynamic range, as previously defined. We also note that a typical display accepts 8-bit magnitude data, with a "perfect" display rendering gray scale luminous flux that increases as the square of the pixel value. So, doubling the applied pixel magnitude value will quadruple the intensity of the image, with 8 bits corresponding to 48 dB of dynamic range. We further note that real displays may exhibit technology-dependent performance limits, and might compress this dynamic range somewhat. We will ignore this particular complication in this report.

The question for us becomes "How do we map the, say, 16 bits of SAR image pixel magnitudes to the 8 bits available in a display?"

Furthermore, we wish to do so in a manner where we are able to discern the natural distributed clutter detail (in the pixel values 1:200, or so, of Figure 1) while maintaining awareness of the brighter discrete RCS values.

We expect this to require a nonlinear scaling of pixel values to adjust the histogram of pixel values. In image processing terms, we expect this to require "Gamma corrections" and some degree of "histogram stretching."[**]

Several mappings are now presented. We define for all of these

p = 16-bit SAR image pixel magnitude value in linear magnitude, and
P = 8-bit "adjusted" pixel value applied to the display device. (10)

For these various mappings we offer both a resulting image as well as the modified histogram. We are well aware of the irony of exemplifying a display monitor with a printed image, but offer that the printed images do convey at least some of the characteristics we wish to illustrate.

[**] "Gamma correction" commonly referes to a non-linear scaling of pixel values usually (but not always) according to some power law, whereas "histogram stretching" more often refers to a linear scaling of pixel values.

3.1 Histogram-Stretched Image

The simplest function to convert 16-bit images to 8-bit images is to simply scale pixel values and crop to the desired range. This is histogram stretching without any gamma correction. For the image of Figure 1, we choose the function

$$P = 255\mu\ p, \tag{11}$$

where we might choose scale factor

$$\mu = \frac{1}{8\ \text{median}(p)}. \tag{12}$$

Of course, the output value P is quantized and limited to the range [0:255]. The resulting image is given in Figure 4, and its histogram is illustrated in Figure 5.

3.2 Quarter-Power Image

Noting that a straightforward mapping that halves the number of required bits is a square-root function, we define the Quarter-Power image as

$$P = 255\beta\sqrt{p}\ . \tag{13}$$

The scale factor β adjusts the brightness of the display. Anecdotal evidence suggests that useful values are in the range

$$\frac{1}{5\ \text{median}\left(\sqrt{p}\right)} \le \beta \le \frac{1}{3\ \text{median}\left(\sqrt{p}\right)}\ . \tag{14}$$

As before, output value P is quantized and limited to the range [0:255]. The resulting image is given in Figure 6, and its histogram is illustrated in Figure 7. The "spikes" result from multiple input pixel values combined into a single display pixel value. Note that the histogram has been substantially stretched, allowing natural clutter to be discerned in greater detail.

The image in Figure 1 was rendered with $\beta = \left[3\ \text{median}\left(\sqrt{p}\right)\right]^{-1}$.

Figure 4. Histogram-stretched image with $\mu = [8\ \mathrm{median}(p)]^{-1}$.

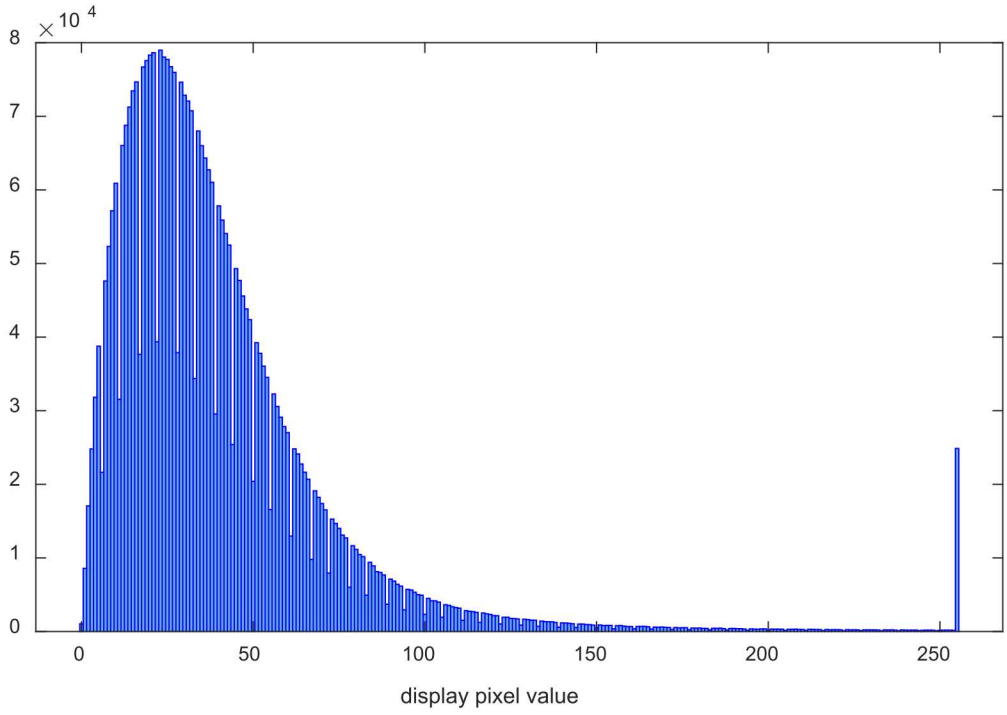

Figure 5. Histogram of display pixel values *P* with $\mu = [8\ \mathrm{median}(p)]^{-1}$.

Figure 6. Quarter-power image with $\beta = \left[3 \ \text{median}\left(\sqrt{p} \right) \right]^{-1}$.

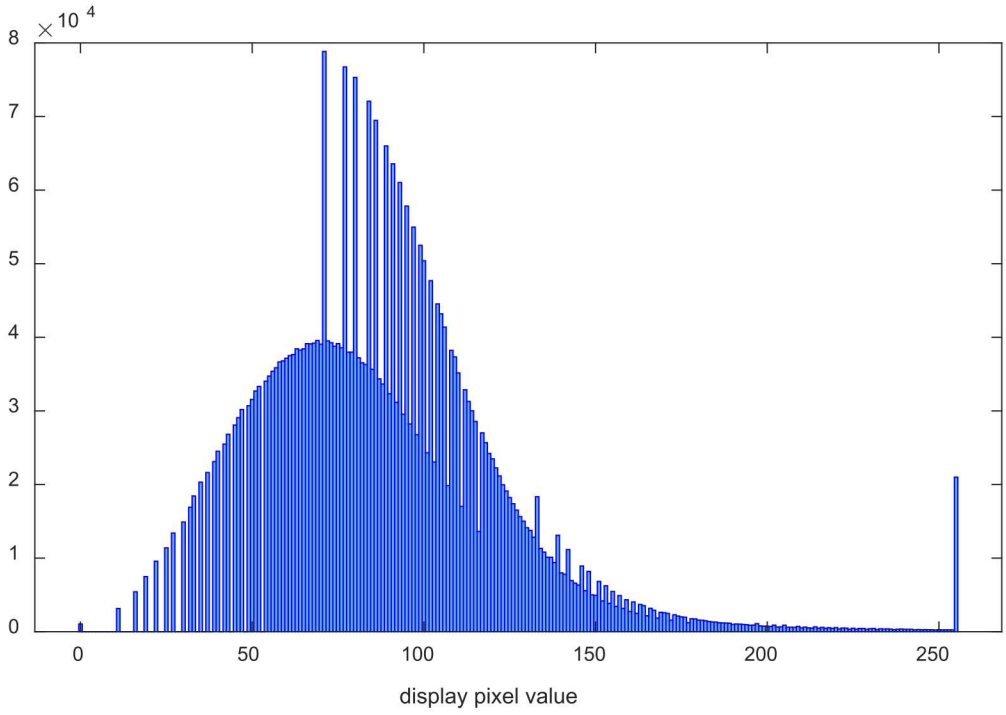

Figure 7. Histogram of display pixel values *P* with $\beta = \left[3 \ \text{median}\left(\sqrt{p} \right) \right]^{-1}$.

