The Effects of misregistration on the dead leaves crosscorrelation texture blur analysis

Robert C Sumner, Ranga Burada; Imatest LLC; Boulder, CO Noah Kram; Rochester Institute of Technology; Rochester, NY

Abstract

The dead leaves image model is often used for measurement of the spatial frequency response (SFR) of digital cameras, where response to fine texture is of interest. It has a power spectral density (PSD) similar to natural images and image features of varying sizes, making it useful for measuring the texture-blurring effects of non-linear noise reduction which may not be well analyzed by traditional methods. The standard approach for analyzing images of this model is to compare observed PSDs to the analytically known one. However, recent works have proposed a cross-correlation based approach which promises more robust measurements via full-reference comparison with the known true pattern. A major assumption of this method is that the observed image and reference image can be aligned (registered) with subpixel accuracy. In this paper we study the effects of registration errors on the calculation of texture-based SFR and its derivative metrics (such as MTF50), in order to determine how accurate this registration must be for reliable results. We also propose a change to the dead leaves cross-correlation algorithm, recommending the use of the absolute value of the transfer function rather than its real part. Simulations of registration error on both real and simulated observed images reveal that small amounts of misregistration (as low as 0.15px) can cause large variability in MTF curves derived using the real part of the transfer function, while MTF curves derived from the absolute value are significantly less affected.

Introduction

The best-established characterization of image sharpness and resolution is the Modulation Transfer Function (MTF). The MTF of a system dictates how spatial frequencies in a scene are boosted or attenuated (relative to a normalized DC gain) by the imaging system—including optics, finite photosensor size, etc. An implicit assumption of MTF determination is that of a linear model relating spatial patterns of light in a scene and the spatial distribution of pixel values recorded in an image. That is—at least locally for a given point in the image field—we assume the pixel values are related to the illumination coming from the scene via a linear, time-invariant (LTI) system (or more appropriately for imaging applications: linear and shift-invariant, or LSI).

When dealing with optics and sensors directly, or simple linear image processing, this LTI assumption is often an appropriate model. The ISO 12233:2014 standard [1] prescribes two powerful methods for determining the MTF of a system, the slanted-edge method and the Siemens Star method, both requiring a linearized image. However, the images we consume are increasingly subjected to substantial non-linear processing, which breaks this assumption. One of the primary applications of such processing is image denoising, which can roughly be described as discriminating between which pixel value variations are supposed to be there (content) and which are not (noise). It is often extremely difficult for even the most advanced denoising algorithms to make this distinction correctly every time. The primary casualties of imperfect (usually over-aggressive) denoising are low-contrast, high-frequency image structures, often known as *image texture*, which are indistinguishable from noise in the eyes of the algorithm.

It is extremely difficult to characterize the effects of modern denoising algorithms on all possible images taken by a camera because they are often very content dependent. In practice this goes even beyond the analysis-confounding effects of being local but non-linear. Edge-preserving smoothing algorithms such as bilateral filtering [2] and anisotropic diffusion [3] perform quite differently in flat areas and near edges. Bilateral filtering, Non-Local Means [4], and BM3D [5] (as well as a host of other denoising algorithms) operate by searching in a region around a pixel for similar content and basing their action on what they find. Moreover, the latter two work at a *patch-wise* level, determining similarity by comparing structure within a small (e.g., 8×8 px) window, compared with other windows in the region. This all means that it is impossible to know how such an algorithm will act on a given (small-scale) image feature in the wild where the content around it in the image may be different every time.

Because of this content dependent action, an MTF measurement derived from a slanted edge (a relatively simple image structure to denoise without degradation) as per ISO 12233 would not faithfully indicate the effect on spatial frequencies in other parts of the image where the true structure is not so obvious. We instead desire an SFR measure which appropriately conveys the effect of the system on different frequencies, whatever context they may appear in in the image. The Siemens Star, as a more complex structure, is more susceptible to non-linear smoothing and thus MTF measurements derived from it may be less optimistically biased. However, since its structure is both fixed and not representative of many real-world image structures, it is not fully adequate for describing the average effect of content-dependent algorithms.

