Measurement of Noise Equivalent Quanta (NEQ) Using the Dead Leaves Technique

Robin Jenkin, NVIDIA, Santa Clara, CA, USA.
Uwe Artmann, Image Engineering GmbH & Co. KG, Kerpen, Germany

Abstract

Noise Equivalent Quanta (NEQ) is an objective Fourier metric which evaluates the performance of an imaging system by detailing the effective equivalent quanta of an exposure versus spatial frequency [1]. Calculated via the modulation transfer function (MTF) and noise power spectrum (NPS), it is a valuable precursor for ranking the detection capabilities of systems and a fundamental metric that combines sharpness and noise performance of an imaging system into a single curve in a physically meaningful way.

The dead leaves measurement technique is able to provide an estimate of the MTF and NPS of an imaging system using a single target [2-7], and therefore a potentially convenient method for the assessment of NEQ. This work validates the use of the dead leaves technique to measure NEQ, firstly through simulation of an imaging system with known MTF and NPS, then via measurement of camera systems, both in the RAW domain and post-ISP. The dead leaves approach is shown to be a highly effective and practical method to estimate NEQ, ranking imaging systems performance both pre- and post-ISP.

Introduction

Many of the processes that occur during the capture of an image may be categorized as distorting, adding noise or blurring content. And while an extraordinary amount of effort is expended on developing objective metrics that measure individual attributes, there still exists difficulties when trying to combine metrics in a meaningful way to yield a single figure to represent the performance of a system.

Many image capture and processing challenges may be reduced to a trade between noise and resolution. Integrating or processing signal over a larger area often reduces noise at the expense of sharpness and vice versa. Binning pixels and noise reduction filters are two trivial examples. Measurements of resolution via the Modulation Transfer Function (MTF), and noise via the Signal-to-Noise Ratio (SNR) or Noise Power Spectrum (NPS) therefore remain important but are often conducted and reported in isolation.

MTF generally does not change with signal level unless some non-linear process is involved and provides no information concerning noise. Noise in the capture process is generally shaped by processes subsequent to its creation and therefore does not account for signal fidelity. Thus, both are needed for a full description of the system. Further, SNR gives an aggregate number for noise and no information concerning frequency content and periodicity which may prove to be more bothersome than evenly distributed 'white noise'.

Noise Equivalent Quanta (NEQ) is a frequency-space metric which combines MTF and NPS to yield the SNR^2 of the image per reciprocal unit spatial frequency squared [1]:

$$NEQ(\omega) = \frac{MTF(\omega)^2}{NPS(\omega)/\mu^2} \tag{1}$$

where u is the mean signal value and ω spatial frequency. Keelan [1], Dainty and Shaw [8] and Barratt and Myers [9] provide good overviews of NEQ. As Keelan remarks, if the spatial frequency unit is chosen to be cycles per pixel, the units of the NEQ become SNR^2 per pixel, or the variance per pixel. For a Poisson distribution with mean q, the variance is also q. Therefore, as photon arrival is discrete and governed by Poisson statistics, NEQ represents the number of equivalent quanta per pixel versus spatial frequency for an idealized detector for a given exposure. NEQ provides a method to combine measures of signal and noise transport that may be interpreted in a useful engineering manner.

The NEQ at the DC frequency can usually not be computed as the NPS is the power of the signal fluctuations and therefore this component is generally removed. If the NEQ curve is projected back to estimate where it would intersect with the DC axis, the effective noise equivalent full well of the sensor can be estimated. The signal level, u, in digital counts can be used to scale the NEQ at this value to estimate the maximum noise equivalent quanta that would be recorded for the maximum bit depth of the sensor. If the bit depth of the pre-HDR combination exposures is known, this will relate more closely to an estimate of the physical actual full well of the pixel.

It should be noted that the noise profile of a modern high dynamic range (HDR) sensor is not monotonic and contains SNR 'holes' [10]. Because of this, if the NPS were measured using a uniform target, the result would not represent the aggregate behavior of the sensor across the dynamic range as the noise power becomes disproportionately larger at high signal levels when compared to conventional single exposure sensors. Further, if the exposure level were close to the noise floor of the sensor or coincident with an SNR hole, it would further deviate from the aggregate behavior. This provides weight to the need to evaluate NPS across a range of exposure levels, such as is afforded by the dead leaves method even when working in the linear signal domain.