3.3 Logarithmic Image

Another function that compresses larger values more than smaller values is the logarithm function. Consequently, our mapping might be

$$P = 255\,\alpha \log_2 \left(p \right) \ , \tag{15}$$

where we might choose scale factor

$$\alpha = \frac{1}{16} \ . \tag{16}$$

As before, the output value P is quantized and limited to the range [0:255]. The resulting image is given in Figure 8, and its histogram is illustrated in Figure 9. While this particular scale factor results in no clipping of large input pixel values, we also note the higher density of mid-scale output values. This might make the image appear somewhat more washed out (smaller apparent display dynamic range).

3.4 Arctangent Image

Another function that compresses larger values more than smaller values is the arctangent function. Consequently, our mapping might be

$$P = 255(2/\pi)\,\mathrm{atan}\left(\frac{\eta\ p}{2^{16}} \right) \ , \tag{17}$$

where we might choose scale factor

$$\eta = 400 \, . \tag{18}$$

As before, the output value P is quantized and limited to the range [0:255]. The resulting image is given in Figure 10, and its histogram is illustrated in Figure 11. Note that this rendering is much like the quarter-power image except that highest discrete clutter values are not hard-limited at value 255. Rather, they are tapered towards 255.

Figure 8. Logarithmic image with $\alpha = 1/16$.

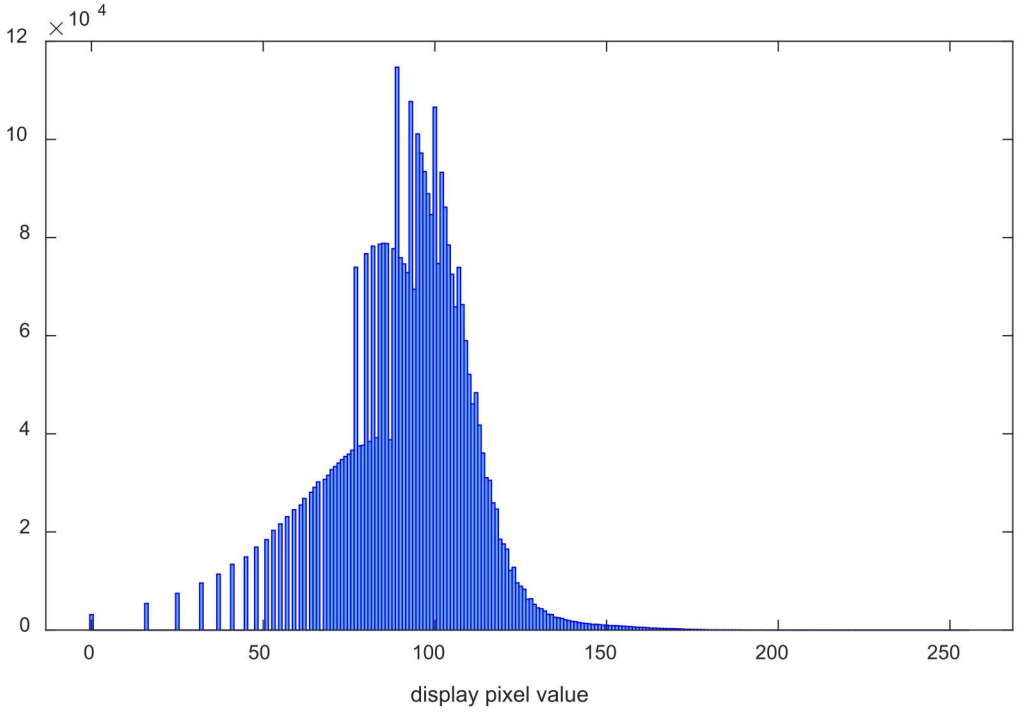

Figure 9. Histogram of display pixel values *P* with $\alpha = 1/16$.

Figure 10. Arctangent image with $\eta = 400$.

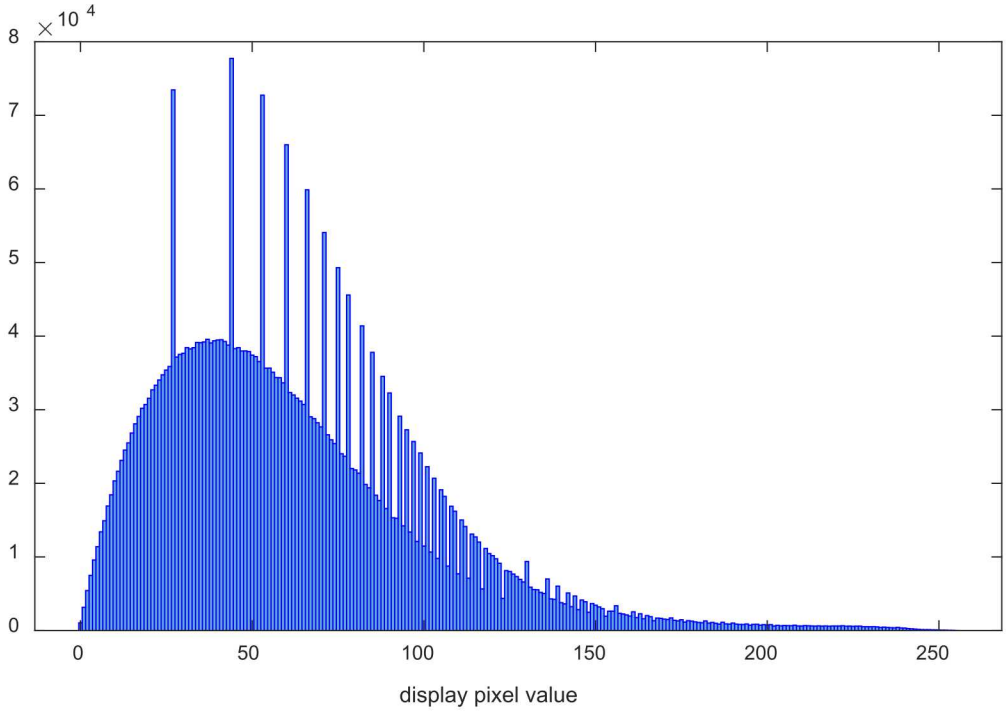

Figure 11. Histogram of display pixel values *P* with $\eta = 400$.

4 Image Display – Color

The utility of color in displaying conventional SAR images is subject of some debate. Clearly, a single SAR image results from a single band of microwave frequencies, so is essentially monochromatic in nature. The paradigm of color representing different passbands of spectral content, as in Electro-Optic (EO) imagery, or Hyper-Spectral (HS) imagery, does not apply to an individual SAR image. Furthermore, any use of color must come with an awareness of the limitations on color cognition in color-blind human observers. Nevertheless, in some communities, the idea of colors being mapped to intensity does find some utility.

With no intent on being rigorous or complete, we do offer here some examples of colorized SAR images, and comment accordingly.

4.1 Single-Color Image Display

For color vision, the human eye is most sensitive to green colors, nominally at a wavelength of 555 nm. Consequently, sometimes SAR images are rendered in a green intensity-based scale. Figure 12 shows the same quarter-power histogram mapping as in Section 3.2, except that the intensity is rendered in the color green instead of gray.

We leave to the reader any subjective assessment of the relative merits.