Dead Leaves Texture Analysis

One effective way to observe the effects of a full system, including processing, on fine details found in real images is to test how it responds to (pseudo-)realistic image content. The "Dead Leaves," or "Spilled Coins," test pattern and analysis technique of Cao et al [6] emulates a $1/f^2$ power spectral density (PSD), which has been empirically observed to characterize natural images, by generating a field of randomly-placed and randomly-sized overlapping circles, as illustrated in Figure 1.



Figure 1. Monochrome dead leaves pattern texture detail.

The original paper described a method for calculating the "Texture MTF" of a camera from an image of this pattern by making use of this *global*, *statistical* analytic assumption about the ground truth pattern, producing a semi-reference metric. Kirk et al [7] introduced a successor to this method which makes use of a known ground truth reference image of the dead leaves pattern. By comparing a test image directly to the reference image, pixel by pixel, this *dead leaves cross correlation* method can measure *local* degradations in the image. This is especially important due to the aforementioned local-content dependent nature of many denoising algorithms.

This tremendous leap in the strength of the assumptions made by this method over its predecessor—which in turn allows a commensurate leap in measurement strength—is dependent entirely upon the assumption of accurate alignment of the *reference image* of the pattern and the *actual observation* of it (i.e. the pattern as it appears in the test image). Without this, the assumptions of the cross-correlation MTF calculation fall apart.

In both [7] and the proposed ISO 19567 Part 2 standard based on it, registration of the two images is achieved by localizing four registration marks on the corners of the patterned area. Using the precise locations of these marks, the reference image is transformed to be in perfect alignment with the observed one. These documents note that these corners must be located with subpixel accuracy, but do not comment on what the effect on the calculation will be if this accuracy is not achieved.

The remainder of this paper is organized as follows. First, we review the two methods for determining a Texture MTF curve using the dead leaves pattern, with special emphasis on the implications and nature of these models. We then propose a change to the calculations used in the cross-correlation method, specifically the use of the absolute value of the transfer function rather than the real part. We then describe an experiment for exploring the effects misregistration of the reference and observed images can have on MTF calculations, and show results for both the existing method and our proposed change. Finally, we offer interpretation and conclusions based on the results.

Dead Leaves Texture MTF Calculations

The dead leaves pattern described in the previous section has many desirable properties, including:

- 1. The desired $1/f^2$ frequency power distribution, which also has the extremely useful characteristic that a signal with such a distribution will be *scale invariant*.
- 2. An image of this pattern will contain many different image features—ranging from busy areas of tiny circles to broad, contrasted edges—which will elicit different behaviors from content-dependent processing.

Power spectral density is well known as describing the *average* power a random signal has at different frequencies. It contains only the magnitude, not the phase, of this information. The PSD of the output of a linear system (in our case, a digital image), characterized by transfer function H(f), is related to the PSD of the input (the test chart target) via the equation

$$PSD_{image}(f) = |H(f)|^2 \cdot PSD_{target}(f)$$
(1)

The classic dead leaves analysis method (which we will refer to as the "direct" method, following [7], as it directly makes use of the image PSD) is to determine the magnitude of the transfer function by taking the square root of the ratio of the observed PSD to the analytically assumed PSD. This was expanded upon by McElvain et al [8] to include compensation for the noise PSD, thereby making it more robust for noisy, real-world images.

$$MTF_{DL-Direct}(f) = \sqrt{\frac{PSD_{image}(f) - PSD_{noise}(f)}{PSD_{target}(f)}}$$
(2)

A major shortcoming of this method is that it uses a relatively weak assumption about the ground truth signal: that its power spectral density has a certain analytical form. This limits the test to only making use of the *statistics* of the reference pattern, not the known pattern instance itself (as essentially all other test chartbased image quality analyses are), which in general lets you determine a lot less about a quantity. Relatedly, the assumption is about global quantities only—the average Fourier component powers—not directly making use of any local structures in the image, such as the edges of the random circles themselves.

