

# Benchmarking of the ISO 12233 Slanted-edge Spatial Frequency Response Plug-in

*Don Williams*

*Imaging Research and Advanced Development, Eastman Kodak Company  
Rochester, New York 14650-1925*

## Abstract

The ISO 12233 Slanted-edge Spatial Frequency Response (SFR) Plug-in (V6.1.3) was independently evaluated in terms of accuracy, precision, and field robustness. Using synthetically generated slanted-edge images with known Gaussian frequency response as input, the SFR (i.e., MTF) results are compared under different combinations of peak SNR, edge angle, region-of-interest (ROI), signal bandwidth, binning, and FIR modulation gains against theoretical aims.

The results show that this plug-in is a very accurate, precise and robust tool for quantifying digital capture device MTFs. It is largely insensitive to edge angle, accurately predicts frequency response to at least 4x Nyquist frequency with proper binning and ROI, gives usable MTF estimates for peak SNRs as low as 10:1, and handles ill-behaved frequency morphologies introduced with FIR sharpening. Best of all, it is easy to use.

## Background

The slanted-edge MTF technique is an edge gradient MTF method specifically suited to MTF calculations for spatially sampled capture devices. Its main feature is the intelligent creation of a 1-D uniformly super-sampled edge profile from sequential lines of a 2-D native-sampled image whose line-to-line edge locations are slightly displaced from one another, as they would be with a slanted edge. Theoretically, this allows for unambiguous MTF estimates beyond a capture device's Nyquist frequency, always a limitation with sampled devices. Another claimed advantage is its alignment insensitivity. Indeed, the method requires that the edge be misaligned for super-sampling to occur. Details of the method can be found in seminal work by Reichenbach, et al.,<sup>1</sup> and complementary efforts of Fischer & Holm.<sup>2</sup>

Briefly, the plug-in works as follows. Once an image of the slanted edge is opened, the user selects a rectangular region-of-interest (ROI) with the marquee tool as illustrated in Figure 1. This defines the region over which the calculations are done. Note, the edge transition must intersect both the top and bottom lines of the ROI.

Once the ROI is selected, the algorithm linearizes and channel weights the ROI data based on user inputs. The following steps then occur. It

- estimates edge location of each scan line.
- regresses a best fit line through the center of the collective edge locations.

- registers each line based on the regression fit
- places data into uniformly sampled bins
- takes derivative of binned data to yield LSF
- Hamming windows the LSF
- performs DFT
- calculates MTF
- yields spreadsheet data of (frequency, modulation)




Figure 1.

While Reichenbach, et al. benchmarked their version of the algorithm, the software for doing so was unavailable for ubiquitous field use. Recently, however, an ISO committee has issued evaluation software in the form of a plug-in that uses the slanted edge method to calculate sampled capture device MTFs. This software is intended to bridge the gap between theory and generalized *en masse* field use and provide a relatively uncluttered path to quickly obtaining MTF measurements. Downloadable versions of the plug-in software and C code can be found at <http://www.pima.net/it10a.htm>.

As with any measurement technique, questions remained on its accuracy, precision, and field robustness. To that end, this independent benchmarking study was completed to evaluate the software. This was not meant as an exhaustive study to test all possible parameter permutations, but rather one that rationally samples these parameters in accordance with common sense engineering judgements.

Note that the term spatial frequency response (SFR) is used in place of MTF in the official title of the plug-in. This terminology choice came about to avoid confusion with photographic MTF as defined in ANSI PH 2.39. In this paper, for better or worse, the abbreviation MTF is used and assumes a linear relationship between digital count value and light intensity except where noted.

## Experimental

Synthetic images generated in Matlab were used as inputs for the plug-in. They were created using repeated error function calls on a line-by-line basis. This dictated that the MTFs of these edges be Gaussian. The edge location in each line was slightly shifted from the previous line's edge location leading to a slanted edge image. The edge angle was tunable by selecting the edge shift increment from line-to-line. The 1-D edge transition rate, or bandwidth, was tuned by selecting different error function increment rates in the Matlab routine. The synthetic images' modulations were set by simple scaling of the 8-bit color planes. For these simulations, the edge modulation was 60%, with a mean of 128. Though the MTF calculations were done on RGB data, each color plane's data was identical to the others.