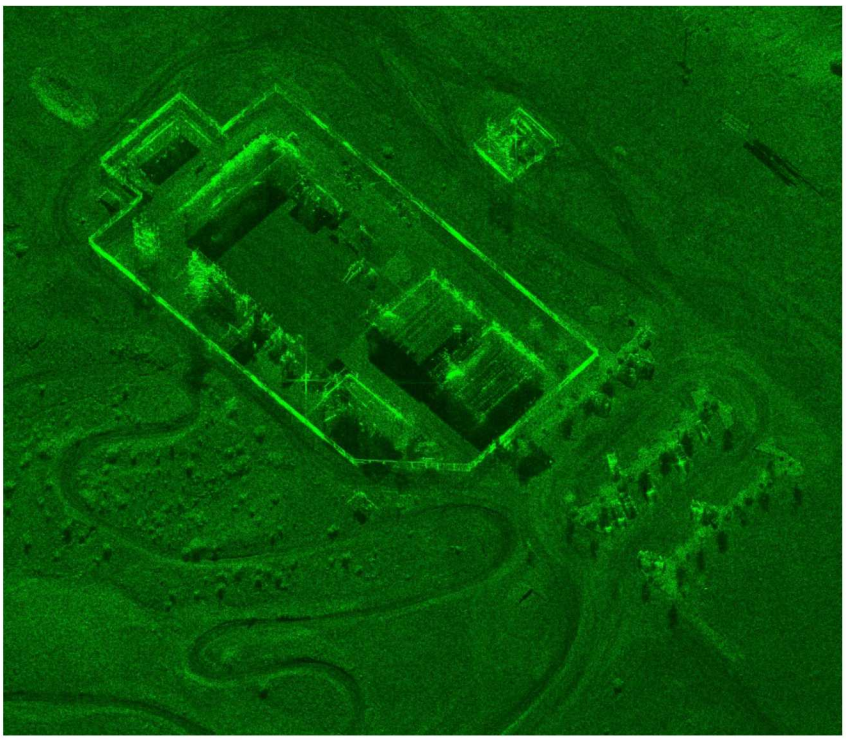

Figure 12. Quarter-power image rendered in green.

4.2 Multi-Color Image Display

In spite of our familiarity with color in our everyday lives, an objective study of color and color perception is incredibly complicated, and far too complex to treat in this report with any sense of completeness. Nevertheless, we offer some 'basics' and attempt to relate them to display of SAR images.

The HVS, color-blindness notwithstanding, is generally able to perceive a combination of optical wavelengths corresponding to colors red, green, and blue; about 7-10 million different colors.

Colors may be represented in several different manners. One popular representation is the chromaticity diagram shown in Figure 13. Note that "white" has coordinates (1/3, 1/3) in this representation. We further identify "hue" as an indication of the direction of a color from white, and "saturation" as the essentially the distance to the color from white. A third dimension to this representation, but not overtly displayed, is intensity. For example, shades of gray are just intensity variations of the color "white."

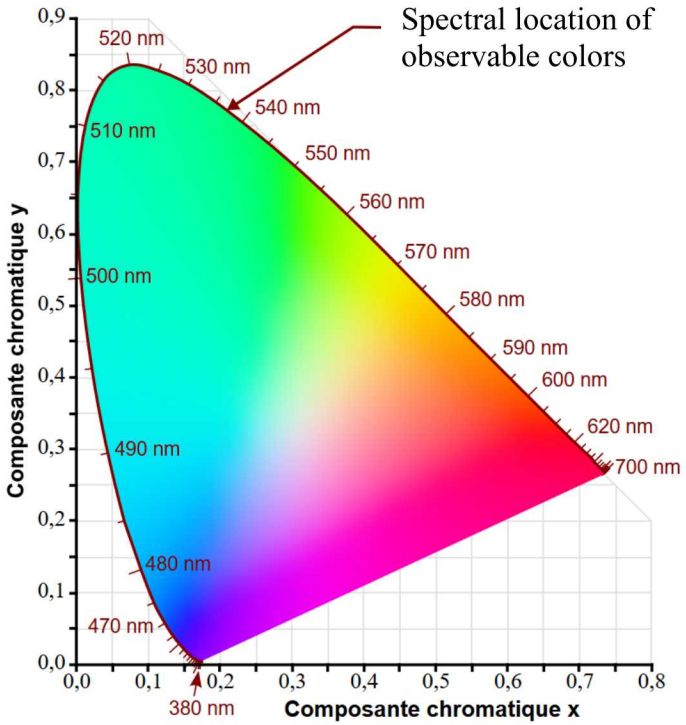

Figure 13. The CIE 1931 color space chromaticity diagram. (Courtesy Wikimedia Commons)

While Figure 13 illustrates the gamut of colors that can be perceived by the HVS, the gamut of colors that can be represented on a display device is a subset of these. For example, a display that can generate pure colors R, G, and B in Figure 14, can combine those colors to create any other color within the triangle defined by them. This is the gamut of the display. Colors outside of this triangle are "outside the gamut" of the display. This display gamut, along with intensity, defines a 3-dimensional color volume that can be represented by a display.

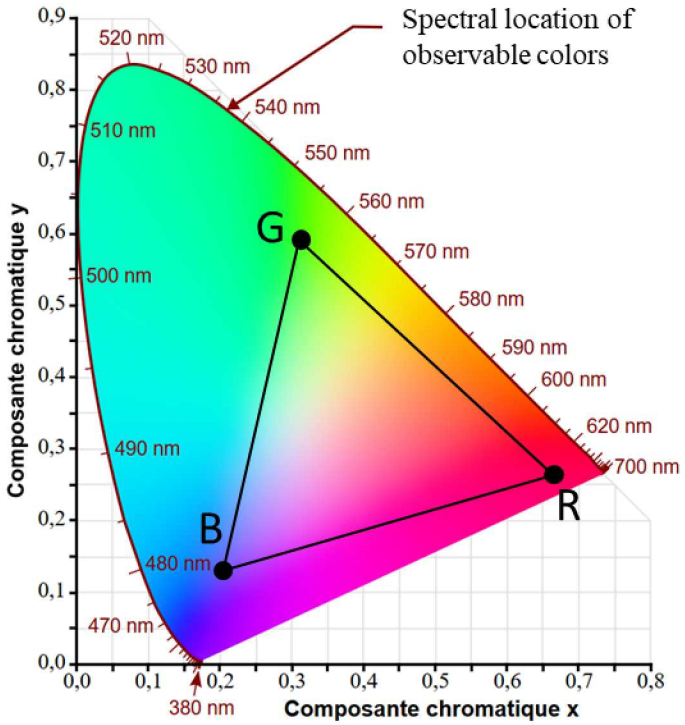

Figure 14. Chromaticity diagram of Figure 13 with notional display gamut superimposed.

The display gamut (including intensity) restricts the set of colors that can be used to render a SAR image. Different colors/intensities can be assigned to different pixel values. If smoothly varying colors/intensities are desired, then pixel values will be assigned to positions along a curve in the 3-dimensional color volume. In some sense, the longer the length of this curve, the more colors/intensities are used, and the greater is the dynamic range of the pixel data that can be discerned. In addition, subtle image features are more apparent with higher color gradients.

We define a specific relationship of colors/intensities to what would be otherwise gray-scale pixel values, as a "colormap." For a typical display monitor, the colormap is a function of how much red, green, and blue are associated with a pixel value. We opine that the utility of any particular colormap will depend on the application and/or intent of the observer. A particular colormap with utility to one application might offer no advantage to another colormap for a different application. Colormap selection has been explored by several researchers.[18,19,20]

We conjecture that useful colormaps for SAR images still begin with black and end with white, but now perhaps arc their way through the color/intensity volume. The choice of a particular colormap might be influenced by any number of factors, including the display gamut, photocopy reproduction limitations, and the gamut available to the observer (e.g. color-blindness).

We now present several examples of the SAR image of the previous examples, with a common quarter-power histogram mapping as discussed in Section 3.2, now with scale factor $\beta = \left[5 \text{ median}\left(\sqrt{p} \right) \right]^{-1}$, but now also employing various colormaps.