Dead Leaves Cross-Correlation MTF

The Dead Leaves Cross-Correlation MTF method was first proposed by Kirk et al [7]. It makes use of a patterned area of overlaid circles, similar to the original dead leaves measurement. The new method, however, effectively moves from using this pattern for a *semi-reference* measurement value—using only statistical knowledge about an image of the pattern—to a *full-reference* calculation wherein we actually directly compare observed pixel values to pixel values of a ground truth image.

This method works by making use of four registration marks at the corners of the patterned field, as seen in Figure 2 (Left). These corner locations are known beforehand in the reference image and are detected at test time in the observed image. By using these coordinates as anchor points, it is possible to determine a homography that transforms the reference image from its intrinsic coordinate space to the observed image space, aligning it pixel for pixel with the observed instance of the pattern.

Once the observed and reference patterns have been aligned, this method makes use of the Cross-Power Spectral Density (CPSD) of the two signals. If a random, wide sense stationary signal x with Fourier spectrum X(f) and PSD $\Phi_{XX}(f)$ is passed through a linear system with frequency response H(f) to produce a similarly defined output signal y, two relations will be true. First, the power spectral density of y is

$$\Phi_{YY}(f) = E[Y(f)Y(f)^*] = E[H(f)X(f)H(f)^*X(f)^*]$$
$$= |H(f)|^2 \Phi_{XX}(f)$$
(3)

Here E[] is the expected value, and * denotes the complex conjugate. This is simply a restatement of Equation 1, and is the origin of the original dead leaves direct method. The other quantity of interest, the CPSD, is given by

$$\Phi_{YX}(f) = E[Y(f)X(f)^*] = H(f)E[X(f)X(f)^*]$$

= $H(f)\Phi_{XX}(f)$ (4)

The cross-correlation approach to determining H(f) is similar to the previous method:

$$H(f) = \frac{\Phi_{YX}(f)}{\Phi_{XX}(f)}$$
(5)

Whereas previously the denominator was an analytically assumed function, which in practice was rather hard to achieve perfectly with a finite pattern of overlaid circles, here the denominator can be achieved simply by taking the square of the Fourier coefficients of the spatially-registered reference image (which is x in the above).

The numerator of Equation 5 is the CPSD, which classically is interpreted as the amount of power shared at a given frequency by both signals, and the phase difference between them at that frequency, *on average*. Since an LTI system behaves the same across an entire image, its average behavior *is* exactly its general behavior, and Equation 5 will not yield any different result from the direct dead leaves method (in the noiseless case).

However, with non-LTI systems this is not true, as some spatial frequencies may be better preserved in some parts of the image (*e,g.* at high contrast edges) and removed from others (low-contrast textures) as discussed in the introduction. So while the calculated CPSD may describe a global average from around the image field, that does not mean it does not make use of local phase information similarities between the observed and reference images.

To get some intuition behind this claim, consider the fact that the CPSD is the same as the Fourier transform of the crosscorrelation function of the observed and reference images. Crosscorrelation can be interpreted as a matching process, wherein the two images are shifted relative to each other by various amounts, and at each shift amount, a measure is taken of how similar the structure of the two images. Two images that have very similar structures and are perfectly aligned will have a very high correlation score. However, if you shift one image even just one pixel to the side, the edges which had so nicely aligned previously will be out of sync and the correlation score for this shift will suffer a significant drop. The sharper this drop is (*i.e.* the more the two images are structurally similar), the closer to a delta function the cross-correlation function will be, and thus the broader its Fourier transform (the transfer function) will be.

When cross-correlation is used, the result incorporates not just the average frequency powers in the observed image are (as we use in the "direct" method), but about how those frequencies align locally with those of the reference image. So while the cross-power spectral density still tells us about averages of common frequency content between two signals, it does so using local areas of similarity.

The Role of Image Registration

Accurate registration of the reference image to the observed image space is of vital importance to the dead leaves crosscorrelation method.