For linear processes or filters applied to the captured image, the NPS is further shaped by the square of the MTF of the filter or process [1]. As NEQ is squared MTF divided by NPS, this means that the NEQ should not be affected by linear processes applied to the image. As Keelan writes, this is intuitively satisfying as, because the process is reversable, there is no change in the fundamental utility of the image [1]. While this has been documented, to the authors knowledge, it has not been demonstrated prior to this publication, see Simulation.

Detection theory describes the statistics of signal separation. If $\Delta f(\omega)$ represents the delta between signal spectra to be determined, the number of mutual standard deviations between the spectra, or d', the idealized observer may be written [11,12]:

$$d'^2 = \int |\Delta f(\omega)|^2 NEQ(\omega) d\omega \tag{2}$$

As may be seen, if the NEQ for an imaging system and spatial power distributions of signals to be found are known, the idealized

detection performance of a system can be determined for particular objects.

Measurement of NEQ Using Dead Leaves

The measurement of NEQ requires the measurement of Modulation Transfer Function (MTF) and the Noise Power Spectrum (NPS). From hereon we use the term MTF and SFR interchangeably.

The dead leaves pattern is a well-established method to measure SFR and derive information concerning texture loss, or low contrast fine details in images, due to noise reduction, compression or other processes. First introduced by Cao *et. al.* [6] and later modified by McElvain *et. al.* [7], the most recent algorithm by Kirk *et. al.* [4] is used here and described in technical specification ISO19567-2 [5].

The approach utilizes cross correlation between the input signal and an output image. The input is a reference image based on known specifications of the dead leaves target and the output is the image obtained by the camera under test.

The complex transfer function, H(f), Figure 1, of the camera system is calculated using Equation 3 [3]:

$$H(f) = \frac{\phi_{YX}(f)}{\phi_{XX}(f)} \qquad (3)$$

where $\phi_{XX}(f)$ is the complex conjugate of the input and $\phi_{YX}(f)$ the cross-power spectrum of input and output. Once *H(f)* has been obtained, the SFR is derived by a rotational averaging process based on the real component of H(f).

The rotational average is computed by aggregating all results with the same spatial frequency, resulting in a conversion of the 2D function into a 1D SFR. Further details to obtain the SFR based on the dead leaves pattern can be found in references [2-5].
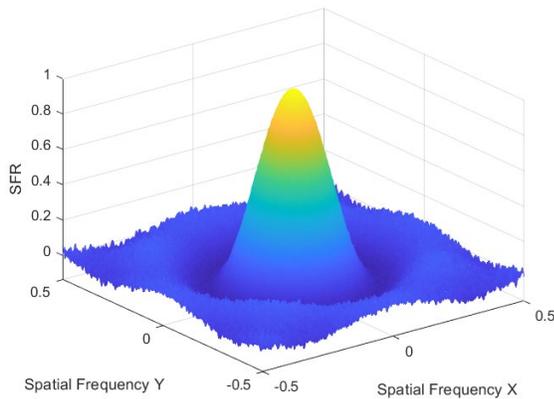
Figure 1 2D SFR obtained from Dead Leaves pattern using the described method.

Evaluation of the NPS from a dead leaves target is less common than that of the SFR, but has been documented previously by Artmann [2]. We assume that image capture consists of four main modifications to the input target:

1. Tone Reproduction
2. Geometric distortion
3. Optical blur
4. Addition of Noise

The effects of tone reproduction may be minimized by utilization of a linearization process based on gray patches surrounding the test pattern. Localized spatial matching may be used to reduce the effects of geometrical distortion and map the output image onto the same area of the input target prior to performing the cross-correlation. Both techniques are used regularly in the calculation of SFR using the dead leaves method [2-7].

Given the above we may now assume that the output image is the input target with optical blur and noise added. As H(f) is complex transfer function of the imaging system, we may multiply it with the Fourier transform of the target to obtain an estimation of the blurred target without noise. Subtracting this result from the noisy and blurred output image results in an estimation of the noise image alone. The noise image is then used to estimate the NPS and subsequently the NEQ according to Equation 1.