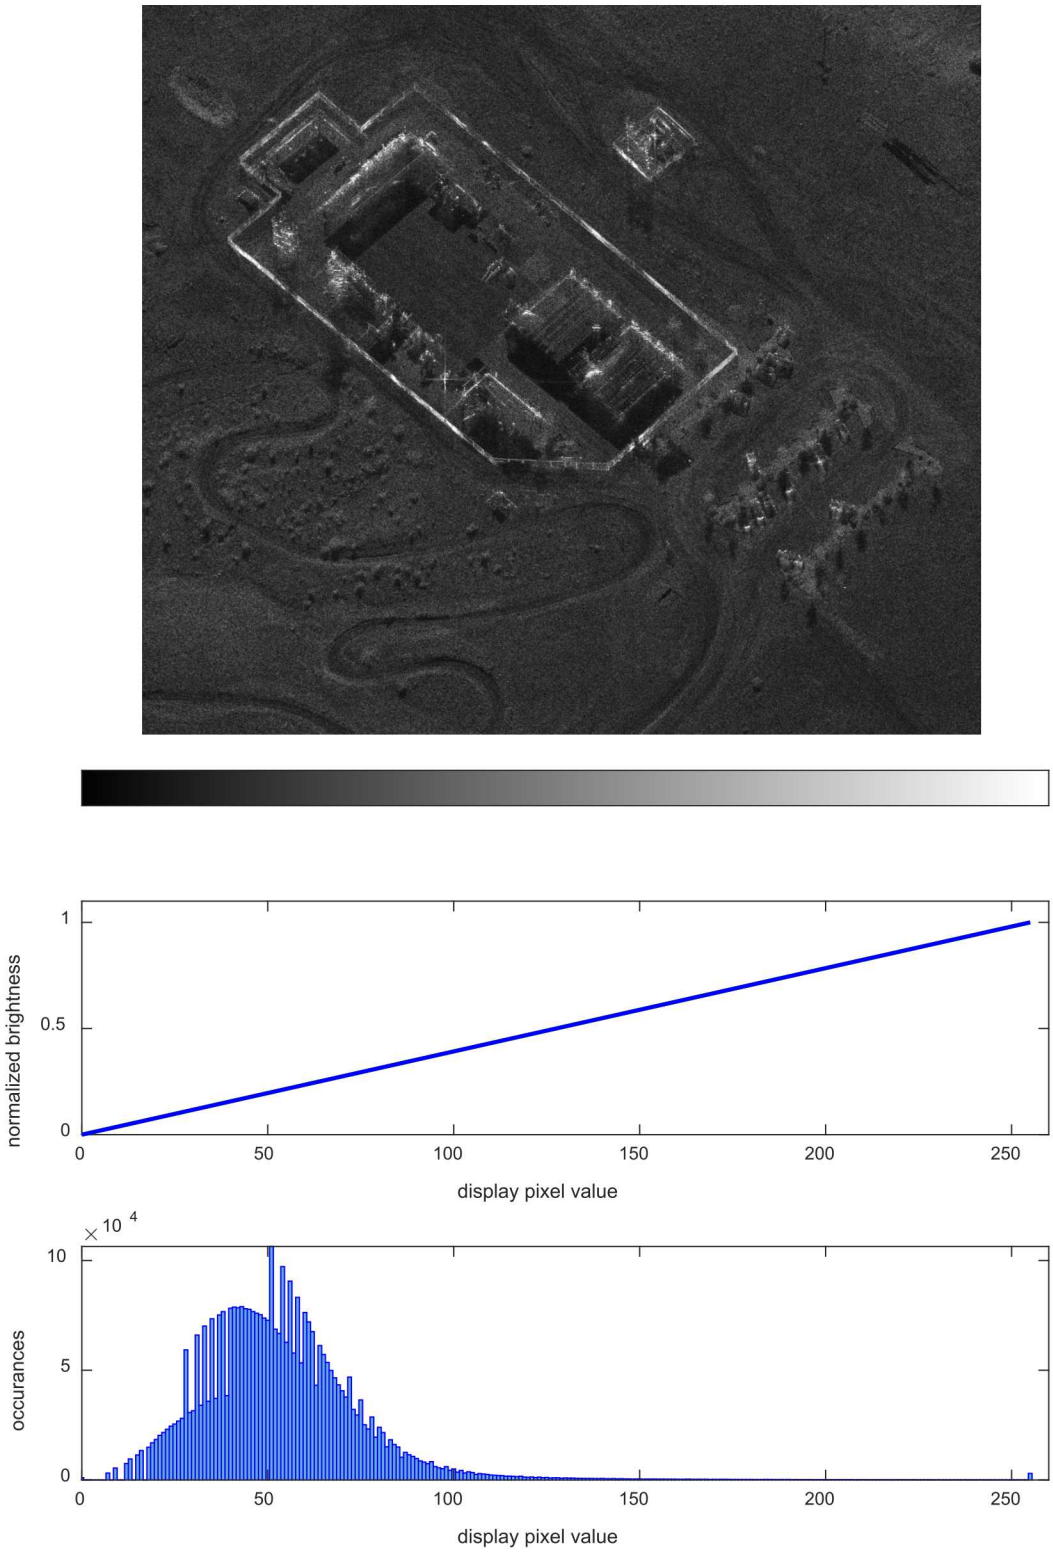

Figure 15. Reference gray-scale image. The underlying distribution of display pixel values is a quarter-power image with $\beta = \left[5 \text{ median}\left(\sqrt{p}\right)\right]^{-1}$.