The standard way to register the two patterns as of this time is to independently detect and localize the special fiducials at the corners of the dead leaves pattern area. If one of these registration marks is localized incorrectly in the observed image, then when the reference image is transformed, it will be transformed incorrectly. This will make the edges of the circles that comprise the ground truth pattern out of alignment with their test image counterparts. It could also change the shape of some circles to be more elliptical, or have otherwise detrimental effects that will make the images dissimilar.

Figure 2 (Right) illustrates how even a small amount of registration error (0.5px on each of the corners, in random directions) can change the values between two otherwise identical images. Here the same (grayscale) reference image was transformed twice with slightly different corner locations, but pixel value differences, indicated by a green-magenta spectrum, abound due to the different samplings that happen.



Figure 2. (Left) Imatest dead leaves chart with four registration marks at the corners of the patterned field. Each corner, once detected in a real image, is used as an anchor to fit a homography from the reference image space to the observed image space. (Right) Detail of two misregistered (black and white) dead leaves pattern areas using the MATLAB imshowpair() function. Differences in pixel values between the two images are seen on a magenta-to-green spectrum, a perfectly registered image would be grayscale. The two compared images differ by having the corners of one misplaced by 0.5px in random directions.

In this paper, we do not study the effectiveness of any particular registration-mark detecting routines. Instead, we ask a more general question about the entire class of such algorithms: how accurate do they need to be so that the Texture MTF results are not affected by their performance?

Any bias in the detection routines—meaning they always tend to misdetect the registration marks in the same direction—will effectively simply add a constant shift to the transformed reference image in the opposite direction of the bias. For the remainder, we presume unbiased detection routines, which may cause misregistration error in any direction, uniformly.

Note that in a few special cases—all registration error happens radially from or towards the center of the pattern area for example—the transformed reference image will have a Fourier spectrum that may be relatable to the expected one analytically. However, in general, independent misregistration of the four corners will cause a slight perspective effect which we will be unable to account for analytically.

Proposed Change to Absolute Value of Cross-Correlation

We propose a change to the dead leaves cross-correlation MTF measurement, specifically the use of the absolute value of the transfer function instead of the real part.

Firstly, the use of the absolute value is more in line with the traditional definition of the MTF of a linear system. For example, if the real part is used, there is nothing in the mathematics that prevents a negative value. Note that in the case of an actual LTI system with a symmetric point spread function, the real part and absolute value of the transfer function are the same. In any other case, they will be different.

We specifically note that even though taking the absolute value discards the phase information of the transfer function, this does not reduce it to a calculation equivalent to the original "direct" dead leaves MTF calculation. In the LTI case it would be, but this measurement is specifically intended to measure non-LTI systems.

The cross-correlation approach, as stated above, still makes use of local comparisons of structure around the image. The calculation is no longer based simply on global frequency content but specifically on the alignment and similarity of structures in the observed and reference images. That is, we are very much using the aligned phase information of the image and the pattern when performing this calculation- it is implicitly incorporated into the cross-correlation procedure, regardless of if we are discarding the transfer function's phase later on by taking the absolute value.

This difference between the direct method and the proposed absolute value method is illustrated by example in Figure 3. A test image was generated directly from the dead leaves reference image, then blurred with a Gaussian blur ($\sigma = 0.7 \text{px}$), and finally a moderate amount of bilateral filtering was applied. The MTF curves generated by the three methods—the direct method of Cao et al, the cross-correlation method with real part used, and the cross-correlation method with absolute value used—are shown. The cross-correlation method with absolute value is distinctly different from the direct method (and in fact is distinctly similar to the real part result), showing that they are *not* the same for non-linear systems despite the fact that the absolute value is used in both.



Figure 3. MTF curves for a non-linearly smoothed image, generated by the original dead leaves method, the cross-correlation method using the real part of the transfer function, and the proposed cross-correlation method using the absolute value of the transfer function. Note that using the absolute value of the cross-correlation does not make it similar to the original ("Direct") method.