Calculation of the SFR and NPS from the same complex input pattern has numerous advantages including reduced capture workload, the ability to assess the impact of non-linear processing and HDR combination techniques and further non-monotonic noise behavior as mentioned in the introduction.

Simulation

Simulation of the technique was conducted to assess viability against known system parameters according to the flow detailed in Figure 2. A dead leaves target was convolved with an FIR filter to simulate optical blur. The FIR filter was generated to mimic a target MTF curve produced by a specified f-number lens (f1.8) at a given wavelength (550nm) in combination with a given pixel size (2.1um) [14]. Poisson noise was then added to the spatially degraded target to simulate exposure such that 100% reflectance had a known number of quanta (200) [15].

This degraded blurred and noisy target was further filtered by a Gaussian spatial filter ($\sigma = 0.66$) to test the hypothesis that the NEQ is not affected by linear processes. In additional the degraded target was filtered by a non-linear median filter (3 x 3) to validate that the NEQ was changed by non-linear processes. No tone or spatial distortion was applied to the images. Figure 3 shows sections of the degraded target, the linearly filtered degraded target and the non-linearly filtered degraded target.
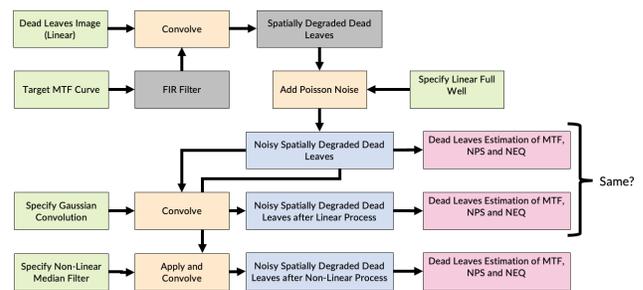
Figure 2 Experimental Simulation flow

Using visual inspection, the Gaussian and Median filtered images clearly contain less noise, though are somewhat degraded in sharpness as would be expected when compared with the original. Depending on the scale of reproduction however, these images may be chosen over the original due to the lack of noise. It should be noted that the patches shown represent about 1/64th the area of the total image. The PSNR of the simulated targets were also evaluated.

The degraded target had a PSNR of 25.7 dB and those of the linearly and non-linearly filtered targets, 28.6 and 27.3 dB respectively.

(a)   (b)   (c)
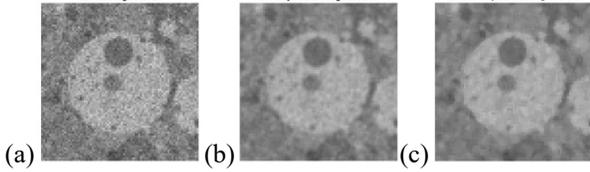
Figure 3 The degraded noisy target, (a), with Gaussian filtering applied, (b), and a non-linear median filter applied (c). The PSNR of the images is 25.7, 28.6 and 27.3 dB respectively.

The resultant SFRs measured for each target are shown in Figure 4. Results are from an average of 128 iterations with 64 rings used in the rotational average for the output. The SFR for the original degraded target matches that expected with a slight uplift in higher spatial frequencies due to noise. This is as expected as documented by Dainty and Shaw [8] and Yeadon *et. al.* [13]. The SFRs after filtering show additional reduction in the response as expected.

Figure 5 shows the NPS measured. The degraded target exhibits white noise as added and the expected and measured degraded NPS levels coincide. After linear filtering, the NPS takes on an exaggerated Gaussian shape as expected as the MTF of the filter was Gaussian and has shaped the noise according to the square of the MTF of the filter. After non-linear filtering however, it may be seen that the NPS is not shaped by the square of the MTF of the filter. While high spatial frequency noise power has been reduced, it also appears that low frequency power has been added. This is due to outlier pixels in the filter region 'snapping' to the median response within the filter window, namely the low-frequency value, as designed.