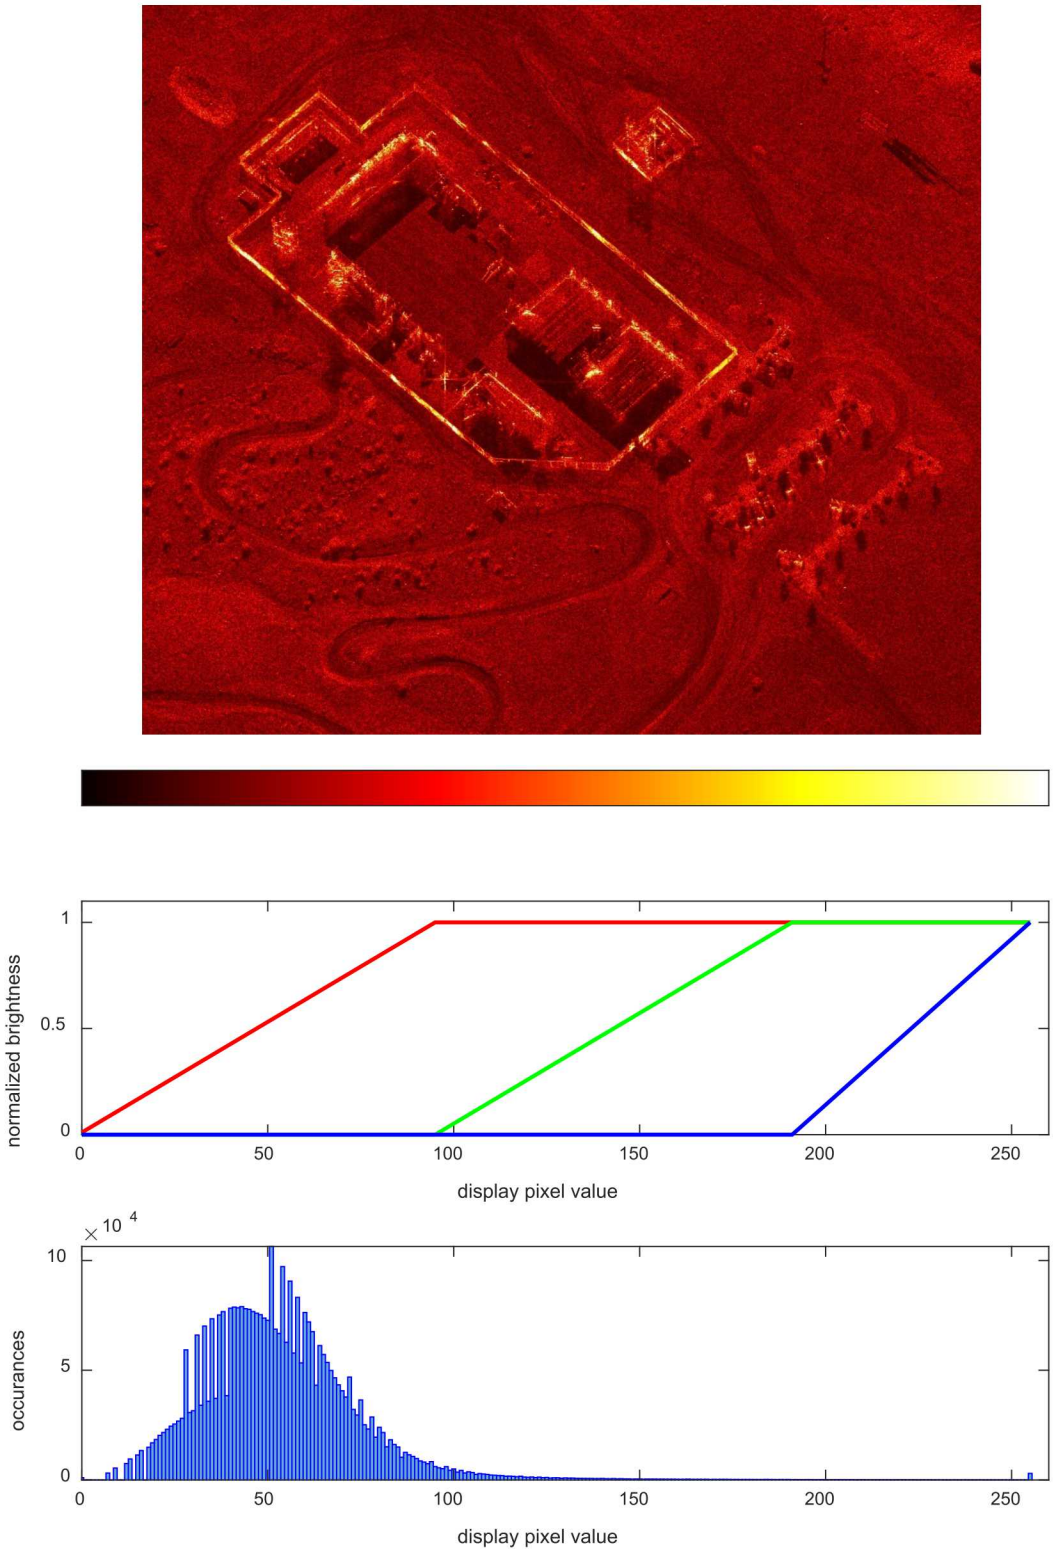

Figure 16. Matlab™ colormap "hot." This colormap is a standard "off-the-shelf" colormap from the Matlab™ computational language. It increases the gradient of the red color for lower pixel values, thereby revealing detail in the natural clutter.

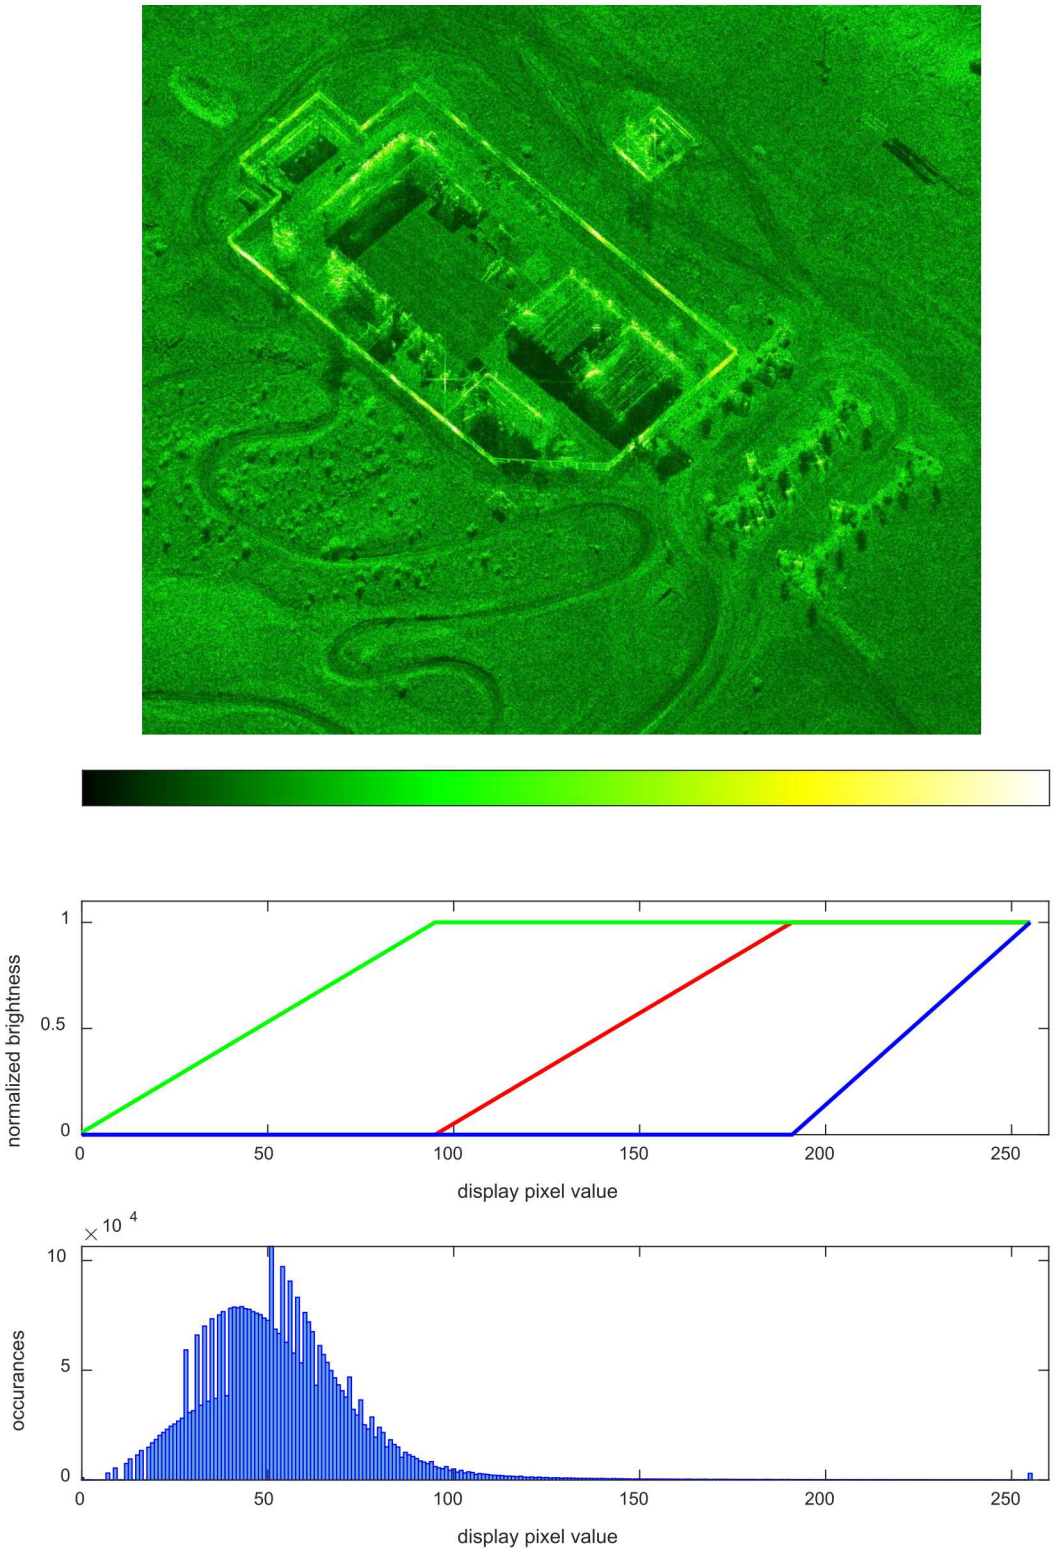

Figure 17. This colormap begins with the Matlab™ colormap "hot" and trades the red and green intensities.