The resulting MTF is still an average over the whole region of the image containing the textured field, but the values which are being averaged do make use of local information shared (or not shared) by the images. Moreover, taking the real part of the transfer function was not necessarily "making use of the phase information" in any particularly well motivated way.

The following toy example illustrates a simple situation where texture MTF derived using the real part of the transfer function deviates significantly from expectation, while the absolute value MTF behaves as expected.

Consider a test image, with Fourier spectrum denoted $Y(f_x, f_y)$, consisting of a dead leaves pattern observation which is exactly the same as the reference image, similarly denoted $X(f_x, f_y)$, but simply shifted to the right by one pixel. Obviously, it has not suffered any loss of sharpness and we would expect a calculated MTF to reflect this. Mathematically, adding a shift to the image is equivalent to adding a linear phase term to its Fourier spectrum, so for this one-pixel right-shift we get

$$Y(f_x, f_y) = e^{j2\pi f_x} \cdot X(f_x, f_y)$$
(6)

Applying the procedure previously described for determining the transfer function (or even from simple inspection of the above equation), we see that the transfer function is $H(f_x, f_y) = e^{j2\pi f_x}$, the real part of which is visualized in Figure 4.



Figure 4. System response of a one pixel shift to the right. (Left) The real part of the transfer function. (Right) The radial mean of the real part.

When taking the radial average of the real part of this function as prescribed in the dead leaves cross procedure, it reduces to the 1-dimensional MTF curve shown in Figure 4. Not only does this MTF imply a significant loss in image sharpness, but it actually has negative values, which is forbidden by most reasonable definitions of MTF.

Using the absolute value of the transfer function on the toy example above produces the expected MTF, which is a perfect MTF of 1 for all frequencies. This toy example illustrates how using the real part of the transfer function can be inconsistent with our expectations of image quality measurements. Moreover, this example actually does describe what would be a compounding effect to a cross-correlation real-part MTF curve in practice if the registration mark detection algorithm exhibits any bias in the direction of its error.

We emphasize that this proposed change is motivated purely by the mathematical nature of determining an equivalent to MTF by means of cross-correlation. It does not stem from particulars of the implementation of the dead leaves method or from real world challenges such as imperfect registration accuracy.

Methodology

In order to isolate and study the effects of misregistration on cross-correlation derived MTF curves, we performed the following experiment. We started with a high resolution raster image of our dead leaves chart (5163×5163px texture region), rescaled it with

bicubic interpolation to a smaller size $(730 \times 730 \text{px} \text{ texture region})$, and applied a simple Gaussian blur ($\sigma = 0.7 \text{px}$) to generate our simulated "observed" image. Since we know exactly the locations of the registration marks at the corners of the pattern field in the reference image and also the exact transformation used to generate the observed image, we can also determine the exact location of the registration marks in the observed image.

If these true registration mark locations were fed along with the observed image to the dead leaves cross-correlation routine, the homography calculated by the routine would be exactly the same as the one we used to generate the simulated image, and so the reference image would be perfectly aligned. In this case we would see a perfect MTF curve corresponding to a Gaussian transfer function, regardless of if the real part or absolute value method was used.

However, in lieu of using the true registration mark locations, we generate slightly perturbed locations and pass these to the calculation routine instead. This causes a slightly different homography to be calculated and thus a slightly differently resampled and aligned reference image to be generated.

Registration error vectors were generated independently for each of the four registration points at the corners of the field. The vectors were generated in random directions with a uniform distribution over all angles. This simulates a registration-mark detection routine which is unbiased in the orientation of its errors. As previously mentioned, if there is a bias in a specific routine (*e.g.* the routine tends to label the marks as 0.5px to the left of where they truly are), this will theoretically lead to a phase shift term in the transfer function which will affect results generated using the real part, but not affect the results generated using the absolute value.

The magnitudes of the random errors were chosen to be the same for all four corners for each given simulation, for simplicity. Thus, for a given simulation, all four registration marks will be, *e.g.*, 0.5px off from the true locations, but in four random directions. We simulated 100 different misregistrations (different sets of random directions) for each magnitude or registration error, ranging from 0 to 10px. Though not exhaustive, we believe this provides a sufficient initial sampling of the space of MTF shape variations.

