# Texture-based measurement of spatial frequency response using the dead leaves target: extensions, and application to real camera systems

Jon McElvain[a], Scott P. Campbell[a], Jonathan Miller[a], Elaine W. Jin[b]
[a]Digital Imaging Systems, 130 N. Brand Blvd. Ste. 303, Glendale, CA, USA 91203;
[b]Aptina Imaging Corporation, 3080 North First Street, San Jose, CA, USA 95134

## ABSTRACT

The dead leaves model was recently introduced as a method for measuring the spatial frequency response (SFR) of camera systems. The target consists of a series of overlapping opaque circles with a uniform gray level distribution and radii distributed as $r^{-3}$. Unlike the traditional knife-edge target, the SFR derived from the dead leaves target will be penalized for systems that employ aggressive noise reduction. Initial studies have shown that the dead leaves SFR correlates well with sharpness/texture blur preference, and thus the target can potentially be used as a surrogate for more expensive subjective image quality evaluations. In this paper, the dead leaves target is analyzed for measurement of camera system spatial frequency response. It was determined that the power spectral density (PSD) of the ideal dead leaves target does not exhibit simple power law dependence, and scale invariance is only loosely obeyed. An extension to the ideal dead leaves PSD model is proposed, including a correction term to account for system noise. With this extended model, the SFR of several camera systems with a variety of formats was measured, ranging from 3 to 10 mega-pixels; the effects of handshake motion blur are also analyzed via the dead leaves target.

**Keywords:** spatial frequency response, dead leaves model, camera image quality

## 1. INTRODUCTION

Traditionally, the spatial frequency response (SFR) of camera systems has been measured using the high contrast knife edge target, as detailed in ISO 12233 [1]. More recently, other methods for SFR quantification have been proposed, such as those using the sinusoidal Siemens Star and the log-F contrast targets [2,3]. It is well known that the knife edge approach, and to a lesser degree the Siemens Star method, are susceptible to adaptive sharpening and noise reduction algorithms and can therefore produce results that may not be representative of the perceived sharpness in real images. This is particularly true for the large fraction of images that contain structure with low contrast edges such as facial regions and foliage, where overly aggressive noise reduction may result in an objectionable loss in texture detail. Recently Cao, Guichard, and Hornung [4] described the use of the "dead leaves" target for measurement of system spatial frequency response in the context of texture detail. This new target consists of a series of overlapping circles with a uniform distribution of gray levels, with radii (r) distributed approximately as $1/r^3$ [5]. This power law distribution suggests near scale-invariance of the target, so that in principle actions such as scaling, cropping, or rotation should not change the power spectral density (PSD). With the knowledge of the "ideal" PSD of the dead leaves distribution, it is possible to compute the SFR of the camera system simply by computing the 2-D Fourier transform and comparing its radial PSD to that of the ideal target. The proposed dead leaves target is part of an ongoing investigation by members of the Camera Phone Image Quality (CPIQ) Initiative, sponsored by the International Imaging Industry Association (I3A), in developing objective metrics to predict texture blur in real digital imaging systems [6,7].

For the purposes of image quality evaluation, the dead leaves target is of particular interest because it uses an occlusion model that is representative of overlapping objects (foliage, etc.) that are commonly found in natural scenery [5]. Phillips, et al [8] performed subjective evaluations on a series of images subjected to sigma filtering. The sigma filters were designed to have different levels of kernel size and strength, and had the visual effect of removing textures from the test images. A softcopy ruler method was used to measure the subjective impressions of these images to the viewers [9]. The results of the preference study correlated well with the results of the dead leaves SFR. Aggressive (and perhaps nonlinear) noise suppression can result in significant reduction in texture detail, and the SFR in the context of the dead leaves target will be penalized. In contrast, for the knife-edge approach, the effects of such texture loss will be

Digital Photography VI, edited by Francisco Imai, Nitin Sampat, Feng Xiao, Proc. of SPIE-IS&T Electronic Imaging,
SPIE Vol. 7537, 75370D · © 2010 SPIE-IS&T · CCC code: 0277-786X/10/$18 · doi: 10.1117/12.838698

minimized and therefore the correlation with preference is greatly reduced. In this respect the dead leaves target is of potential high value, as this relatively straightforward objective measurement could be used in place of a much more time consuming subjective image quality evaluation.

In this paper we analyze the use of the dead leaves target for measurement of camera system spatial frequency response. We show that the PSD of the ideal dead leaves target does not exhibit a simple power law dependence over the entire spatial frequency range, and furthermore we find that the ideal PSD is not truly scale invariant. To force consistency of SFR measurements across different scales, we propose an extension of the ideal dead leaves PSD model that includes a correction for captured image size. Additionally, we discuss the contribution of noise introduced by the camera system, and how it impacts the computed SFR results. A further extension is proposed to approximate (via measurement) and subtract the PSD contributions resulting from camera system noise, and we show that the computed SFR produce more reasonable behavior at higher spatial frequencies. With these extensions, we were able to compare the SFR of camera systems with a variety of formats, ranging from 3 to 10 mega-pixels, and we also use the dead leaves target to analyze the effects of handshake motion blur.