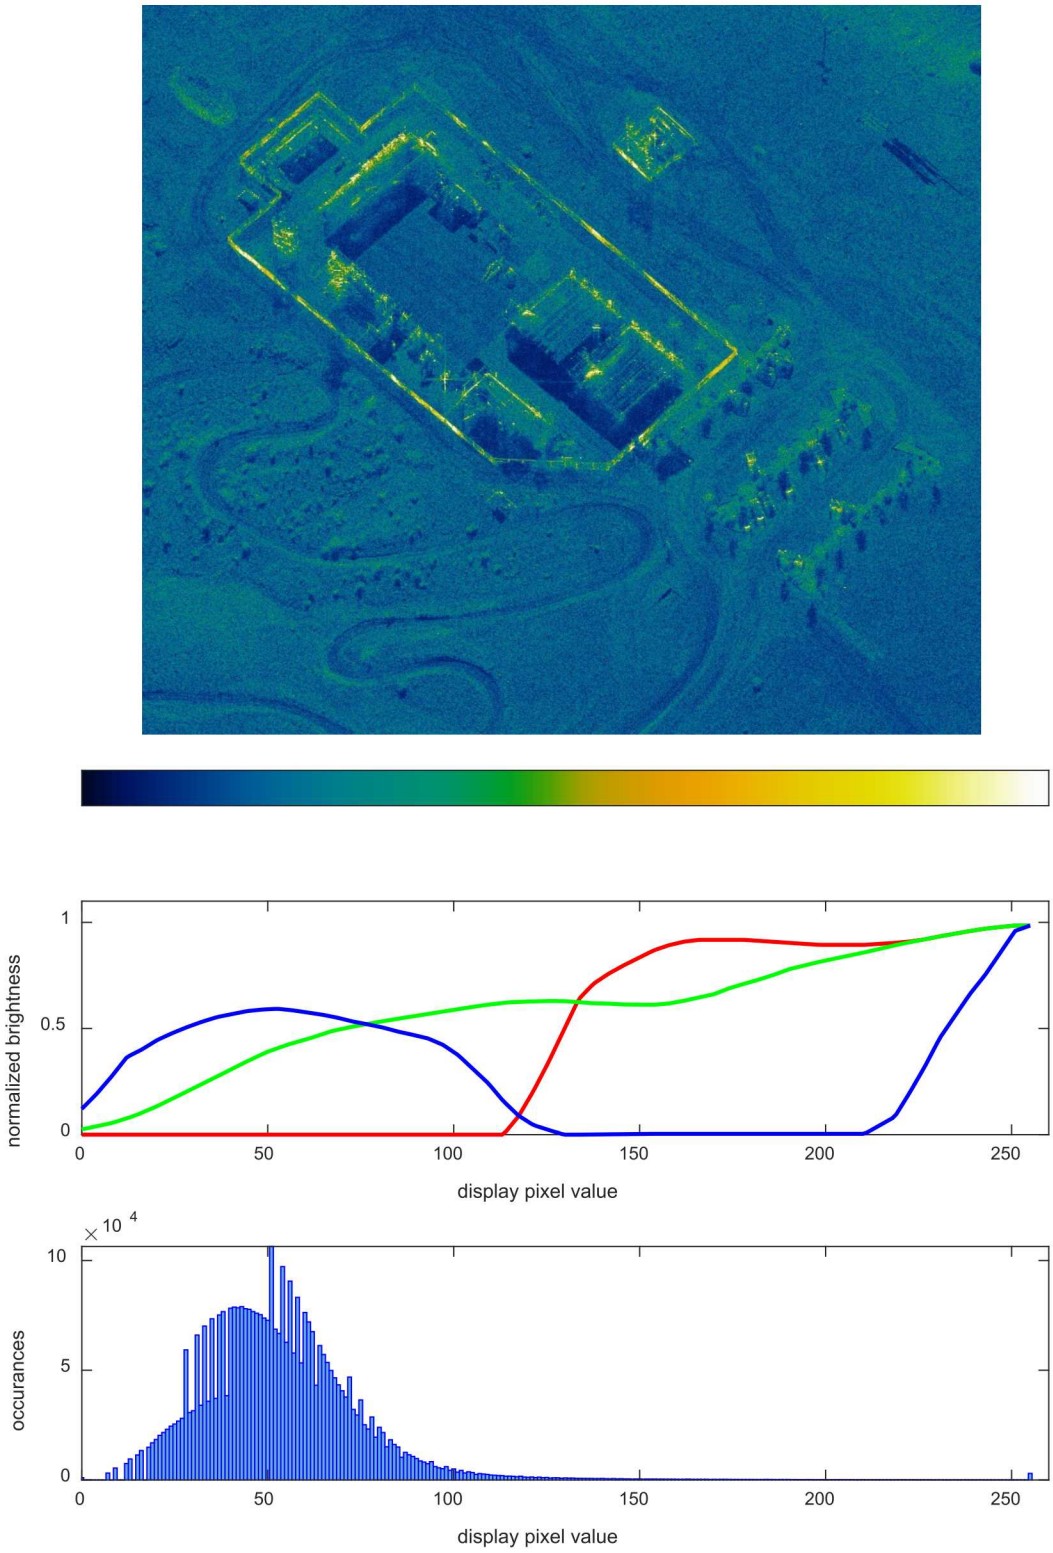

Figure 18. This colormap was created by W. H. Hensley and is loosely based on the "Scientific American" colormap from the Khoros™ integrated software development tool. This map bears some similarity to the Matlab™ "parula" colormap, although significant differences still exist between them.

5 Miscellaneous Topics

We now address some miscellaneous topics.

5.1 Mosaicked Images (including Stripmaps)

In earlier sections we investigated a single spotlight SAR image, created from a single synthetic aperture. Oftentimes, multiple SAR images are mosaicked together to form larger composite images. This includes SAR stripmap images composed of mosaicked patches, and ScanSAR images composed of mosaicked stripmap SAR images.

To avoid obvious seams in the mosaicked result, particularly in regions of distributed natural clutter, individual SAR image patches need to exhibit the same relationship between pixel reflectivity and displayed color/intensity. If a displayed rendering is data dependent, then it needs to adapt to perhaps one representative image, perhaps the first in a sequence, and then the mapping function needs to be held constant for all subsequent images of the larger mosaic.

5.2 Image Phase

Although SAR images are naturally complex-valued, typically only the magnitude information is displayed (after suitable mapping). One might ask "What about the phase?"

Pixel phase for a single SAR image is typically uniformly distributed over all angles. For example, this histogram for the SAR image of Figure 1 is displayed in Figure 19. Note the essentially uniform distribution of phase values.