We emphasize that this is only a study of the effects of misregistration on the algorithm's output, and so we primarily look at simulated observations with zero noise. Of course, real world effects such as noise, sharpening, and denoising will affect the shapes of these curves in various ways, but we consider this registration error effect to be orthogonal to and compounding on top of those. However, to observe the effects on a real, non-linearly processed image, we repeated this process (without the simulated blur) on a test image from a Nexus 5X mobile device. In this case, the "true" registration mark locations were identified with sub-pixel accuracy by eye.

Results and Discussion

For the sake of understanding the consequences of registration error on both the method currently in use by practitioners as well as the method we have proposed, all figure here will display similar graphs for each method, side by side. So that we need not repeat the observation in the discussion of each plot, we note up front the general trend that the variation in MTF curves is less (often significantly so) for those derived with the absolute value of the transfer function, compared to those using the real part.

Figure 5 shows the effects of increasing amounts of registration error for a single random instance of misregistration directions of the four registration points on a simulated image. Note that the 0px miregistration case is the "ideal" MTF curve for this system which was modeled with a Gaussian blur of 0.7px standard deviation. This representative example shows an obvious general trend: as you increase the magnitude of the registration error, MTF curves drop at all frequencies. This makes intuitive sense, as we are attempting to correlate the observation with a warped reference that looks increasingly dissimilar. Though it is reasonable to expect no more than one or two pixels' worth of registration error from a localization routine, we include here unrealistically large amounts as well to help illustrate the trend.

Whereas Figure 5 displays MTF curves corresponding to different amounts of misregistration in the same direction, each plot in Figure 6 shows entire random populations of MTF curves for a single magnitude of misregistration. Each population consists of 100 MTF curves generated by misregistration in a random direction, and so represents the range of possible MTF curve shapes that might be generated from any image where that amount of registration error is expected. Obviously, a narrower spread is better and a wide spread indicates that there should be less confidence in any individual MTF curve observed in practice.



Figure 5. Simulated observation MTF curves generated using a reference image misregistered to the test image by different amounts. The test image was a simulated observation with Gaussian blur, $\sigma = 0.7$ px.



Figure 6. Simulated image populations of MTF curves generated by registration errors of various magnitudes in 100 random directions.

Real Image Example

Figures 7 and 8 display the same families of results, but generated from a real, processed image taken on a Nexus 5X. Interestingly, Figure 7 shows minimal amounts of change in the MTF curve shape for small amounts of misregistration (≤ 1 px in magnitude) using either calculation method. However, Figure 8 indicates that registration error in different directions (of the same

magnitude) can still lead to a nontrivial spread in overall MTF curve shape. Interestingly, when using the absolute value method in the real image case, the MTF curve variations tend to converge at higher frequencies closer to the "true" MTF curve (when the registration is accurate, in this case chosen by eye).

These results, both on the simulated image with simple linear Gaussian blur and the real image with non-linear processing, indicate why in [7] the authors note, "The spatial matching process



Figure 7. Real observation MTF curves generated using a reference image misregistered to the test image by different amounts. The test image was taken with a Nexus 5X, with true registration mark locations identified by eye.



Figure 8. Real image (Nexus 5) populations of MTF curves generated by registration errors of various magnitudes in 100 random directions.

results." The method's sensitivity to inexact registration are clear from this investigation, and should be known to practitioners. Fortunately, this effect is lessened fairly significantly by the change to using the absolute value of the transfer function.

Registration Error's Effect on MTF50

In many applications, MTF curves are often reduced to single value summary metrics such as MTF50, MTF10, MTF at Nyquist/4, etc. These often act as very rough correlates to quantities of interest, *e.g.*, MTF50 is often used as an indicator of sharpness and MTF10 is often used as an indicator of "resolution." Figure 9 shows the effects of misregistration on the common summary metric MTF50, the frequency value at which the MTF curve first drops to 50 percent modulation relative to the DC gain.