Examining the NEQ, Figure 6, the degraded target measurement is as expected. The slight uplift from the SFR result is transferred into the NEQ. Because of the mathematical combination of the curves, the noise in both the SFR and NPS measurements has resulted slightly increased noise in the NEQ measurement. The mean value of the target was 0.5 and exposed such that 100% was 200 quanta. It may be seen that the NEQ curve 'points' to a value of 100 at the DC Frequency and would lead to an accurate estimate of the linear full well of 200 as a linear process was modelled. After applying the linear filter, it may be seen that the NEQ is unchanged. Despite the reduction in the SFR of the system, the NPS has also reduced by the square of the filter response and effectively cancelled the effect. As Keelan indicated, though the appearance of the image has changed, the utility of the image has not as the filtering process may be reversed [1]. The application of the non-linear filter shows a severe degradation in the NEQ and thus the utility of the image. Not only has the increase in the noise-power in the low frequency regions caused a corresponding drop in the NEQ, but the decrease in the noise power throughout the spatial frequency range was too expensive in terms of the SFR drop in the image, leading to an overall decrease in NEQ. Analysis of filter behavior in this manner can yield more sophisticated insights as to the comparative efficiency of algorithms,

It should be noted that the PSNR of both the filtered images was higher than that of the original degraded image. If PSNR was used to guide filter design or image enhancement, it would have guided development in the wrong direction. PSNR is generally a poor proxy metric for any image science analysis.
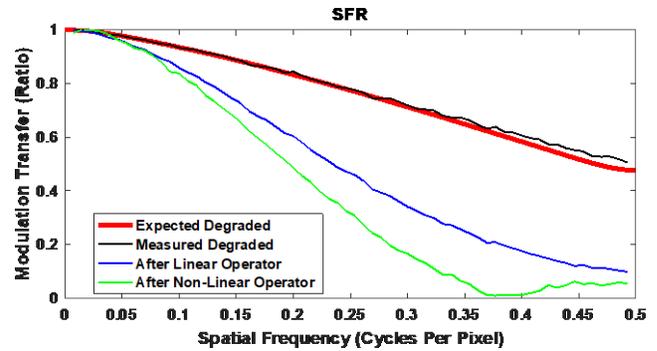
Figure 4 SFR as measured from the simulated targets using the dead leaves technique.
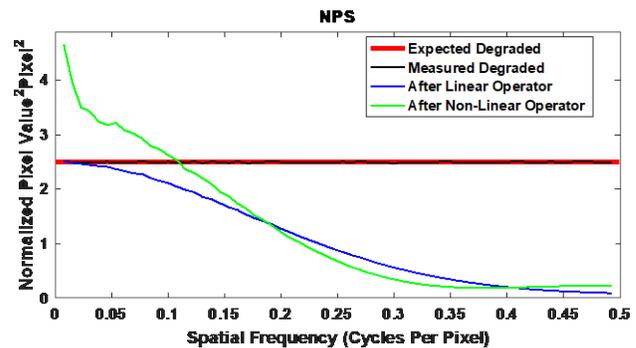
Figure 5 NPS as measured from the simulated targets using the dead leaves technique.
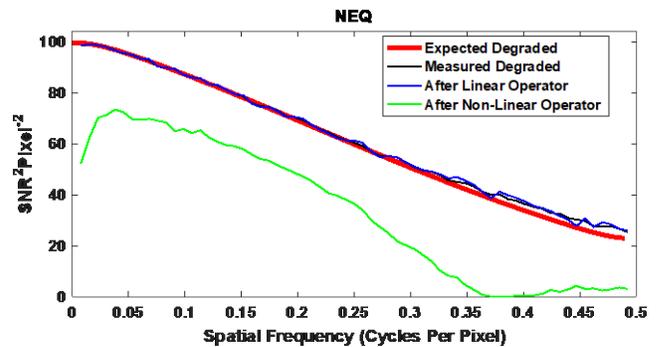
Figure 6 NEQ of the simulated targets as measured by the dead leaves technique.

Application to Real Systems

The dead leaves technique was also applied to several commercial camera systems to reveal any shortcomings. Images of test targets under different lighting conditions were captured and previous process used to calculate SFR, NPS and NEQ.