The performance of the software was tested with respect to the following variables, some more extensively over others.

- Edge Angle
  - 1° thru 60°
- Peak Signal-to-Noise Ratio (SNR)
  - 5:1 thru 50:1
- Region-of-Interest (ROI)
  - (64h x 128v) and 18h x 128v)
- Signal Bandwidth
  - 50% frequency modulation at 0.20, 0.60, 1.20 pixel<sup>-1</sup>
- Bins/Native-Pixel
  - 4, 8, 16
- Frequency Modulation Gain
  - 5x raised cosine gain at 0.25 pixel<sup>-1</sup>
- Clipping
  - 120% and 110% modulation

The results and observations for the study follow.

## Results and Observations

The interpretation of algorithm and software goodness in the following section is based on a visual comparison with the theoretical aim MTF plot.

### Edge Angle

Edge angles of 1°, 2°, 5°, 10°, 20°, 30°, 45°, and 60° were evaluated for a Gaussian noiseless signal having 50% frequency modulation at 0.20 pixel<sup>-1</sup> (referred to as 20% bandwidth). This MTF might be considered typical for digital capture devices since its response just falls to zero at the Nyquist frequency of 0.50 pixel<sup>-1</sup>. A (64h x 128v) ROI window was used except for the larger angles, in which case the horizontal dimension was maintained with the vertical dimension being adjusted appropriately. The results are shown in Figure 2.

The software's robustness in handling edge angle variations was very good. The extreme angles attempted in this simulation were actually meant to break the analysis software and are outside the Plug-in User Guide's recommendations. As the reader can see from Figure 2, except for the 45° case, the other estimates lie virtually on

top of each other. In fairness, the 45° provides no unique super-resolution data from line to line, so poorer estimates in the MTF are expected.

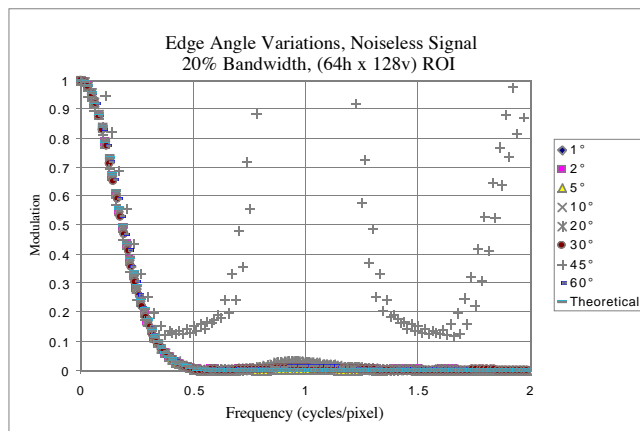


Figure 2.

A spurious response was noted for the 10° and 20° cases at 1.0 pixel<sup>-1</sup>. The source of this response was not traced. It was observed, however, that increasing the number of bins diminished its amplitude to the same levels as other edge angles without affecting other frequency responses.

### Peak Signal-to-RMS Noise Ratio (SNR)

Because of its good agreement with the aim theoretical curve, the 5°, 20% bandwidth signal from the previous section using noiseless signals was used as a reference case to compare SNR levels of 50:1 (34 dB), 40:1 (32 dB), 20:1 (26 dB), 10:1 (20 dB), and 5:1 (14 dB). As of this writing, worst case digital cameras generally rank about 30 dB.

Using a ROI window of (64h x 128v), Good results were achieved down to SNR of 20:1. The results are shown in Figure 3.

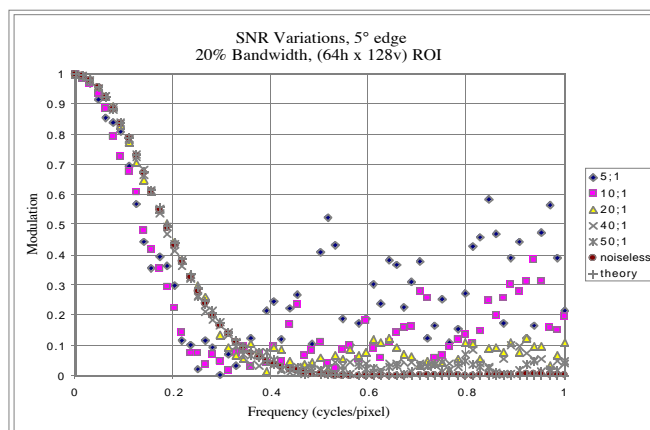


Figure 3.