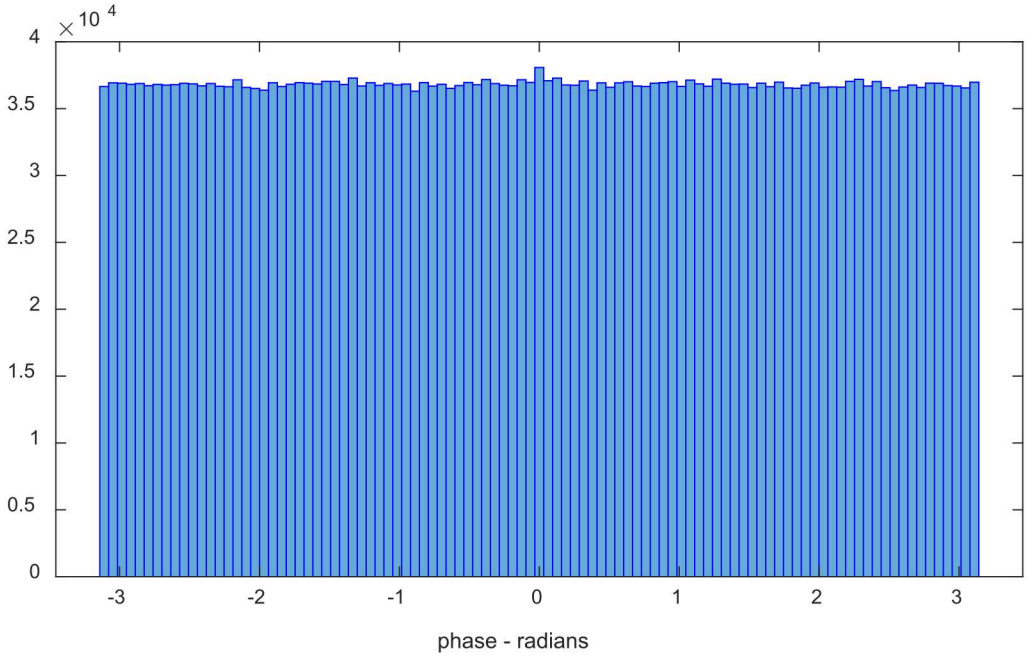

Figure 19. Histogram of pixel phase for image of Figure 1.

"The aim of art is to represent not the outward appearance of things,
but their inward significance."
-- Aristotle

6 Conclusions

We repeat some key points.

- The SAR "image" file format can be a dynamic range choke-point. Proper scaling of the image data is required to maximize the dynamic range that can be represented, especially for integer-based pixel values.

- Most SAR images have a magnitude histogram that is approximately Rayleigh distributed, albeit with stronger "tails" to accommodate clutter discretes.

- SAR images often have dynamic ranges in the many of tens of dB, and may even manifest well over 100 dB.

- The Human Visual System is typically limited to discerning 30-40 dB of intensity dynamic range in a single image. Consequently, SAR image data will typically need some sort of dynamic range compression to be exploited by a human observer, especially for gray-scale and monochromatic image display. Several functions are identified to help achieve this.

- Colors can be employed to help display an extended dynamic range of the underlying image. Higher color gradients can be used to enhance the observability of subtle clutter features in the SAR image.

"There are things known and there are things unknown,
and in between are the doors of perception."
-- Aldous Huxley

Appendix A – The Rayleigh Distribution

Here we examine characteristics of the Rayleigh distribution.

Consider a complex random variable, composed of independent real and imaginary components, each an independent zero-mean Gaussian (Normal) distribution, albeit with identical variance. That is, let

$$Z = X + jY\,, \tag{A1}$$

where X and Y distributions are each described with Probability Density Function (PDF)

$$f_X(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{x^2}{2\sigma^2}}\,, \text{ and}$$

$$f_Y(y) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{y^2}{2\sigma^2}}\,, \tag{A2}$$

where

$$\sigma^2 \text{ = variance common to components } X \text{ and } Y. \tag{A3}$$

The magnitude of Z is calculated as

$$R = |Z| = \sqrt{X^2 + Y^2}\,, \tag{A4}$$

and is Rayleigh distributed with PDF

$$f_R(r) = \frac{r}{\sigma^2} e^{-\frac{r^2}{2\sigma^2}}\,. \tag{A5}$$

The Cumulative Density Function (CDF) is identified as

$$F_R(r) = 1 - e^{-\frac{r^2}{2\sigma^2}}\,. \tag{A6}$$

The statistics of the Rayleigh distribution are as follows.

Table 2. Statistics of Rayleigh distribution.

Mean	$\sigma\sqrt{\frac{\pi}{2}}$
Median	$\sigma\sqrt{2\ln(2)}$
Mode	σ
Variance	$\frac{4-\pi}{2}\sigma^2$
RMS value	$\sqrt{2}\,\sigma$
90% limit	2.146σ
99% limit	3.035σ
99.9% limit	3.717σ

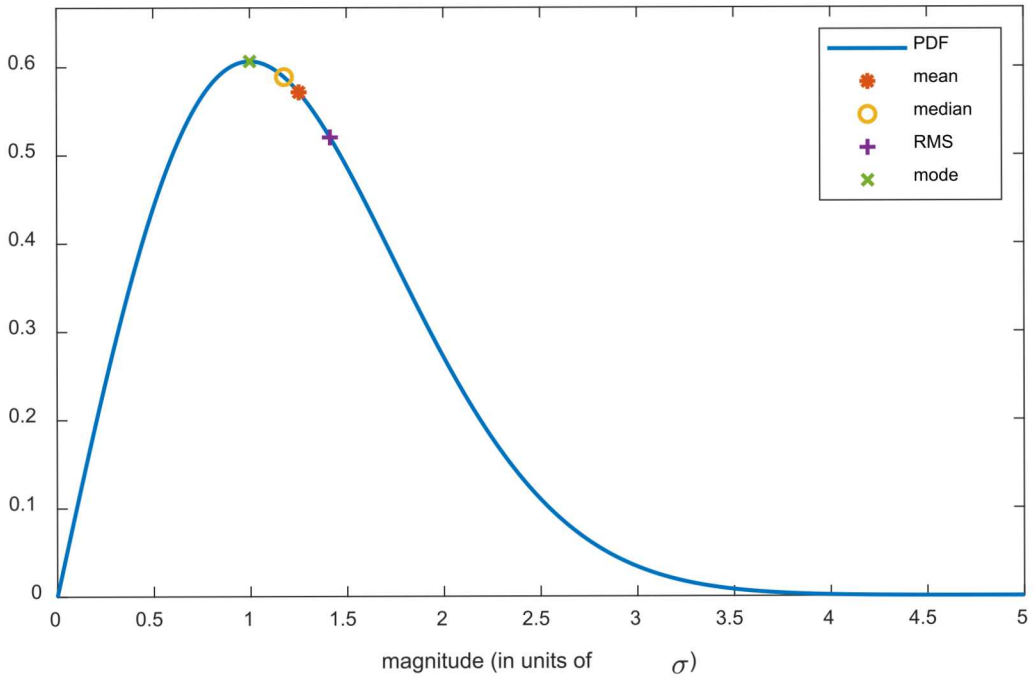

Figure 20. Rayleigh PDF, with statistics identified.

Appendix B – Comments on SICD Image Format

The National Geospatial-Intelligence Agency (NGA) has created and adopted a standard format for SAR images known as the Sensor Independent Complex Data (SICD) format. We briefly examine here some aspects of the SICD format applicable to the topic of this report. Details can be found in the SICD documentation by NGA.[21]

Pixel Formats

The current SICD format allows for three distinct pixel formats, given in Table 3.

Table 3. SICD Pixel Types.

Pixel Type	Format description
RE32F_IM32F	Each pixel is stored as a pair of numbers that represent the real and imaginary components. Each component is stored in a 32-bit IEEE binary32 floating point format (4 bytes per component, 8 bytes per pixel).
RE16I_IM16I	Each pixel is stored as a pair of numbers that represent the real and imaginary components. Each component is stored in a 16-bit signed integer in 2's complement format (2 bytes per component, 4 bytes per pixel).
AMP8I_PHS8I	Each pixel is stored as a pair of numbers that represent the amplitude and phase components. Each component is stored in an 8-bit unsigned integer (1 byte per component, 2 bytes per pixel).

The pixel type RE32F_IM32F, by virtue of it being floating point numbers, offers the largest dynamic range.

Note that the pixel type RE16I_IM16I uses signed integers, meaning that the dynamic range is limited to 15 bits, instead of 16 bits, accounting for the sign-bit.