A significant difference between the simulation and real systems is the complexity of the code needed to correct for tonal non-linearity and geometric distortion. The existing technique established for dead leaves SFR measurement may be used for this purpose [3].

A camera with a 2.3 Mp automotive grade sensor and free running non-optimized ISP was tested as shown in Figure 7. The

target was illuminated at D50 with 2500, 250, 25 and 10 lux, Figure 8. Visual inspection shows a clear degradation in sharpness, details and noise with decreasing illuminance. The results, Figure 9, illustrate that the camera uses noise reduction techniques to reduce the image noise with a clear influence on the SFR. While the SFR reduces with decreasing illumination, the noise power also changes frequency distribution due the filtering, reducing further in high spatial frequency and increasing in low spatial frequency regions. As a result, we can observe a clear ranking of the different illumination levels in the results, demonstrating the ability of NEQ to discriminate the tradeoff between noise in the image and loss of details due to noise reduction.

A further test was conducted using a mobile phone with default settings and processing. The device captured an image of a test target containing the dead leaves pattern and a number of other structures [16] under D55 illumination with an intensity of 2000, 10, 5 and 1 lux, Figure 10. It can be observed that the device uses pixel binning dependent on the scene illumination level. While the image at 10 lux has the same pixel count as that at 2000 lux, it shows significant texture loss due to noise reduction. The images captured at 5 and 1 lux contain less pixels but show more apparent detail.

The results as shown in Figure 11 are plotted using the unit line pairs per picture height (LP/PH) for spatial frequencies on the x-axis as this enables the direct comparison of the results including the change in the image size.

The observed behavior is replicated in the objective results. The image at 10 lux has a good noise performance while is has a lower SFR. The image at 5 lux has more noise and a better SFR, both results in a higher NEQ for 5 and 1 lux compared to those for the 10 lux image.
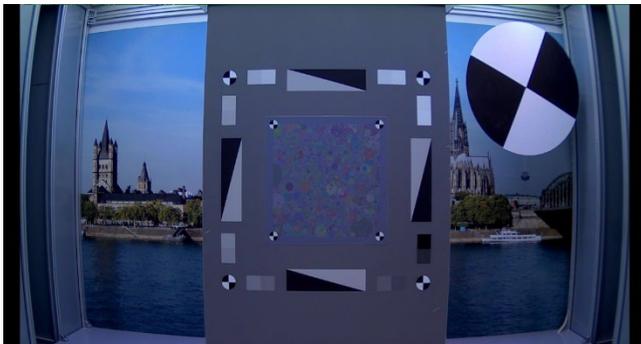
Figure 7 Experimental setup showing the dead leaves target inside a light booth, used for real camera test with an automotive grade sensor.

Conclusions

The authors have successfully demonstrated measurement of NEQ from a single target and verified performance of the technique through simulation and application to camera systems. In simulation measurement result matched expectation and NEQ was shown not to change under after the application of linear processes.

Applied in the RAW domain, NEQ can estimate the objective quality of exposures entering the ISP for processing. If the correct units are used NEQ can estimate the effective noise equivalent linear full well of the sensor. The technique can also be used as a pre-cursor for detection theory calculations.

If applied post-ISP, NEQ can objectively assess the benefit of image processing on signal preservation and the efficiency of image processing algorithms.

References

[1] B. Keelan, "Imaging Applications of Noise Equivalent Quanta", Proc. Electronic Imaging, Image Quality and System Performance XIII, Imaging Sci. and Technol. 2016.

[2] Uwe Artmann, "Measurement of Noise using the dead leaves pattern" in Proc. IS&T Int'l. Symp. on Electronic Imaging: Image Quality and System Performance XV, 2018, pp 341-1 - 341-6 https://doi.org/10.2352/ISSN.2470-1173.2018.12.IQSP-341

[3] Uwe Artmann, "Image quality assessment using the dead leaves target: experience with the latest approach and further investigations," Proc. SPIE 9404, Digital Photography XI, 94040J (27 February 2015); https://doi.org/10.1117/12.2079609