Though the estimates became noisy beyond the cutoff frequency for the 20:1 SNR, the random nature of the MTF estimates beyond the cutoff frequency actually helped to classify these frequencies as void of any true signal content. Below the cutoff frequency, the data was well-behaved and accurate.

The 10:1 and 5:1 SNR cases revealed under-estimates of the frequency response with extremely noisy estimates beyond the Nyquist frequency. By using a much narrower ROI though (18h x 128v), the MTF estimates for these low SNR cases agreed well with the theoretical aim. The noisy estimates beyond the Nyquist frequency were also reduced, albeit at a reduced frequency resolution. This indicates that selecting a ROI width far beyond the zero-slope portions of the edge transition may actually lead to poor estimates of the frequency response. This behavior has, in fact, been observed in the past for other methods of frequency response characterization<sup>3</sup>. The results of low SNR signals using the reduced ROI are shown in Figure 4. MTF estimates using the narrower ROI for higher SNR remained unchanged from those of Figure 3.

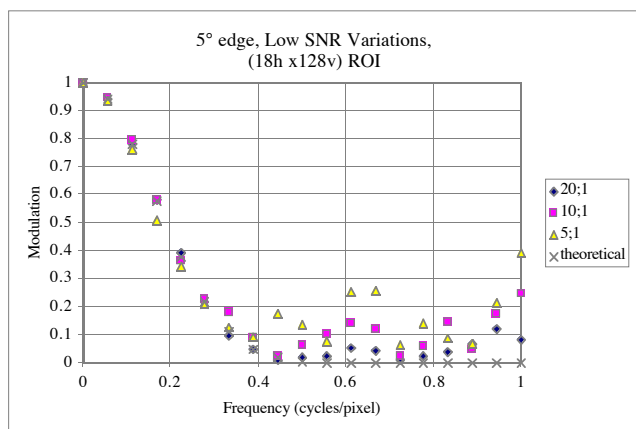


Figure 4.

### Signal Bandwidth, and Bins/Native-Pixel

In addition to the 20% bandwidth signal used so far, 60% and 120% bandwidth signals were also tested for both noiseless and noisy signals. The results for the noiseless signals with the standard software are shown in Figure 5.

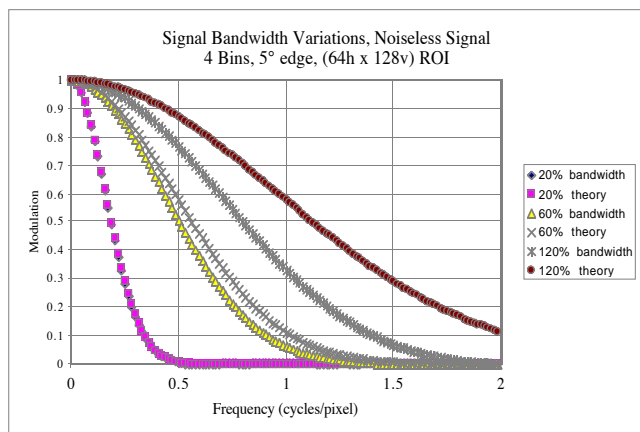


Figure 5.

As coded, the software underestimates the true MTF responses at these higher bandwidths. This behavior was observed despite edge-angle and SNR changes. The reason for this is traceable to the number of bins/native-pixel used in the existing code. This binning is a means of spatially

quantizing the super-sampled edge locations, which are real-valued, back to some uniformly sampled integer fraction of the original natively sampled locations. In this version of the software, the number of bins is hard coded to four. In other words, the effective super-sampled edge is uniformly sampled four times higher than in the native sampling. Heuristically, one can imagine that, to some extent, the binning process acts as an averaging operator to help reduce noise in the MTF estimate. This result was observed in testing albeit with a tradeoff in accuracy.