These plots are related to cross-sections of the plots in Figures 6 and 8 at the 0.5 modulation line. Thus, these plots reveal the same trends as the previous plots showing the full MTF curves, though more obviously quantized: variance in MTF50 value caused by slight misregistration increases with the magnitude of the registration error and is generally less when using the absolute value of the transfer function than the real part.

Conclusions

We have presented a study of the effects that small, random amounts of error in the registration process between an observed image and the ground truth reference image can have on the dead leaves cross-correlation measurement. By simulating random registration errors, both on simulated images processed with a simple linear Gaussian blur and on real images from a Nexus 5X device, we have shown that MTF curves using this method can vary significantly from one random observation to another, even when there is as little as half a pixel of registration error at each corner of the field.

We have also proposed a change to the cross-correlation method of calculating Texture MTF, namely the use of the absolute value of the transfer function as opposed to the real part. We promote this change primarily on mathematical justification and consistency with other accepted definitions of MTF. We suggest that using the absolute value does not discard the advantages of the cross-correlation method (*i.e.*, making use of the phase information) because the advantage actually comes from the use of a full reference instead of a semi-reference ground truth, and the cross-correlation intrinsically makes use of local phase information around the image. The fact that this change also produces results which are much more stable in the face of realworld difficulties such as accurate registration is serendipitous.

Further studies are needed to verify that the absolute value method exhibits the same robustness to various amounts of noise, sharpening, and non-linear processing as its predecessor, but these initial results are promising. Other future work in this area may include a similar analysis of the effect of various amounts of geometric distortion on the MTF calculations.



Figure 9: MTF50 error distributions due to misregistration errors of varying magnitude, represented as boxplots. Each box represents the distribution of error values between the true MTF50 value and the ones derived from MTF curves generated from a misregistered reference image. Each box is centered on the mean of the distribution, with the width including 50% of the values and the black whiskers representing 90%. The red line indicates the median.

References

- ISO 12233:2014 Photography Electronic still picture imaging Resolution and spatial frequency responses, Geneva: ISO, 2014.
- [2] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *Computer Vision. Sixth International Conference on*, 1998.
- [3] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 7, pp. 629-639, 1990.
- [4] A. Buades, B. Coll and J.-M. Morel, "A review of image denoising algorithms, with a new one," *Multiscale Modeling & Simulation*, vol. 4.2, pp. 490-530, 2005.
- [5] K. Dabov, A. Foi, V. Katkovnik and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Transactions on Image Processing*, vol. 16, no. 8, pp. 2080-2095, 2007.
- [6] F. Cao, F. Guichard and H. Hornung, "Dead leaves model for measuring texture quality on a digital camera," in *Digital Photography VI*, SPIE, San Jose, 2010.
- [7] L. Kirk, P. Herzer, U. Artmann and D. Kunz, "Description of texture loss using the dead leaves target: Current issues and a new intrinsic approach," in *Digital Photography X*, SPIE, San Francisco, 2014.
- [8] J. McElvain, S. P. Campbell, J. Miller and E. W. Jin, "Texture-based measurement of spatial frequency response using the dead leaves target: extensions, and application to real camera systems," in *Digital Photography VI*, SPIE, San Jose, 2010.

Author Biography

Robert C Sumner holds a BS in Cinema and Photography from Ithaca College and a MS in Electrical Engineering from Boston University. He joined Imatest LLC as an Imaging Science Engineer in 2016.

Ranga Burada received his Bachelor's degree in Electrical and Communications Engineering in 2011 from JNTU-HYD, India. He received his Master's degree in Electrical Engineering in Fall 2015 from University of Dayton, Ohio. He is currently working as an Imaging Science Engineer at Imatest LLC. His research interests are in developing image quality metrics and image quality analysis.

Noah Kram is a student at the Rochester Institute of Technology, planning to graduate with a BS in Motion Picture Science in the spring of 2017. He was an imaging science engineering intern at Imatest LLC during the summer of 2016.