[4] Leonie Kirk, Philip Herzer, Uwe Artmann, Dietmar Kunz, "Description of texture loss using the dead leaves target: current issues and a new intrinsic approach," Proc. SPIE 9023, Digital Photography X, 90230C (7 March 2014); https://doi.org/10.1117/12.2039689

[5] ISO/TS 19567-2:2019 – Photography - Digital cameras- Part 2: Texture analysis using stochastic pattern

[6] Frédéric Cao, Frédéric Guichard, Hervé Hornung, "Dead leaves model for measuring texture quality on a digital camera," Proc. SPIE 7537, Digital Photography VI, 75370E (18 January 2010); https://doi.org/10.1117/12.838902J

[7] Jon McElvain, Scott P. Campbell, Jonathan Miller and Elaine W. Jin, "Texture based measurement of spatial frequency response using the dead leaves target: extensions, and application to real camera systems", Proc. SPIE 7537, 75370D (2010); doi:10.1117/12.838698

[8] J. C. Dainty and R. Shaw, Image Science: Principles, Analysis and Evaluation of Photographic-type Imaging Processes. London: Academic Press Ltd., 1974.

[9] H. H. Barrett and K. Myers, Foundations of Image Science (Wiley, Hoboken, NJ, 2004).

[10] M. Innocent et al., "Automotive 8.3 MP CMOS Image Sensor with 150 dB Dynamic Range and Light Flicker Mitigation," 2021 IEEE International Electron Devices Meeting (IEDM), San Francisco, CA, USA, 2021, pp. 30.2.1-30.2.4, doi: 10.1109/IEDM19574.2021.9720683.

[11] ICRU Report 54, Medical Imaging – The Assessment of Image Quality, Bethesda, MD, International Commission on Radiation Units and Measurements, 1996. (a) Eqs. 3.18 and 3.10. (b) Appendix E.

[12] R. Jenkin and P. Kane, "Fundamental Imaging System Analysis for Autonomous Vehicles", Proc. IS&T Electronic Imaging, Autonomous Vehicles and Machines 2018, Burlingame, California, 2018.

[13] E. C. Yeadon, R. A. Jones and J. T. Kelly, Confidence Limits for Individual Modulation Transfer Function Measurements Based Upon the Phase Transfer Function, Photogr. Sci. Eng. 14(2), 153-156 (1970).

[14] R. Jenkin, The Influence of CFA Choice on Automotive and Other Critical Systems, AIC 14th Congress Milano 2021.

[15] R. Jenkin and C. Zhao, "Radiometry and Photometry for Autonomous Vehicles and Machines: Fundamental Performance Limits", Proc.

IS&T Electronic Imaging, Autonomous Vehicles and Machines 2021, Burlingame, California, 2021.

[16] Multi Purpose Test Charts, https://www.image-engineering.de/library/blog/articles/1236-multipurpose-test-charts

Author Biographies

Robin Jenkin is a Distinguished Engineer at NVIDIA where he models image systems and performance for autonomous vehicles and other applications. He is also a Visiting Professor at University of Westminster within the Computer Vision and Imaging Technology Research Group. Robin received BSc(Hons) Photographic and Electronic Imaging Science (1995) and his PhD (2001) in the field of image science from University of Westminster. He also holds a M.Res Computer Vision and Image Processing from University College London (1996). Robin is a Fellow of The Royal Photographic Society, UK, and a previous board member and VP Publications of IS&T. He is also Secretary of IEEE P2020 WG on Image Quality for Autonomous Vehicles.