In addition, for a pixel type AMP8I_PHS8I, the stored pixel value might optionally be used as an index into a lookup table to determine the actual pixel magnitude value. If used, this amplitude lookup table is stored in an SICD parameter identified as AmpTable. This allows for companding of magnitude data to enhance dynamic range of the stored data.[3,4]

Radiometric Calibration

The current SICD format allows several optional mechanisms to describe the correspondence of pixel values to radiometric parameters. The relationship is generally expressed as a 2-dimensional polynomial dependent on image row and column indices. i.e. allowing for spatial variance. The SICD parameters are described in Table 4.

Table 4. SICD Radiometric Correction Parameters.

Parameter	Description
RCSSFPoly	Polynomial coefficients that yield a scale factor to convert pixel power to RCS (m^2) as a function of image row coordinate (variable 1) and column coordinate (variable 2).
SigmaZeroSFPoly	Polynomial coefficients that yield a scale factor to convert pixel power to clutter parameter Sigma-Zero (σ_0) as a function of image row coordinate (variable 1) and column coordinate (variable 2).
BetaZeroSFPoly	Polynomial coefficients that yield a scale factor to convert pixel power to radar brightness or Beta-Zero (β_0) as a function of image row coordinate (variable 1) and column coordinate (variable 2).
GammaZeroSFPoly	Polynomial coefficients that yield a scale factor to convert pixel power to clutter parameter Gamma-Zero (γ_0) as a function of image row coordinate (variable 1) and column coordinate (variable 2).

Each of these describes the nature of a 2-dimensional polynomial, where coefficients are specified for specific exponent combinations, with the overall orders of the polynomial specified as attributes. The spatially-variant relationship allows for images that are not corrected for variations in antenna beam patterns and range loss. These various parameters are not entirely independent of each other; related via imaging geometry and resolution, etc.

Noise Floor

The current SICD format requires an indication of the noise level in the pixels of the image, expressed as a 2-dimensional polynomial dependent on image row and column indices. i.e. allowing for spatial variance. The SICD parameters describing noise are given in Table 5.

Table 5. SICD Noise Description Parameters.

Parameter	Description
NoiseLevelType	Parameter to indicate that the noise power polynomial yields either absolute power level or power level relative to the Scene Center Point (SCP) pixel location.
NoisePoly	Polynomial coefficients that yield thermal noise power (in dB) in a pixel as a function of image row coordinate (variable 1) and column coordinate (variable 2).

NoisePoly describes the nature of a 2-dimensional polynomial, where coefficients are specified for specific exponent combinations. The spatially variant relationship allows for images that might have been corrected for variations in antenna beam patterns and range loss.

“Change the way you look at things and the things you look at change.”
— Wayne W. Dyer

References

[1] Ana Radonjić, Sarah R. Allred, Alan L. Gilchrist, David H. Brainard, "The Dynamic Range of Human Lightness Perception," Current Biology, Vol. 21, pp. 1931–1936, November 22, 2011.

[2] Timo Kunkel, Erik Reinhard, "A Reassessment of the Simultaneous Dynamic Range of the Human Visual System," Proceedings of the 7th Symposium on Applied Perception in Graphics and Visualization, APGV '10, pp. 17–24, Los Angeles, California, 23-24 July 2010.

[3] Armin W. Doerry, *SAR Image Complex Pixel Representations*, Sandia National Laboratories Report SAND2015-2309, Unlimited Release, March 2015.

[4] A. W. Doerry, "Representing SAR complex image pixels," SPIE 2016 Defense & Security Symposium, Radar Sensor Technology XX, Vol. 9829, Baltimore, MD, 17-21 April 2016.

[5] Armin W. Doerry, *Performance Limits for Synthetic Aperture Radar – second edition*, Sandia National Laboratories Report SAND2006-0821, Unlimited Release, February 2006.

[6] Armin W. Doerry, *Collecting and Processing Data for High Quality CCD Images*, Sandia National Laboratories Report SAND2007-1545, Unlimited Release, March 2007.

[7] G. L. Hover, G. A. Mastin, R. M. Axline, J. D. Bradley, *Evaluation of Synthetic Aperture Radar for Oil-Spill Response*, U.S. Coast Guard Final Report CG-D-02-94, Unclassified, October 1993.

[8] Maurice W. Long, *Radar Reflectivity of Land and Sea*, 2nd edition, ISBN 0-89006-130-0, Artech House Inc., 1983.

[9] Fawwaz T. Ulaby, M. Craig Dobson, *Handbook of Radar Scattering Statistics for Terrain*, ISBN 0-89006-336-2, Artech House Inc., 1989.

[10] Ann Marie Raynal, Douglas L. Bickel, Dale F. Dubbert, Tobias J. Verge, Bryan L. Burns, Ralf Dunkel, Armin W. Doerry, "Radar cross section statistics of cultural clutter at Ku-band," SPIE 2012 Defense, Security & Sensing Symposium, Radar Sensor Technology XVI, Vol. 8361, Baltimore MD, 23-27 April 2012.

[11] Armin W. Doerry, *Reflectors for SAR Performance Testing – second edition*, Sandia National Laboratories Report SAND2014-0882, Unlimited Release, Supersedes SAND2008-0396, February 2014.

[12] Rafael C. Gonzalez, Paul Wintz, *Digital Image Processing – second edition*, ISBN 0-201-11026-1, Addison-Wesley Publishing Company, Inc., 1987.

[13] A. W. Doerry, "Radiometric calibration of range-Doppler radar data," SPIE 2019 Defense & Commercial Sensing Symposium, Radar Sensor Technology XXIII, Vol. 11003, Baltimore, MD, 14-18 April 2019.

[14] *IEEE Standard for Floating-Point Arithmetic*, IEEE Std 754™-2008, IEEE Computer Society, 29 August 2008.

[15] Fred E. Nathanson, *RADAR Design Principles – second edition*, ISBN 0-07-046052-3, McGraw-Hill, Inc., 1991.

[16] K. Krishen, "Detection of oil spills using a 13.3-GHz radar scatterometer," *Journal of Geophysical Research*, Vol. 78, No. 12, pp. 1952–1963, 20 April 1973.

[17] J. W. Johnson, W. F. Croswell, "Characteristics of 13.9 GHz radar scattering from oil films on the sea surface," *Radio Science*, Vol. 17, No. 3, pp. 611–617, May-June 1982.

[18] Bernice E. Rogowitz, Lloyd A. Treinish, "Data visualization: the end of the rainbow," *IEEE Spectrum*, December 1998.

[19] Alan D. Kalvin, Bernice E. Rogowitz, Adar Pelah, Aron Cohen, "Building perceptual color maps for visualizing interval data," SPIE 2000 Electronic Imaging, Human Vision and Electronic Imaging V, Vol. 3959, San Jose, CA, USA, 2 June 2000.

[20] Matthias Geissbuehler, Theo Lasser, "How to display data by color schemes compatible with red-green color perception deficiencies," Optics Express, Vol. 21, Issue 8, pp. 9862-9874, 2013.

[21] "SENSOR INDEPENDENT COMPLEX DATA (SICD)," Volume 1, Design & Implementation Description Document, NGA STANDARDIZATION DOCUMENT, NGA.STND.0024-1_1.1, Version 1.1, 30 September 2014.

Distribution

Unlimited Release

Hardcopy Internal

1	A. W. Doerry	5349	MS 0519
1	L. Klein	5349	MS 0519
1	M. R. Lewis	5349	MS 0519
1	S. P. Castillo	5340	MS 0532

Email—External

Brandeis Marquette	Brandeis.Marquette@ga-asi.com	General Atomics ASI
Jean Valentine	Jean.Valentine@ga-asi.com	General Atomics ASI
John Fanelle	John.Fanelle@ga-asi.com	General Atomics ASI

Email—Internal

Technical Library	9536	libref@sandia.gov

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.