## 2. DEAD LEAVES TARGET GENERATION AND PRINTING

The procedure for generating a list of circles for the dead leaves target is fairly straightforward, and is in part outlined in [5]. It is based on an occlusion model, with circles generated with a uniform distribution in gray level in the range [0.25,0.75]. The circle radii are chosen according to a $1/r^3$ probability distribution (uniform in $1/r^2$), although there must be an upper and lower bound for the radii to avoid full coverage by the smallest or largest circles, as discussed in [5]. Typically, $r_{min}$ is chosen such that the size of the smallest circles will be much less than the image pixels when the target is projected onto the camera sensor array; $r_{max}$ is selected to be much larger than $r_{min}$, but generally not larger than the width of the image, W. The center of each circle is randomly placed on a square canvas of width $W+2r_{max}$; during the actual rendering (printing) phase it is assumed that only the center WxW area is reproduced. Circles whose entire area falls outside of the inner WxW image are obviously excluded from the final list.

The occlusion model assumes that each object is opaque, so that each circle in the list will overwrite any pixels from other circles that are below its local area on the rendered image. The most obvious procedure ("back to front") for generating a list of circles involves a loop that creates a Z-stack of circles (bottom first) until the entire image area WxW is filled. This can be accomplished by creating a high resolution image map (e.g. 32768x32768 pixels, initializing to zero) to subdivide the WxW image area, and then sequentially rendering each circle onto this pixel map; a list of circles is created with each new circle added to the end of the list. Following each circle (or after each 1000 circles or so), the image map is checked for any null entries, and the loop is exited once all pixels have been touched by at least one circle in the list. In the back-to-front approach, a relatively large list is created with an $r^{-3}$ distribution, and can exceed 7-8 million entries for large $r_{max}/r_{min}$ ratios. However, many of the circles created early in the process become fully occluded, and so the list may be pruned to a significantly lower number of entries, perhaps in the range of 500,000 - 700,000. Since smaller circles have a higher probability of becoming fully occluded, the final distribution of radii will change to some other power law, e.g. $r^{-2.7}$.

An alternative method for generation of the circles list is the "front-to-back" approach mentioned by Lee [5]. Instead, the topmost circles are generated first, and subsequent circles will be added to the list depending on whether or not the entire area of the image map pertaining to the circle has already been populated. With this method, each new valid circle is inserted at the beginning of the list (the topmost list elements will be rendered first), and like the back-to-front method, the loop continues until every pixel in the image map has been rendered by at least one circle in the list. By virtue of the in-loop exclusion, the final list will have already been pruned of fully occluded circles. As the list grows, smaller candidate circles will have a lower probability of becoming "valid", so the final distribution will be changed. In fact, the final distribution will be equivalent to the back-to-front method assuming the same $r_{min}$ and $r_{max}$. In both cases, approximately the same number of iterations is required to achieve a complete circles list. However, using the front-to-back method does not require the additional pruning step, but does necessitate more frequent checks of full image coverage. For the analysis that follows in Sections 3 and 4, we employed the front-to-back method for our circles list generation, an example of which is shown in Figure 1.

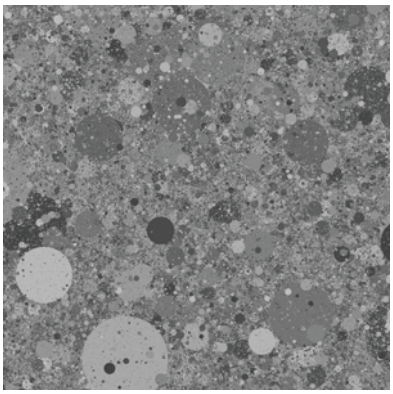

Figure 1. Example dead leaves target generated using a $r^{-3}$ distribution of radii and a uniform distribution of gray scales in the range [0.25,0.75].

Once the list of circles has been created, it is necessary to place the target description in a format that is consumable by a printer rendering engine. One possibility is rendering the target at a fixed resolution in a format such as TIFF, either at 8 bits/pixel or preferably 16 bits/pixel. However, this approach can suffer from resampling artifacts due to differences in resolution between the original TIFF and the printer physical raster. A more scalable and compact method is to instead represent the list of circles as abstract vector objects in a language such as Postscript or PCL. In this fashion, the printer raster image processor (RIP) will optimally render each of the circles onto the native raster of the printer without resampling. For our purposes, we created the printer-ready target in the form of an encapsulated Postscript (EPS). This format allows for the circles list to be easily imported into page layout applications such as Adobe Illustrator or Freehand, and it is straightforward to hand-edit Postscript code if desired. Also important to include is a step wedge intended to derive the opto-electronic conversion function (OECF). The OECF patches can ramp in any fashion, as long as the luminance values (relative to paper white) are known for each patch. In our case, we chose a linear OECF series, since we are most interested in linearizing the midtones in the captured images.