Uwe Artmann studied Photo Technology at the University of Applied Sciences in Cologne following an apprenticeship as a photographer and finished with the German 'Diploma Engineer'. He is the CTO at Image Engineering, an independent test lab for imaging devices and manufacturer of all kinds of test and calibration equipment for these devices. He is also the head of the standards department within VCX-Forum e.V. and member of various international workgroups regarding standardization of image quality measurement including IEEE and EMVA.
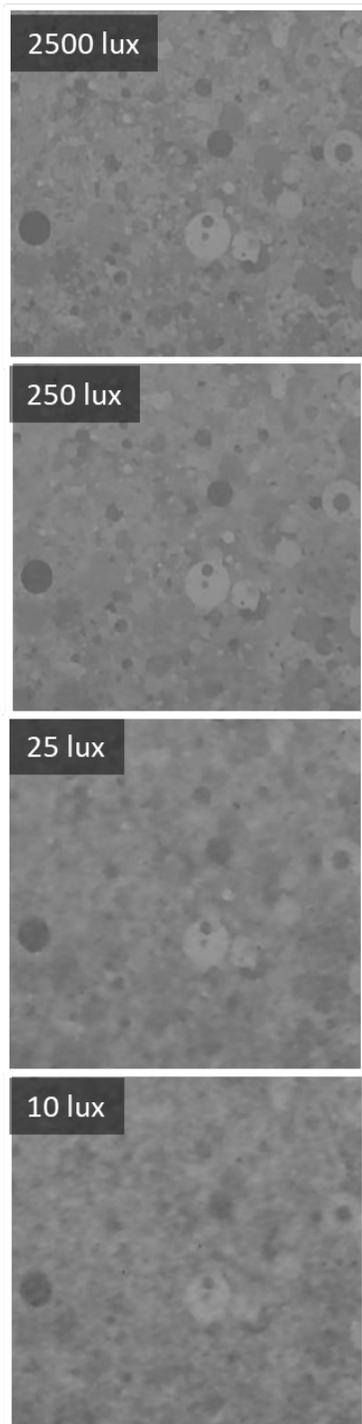