The original bin value of 4 is probably a good choice based on classically designed digital capture systems where MTF modulations much beyond Nyquist frequency are not desired in order to reduce aliasing. Where aliasing is not a concern though, and significant modulation beyond Nyquist is likely, an increase in the number of bins is suggested, and can be modified in the source code (the “alpha” variable). This was done, with the results shown in Figure 6. It reveals the extent to which the MTF estimates align with the theoretical aims for the noiseless higher bandwidth signals for increased number of bins. With eight bins the 60% bandwidth MTF is accurate. Sixteen bins were necessary to align the 120% signal with its aim.

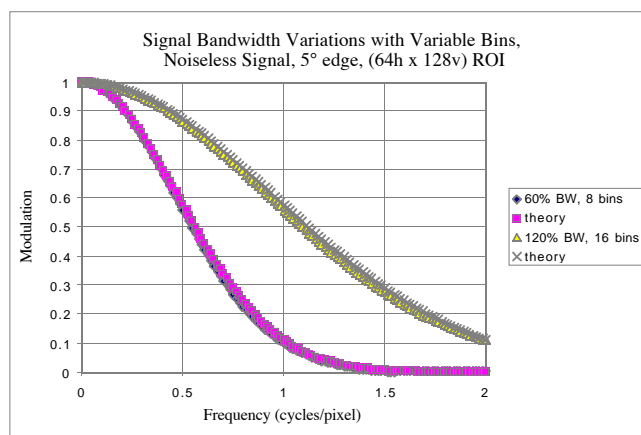


Figure 6.

Similar to the 20% bandwidth case of Figure 3, low SNR cases did not align well with their theoretical aims for the higher bandwidth signals either. But by using low aspect ratio ROI windows of (18h x 128v) these estimates did track well, similar to the 20% bandwidth signal in Figure 4. These similar results are shown in Figure 7 and Figure 8 and indicate that with a greater number of bins and low aspect ratio ROI, very reasonable MTF estimates can be achieved with low SNR signals—down to 20:1 for high bandwidth systems.

### Frequency Modulation Gains and Clipping

A limited set of data was used to test the software’s ability to predict frequency modulation gains often imposed by FIR sharpening filters. Because these filters often clip the resulting signal, the effect of this clipping on the MTF estimate was also tested.

Naturally, clipping is considered a non-linearity, and as such is inconsistent with true MTF estimates. Nevertheless, it is good to know the frequency signature that clipping

imposes in order to diagnose ill-behaved MTF morphologies.

A noiseless, 20% bandwidth Gaussian signal with a raised cosine enhancement was used. Unclipped, and two clipped signals with effective modulations of 110% and 120% were tested. The results are shown in Figure 9.

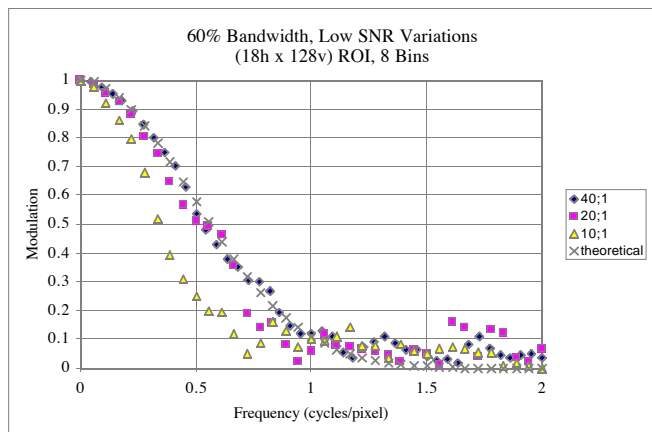


Figure 7.

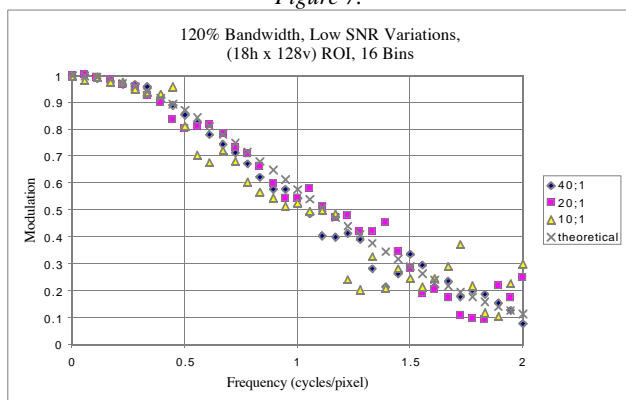


Figure 8.