Care must be taken when printing the dead leaves target, as it is important to ensure a linear gray scale output with respect to paper-white luminance. As mentioned previously, the gray levels of the circles list are confined to the range [0.25,0.75], so that there will be no issues with respect to limitations imposed by $D_{max}$ of the printer. However, the printer rendering engine must correctly convert the linear gray color space of the incoming dead leaves image to an output with the same linearity in luminance. In practice, most printer color management systems do not provide the level of control over the color rendering to achieve this goal. For standard ICC color paths, often various forms of gamma correction and tone corrections are applied to the gray levels so that the output may not be completely linear. Likewise, for some printer systems it may be possible to linearize, but only between $D_{max}$ and paper white. Therefore, it is necessary to have a RIP/printer that provides the low level color rendering control required for this task. In our case we used a Serendipity Blackmagic RIP to drive an HP Z3100, which produced the desired linear output.

## 3. TARGET ANALYSIS

### 3.1 Scale-invariant analysis

According to Cao, et al [4] the power spectral density of the dead leaves target can be approximated by a power law dependence, represented for an LxL image as follows:

$$\left|\hat{U}_{mn}\right|^2 = \begin{cases} L^4\langle u\rangle^2 & m=0, n=0 \\ \dfrac{A(L)}{\left(\sqrt{m^2+n^2}/L\right)^p} & otherwise \end{cases} \tag{1}$$

where p is the empirical power law coefficient, $\langle u \rangle$ is the average digital level of the image, and $A(L)$ represents the normalization coefficient, which is dependent only on the image size L. It is assumed that the Fourier spectrum is shifted such that (m,n)=(0,0) appear at the center of the transformed image. $A(L)$ can be related to the image variance by employing Parseval's theorem,

$$\mathrm{var}(u) = \frac{1}{L^2}\sum_{i,j}\left|u_{ij}\right|^2 - \langle u \rangle^2 = \frac{1}{L^4}\sum_{m,n}\left|\hat{U}_{mn}\right|^2 - \langle u \rangle^2 \tag{2}$$

or

$$A(L) = \frac{L^4\,\mathrm{var}(u)}{\sum_{m,n}\frac{1}{\left(\sqrt{m^2+n^2}/L\right)^p}}, \quad m^2+n^2 \neq 0 \tag{3}$$

As reported in [4], the power law coefficient was measured as p≈1.857; assuming the gray levels are scaled to lie within [64,196] out of 255 (i.e. [0.25,0.75]*255), the normalization coefficient can be fairly well approximated by the following expression:

$$A(L) = 71.0156 * L^{1.8905} \tag{4}$$

The spatial frequency response of the camera system is then computed by 1) calculating the (centered) Fourier transform of the captured image and subsequently the power spectral density (PSD); 2) dividing the measured PSD by the ideal PSD from Equation 1 for image size L; 3) taking the square root of the result from (2) and averaging over all orientations within radial spatial frequency bins out to the square bounded by (±L/2,0) and (0,±L/2). The spatial frequency of the outer bin will be normalized to 0.5 cy/pix, and thus a list of SFR versus spatial frequency will be created. Alternatively, in step (2) above, the measured PSD can be first averaged over angle within spatial frequency rings [f, f+Δf] and divided by the appropriately averaged ideal PSD. Using the dead leaves method, the SFR accuracy at lower spatial frequencies is reduced since there are fewer image samples over which the averaging takes place. In Cao [4] it is suggested that traditional knife-edge SFR measurements according to ISO 12233 can be used to populate the low spatial frequencies, and perhaps to re-normalize the SFR curve as well.

### 3.2 Extensions

We investigated several dead circles lists generated according to the procedure described in Section 2. Each list was created with uniformly distributed (double precision) gray levels in the range [0.25,0.75]*255; $r_{min}$=(1/4096)*N and $r_{max}$ =497*$r_{min}$, where N was the size of the rendered image. We rendered the circles to create a large initial image of dimensions 32768 x 32768, so that the radius of the smallest circle was equal to 8 pixels. This large image was used to reduce pixelization artifacts and simulate the dead circles target as if it were printed on a calibrated high-resolution output device.