Figure 8 Examples of images captured using D50 at 2500, 250, 25 and 10 lux using a camera with a 2.3Mp automotive grade sensor and a free-running, non-optimized ISP.
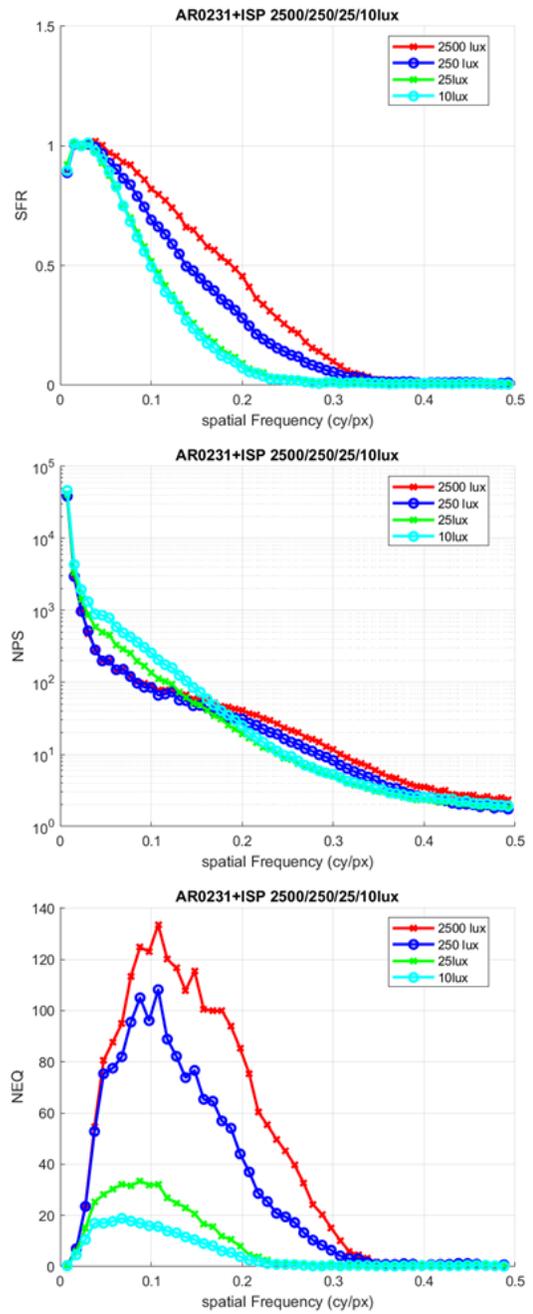
Figure 9 SFR, NPS and NEQ measured using a 2.3Mp camera with automotive grade sensor and free running, non-optimized ISP
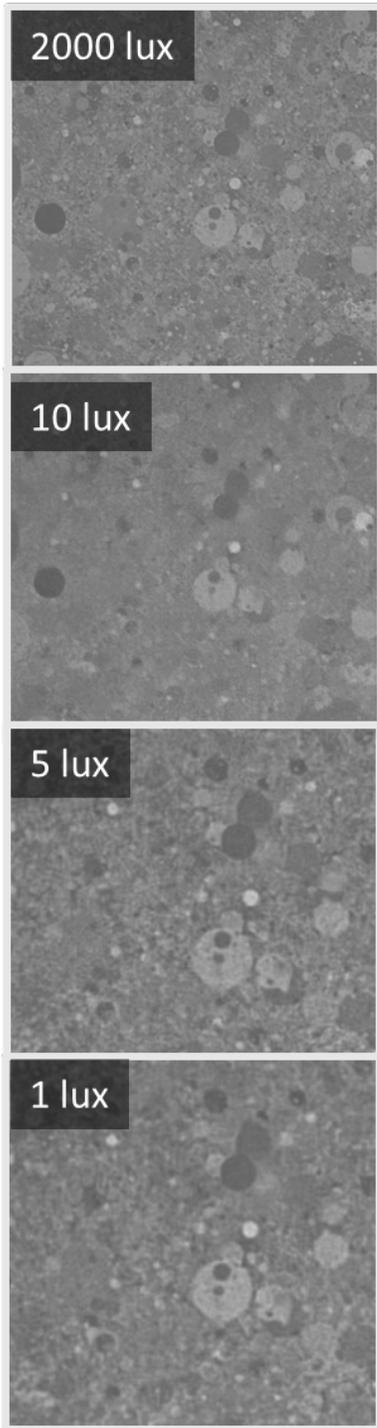
Figure 10 Detail of images captured at 2000, 10, 5, and 1 lux using a mobile phone in default settings.
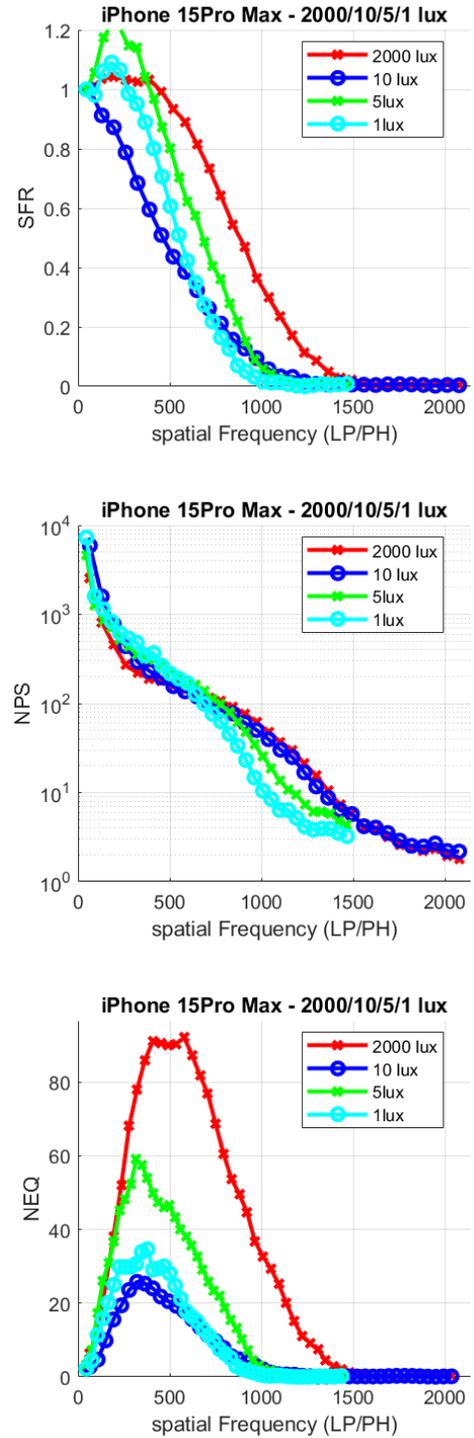
Figure 11 SFR, NPS and NEQ measured using a mobile phone camera in default settings and decreasing illumination.