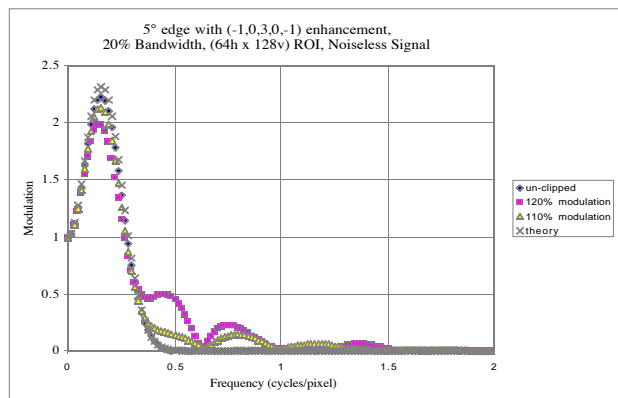


Figure 9.

The MTF estimates for the unclipped data compare well with the theoretical aim. The clipped data, however, introduces a sinc-like ringing into the MTF estimate. This same type behavior was noted, though not shown here, with unclipped, but quantized, data (4-6 bits/pixel). This is not considered to be a problem with the plug-in software itself,

but rather an artifact of nonlinearities introduced by the clipping process.

## Conclusions

This study has shown that for typical digital capture devices, the ISO 12233 Slanted-edge SFR Plug-in (V6.1.3) is a very accurate, precise, and robust tool for quantifying these devices' MTFs with a linear data assumption. Though Gaussian functions were used in this study, one could probably generalize these results to any continuous MTF functional form.

The MTF estimates were largely insensitive to edge angle. A range from 1° to 60° was tested with only the 45° angle being unacceptable, unsurprisingly. A slight spurious increase in modulation was detected at the sampling frequency for certain edge angles. It is recommended that angles be kept between 3° and 30° from the vertical.

Only for SNRs (peak signal-to-rms dark noise) below 20:1 and at bandwidths starting to exceed 60% (half modulation at 0.60 pixel<sup>-1</sup>) do the MTF errors become unacceptable. The SNRs and bandwidths where failure occurs are unusual for typically designed digital capture devices. It was shown that increasing the number of bins improved the MTF estimates for higher bandwidth signals, albeit with a slightly noisier MTF estimate.

As the vertical-to-horizontal (v:h) region-of-interest (ROI) aspect ratio increased, less noisy and more accurate MTF estimates were observed. The v:h ROI ratio for the edge image should be as small as possible while still maintaining the desired frequency resolution. Good engineering judgement should be used when selecting the ROI to exclude an excessive number of samples/line. Noisier MTF estimates occur when the ROI is excessive in the horizontal dimension.

The plug-in software also accurately predicted boost behavior introduced with the FIR kernel tested. However, the user is cautioned on introducing clipping when using high gain boosts with these kernels because of the ill-behaved MTF estimates.

## Gratitudes

Thanks to members of the ISO/TC42/WG18 "Electronic Picture Still Imaging", for providing the evaluation copy of plug-in software. Special thanks to Ken Parulski of Eastman Kodak Company, and Andrew Juenger and Eric Higgins of Polaroid Corp. for the execution and follow through on the concept of this useful tool. Thanks also to Stacy Mayo for providing Matlab programming support for synthetic image data generation.

## References

1. Reichenbach, S.E. Park, S.K. and Narayanswamy, R., "Characterizing digital image acquisition devices", *Opt. Eng.*, Vol. 30, No. 2, Feb. 1991, pp. 170-176.
2. Fischer, T.A. and Holm, J., "Electronic still picture camera spatial frequency response measurement", *IS&T's 47th Annual Conference/ICPS 1994*, Vol. II, pp. 626-630.
3. P. Burns, personal communication.