Now consider capturing this printed target with a digital camera system. For practical purposes, the entire target should not consume more than about 30% of the horizontal field of view, in order to reduce effects such as variations in lens point spread function, lens shading, aberrations, pixel cross talk, etc. Therefore, for sensors below 30 megapixels, captured images will posses a much lower resolution relative to the image generated by the printing device, with capture image sizes ranging from 128 x 128 up to 2048 x 2048 pixels. Consider further the case of an ideal camera system, with an infinitely sharp lens point spread function (PSF) and constant response across a given sensor pixel. In this case the sensor will average the higher resolution printer image across the individual pixels (where the smallest circles may have a radius equal to 1/8 of a sensor pixel, or smaller), and subsequently will create the captured image with a smaller pixel count. Introducing anything other than a delta function PSF or a constant intra-pixel response will necessarily blur this image further. To simulate this ideal capture process for our purposes, we applied an *s* by *s* box filter to the large "printer" image, and then subsampled by s to create a (N/s x N/s) image. In principle, this produces "ideal" images in the context of the capture system, and is consistent with the processing approach described in [5].

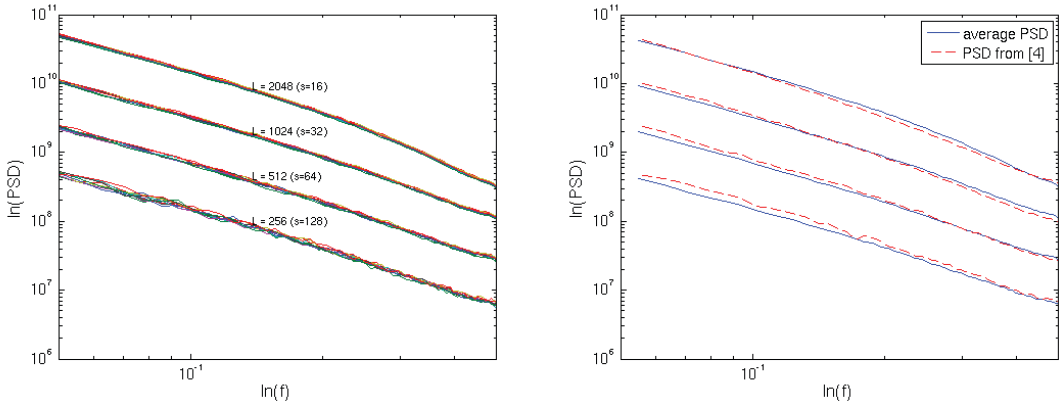


a) b)

Figure 2. a) Dead leaves PSD vs. spatial frequency for different scales for 10 different lists generated according to the procedure of Section 2. Scale invariance is only loosely obeyed, and each curve shows a fair degree of curvature; b) the average PSD vs. f curves for the 10 lists generated, in comparison to the PSD curves generated by a circles list provided by the authors of [4].

Next, we computed the DFT of these ideal capture images, and subsequently the radially binned power spectral density was calculated according to the procedure described in Section 3.1. Figure 2a shows a log-log plot of the PSD vs. spatial frequency for different values of s (10 different circle lists), and it can be seen that the ideal camera system PSD does *not* exhibit simple power law behavior that is independent of the scale factor *s* across the entire spatial frequency range. In fact, for larger "ideal" capture images (smaller s), there is a substantial curvature in these log-log plots. Assuming a simple power law functional form would therefore result in significant error when computing the system spatial frequency response. The deviation from a simple power law is attributed to the bounding of the radii in the distribution, as discussed in [4,5]. It is important to note that we also analyzed a list provided by the authors of [4] (final radius distribution of approximately $r^{-2.5}$ versus our distribution of $r^{-2.75}$), and a departure from scale-invariance is also observed in that case, as shown in Figure 2b. Although, as expected the PSD vs. *f* curves behave slightly differently due the change in distribution of radii.

Considering the substantial departure from a simple power law, we propose modification to the dead circles ideal PSD model to include a quadratic dependence on $\ln(f)$. The coefficients fundamentally depend on the captured image width (L=N/s), and can be empirically determined at discrete capture image sizes in $2^k$ steps, e.g. 128x128 up to 2048x2048. The modified model for the PSD is obtained replacing *p* with $B(s)+C(s)\ln(f)$,

$$\ln\left[PSD_{ideal}(f)\right] = \ln\left[A(L)\right] - B(s)\ln(f) - C(s)\left[\ln(f)\right]^2 \tag{5}$$

Using the 10 generated lists of circles, the average value of the coefficients A, B and C were determined for different scale factors (Table 1). Note that these coefficients are dependent on the specific choice of $r_{min}$, $r_{max}$, as well as the final distribution of radii resulting from the circles generation procedure. Nevertheless, if the circles generation procedure and selection of min/max radii were to be standardized, then these coefficients would be valid for computation of the spatial frequency response of all camera systems in the context of the dead leaves target.

It must be noted that the PSD function is approximately invariant with respect to cropping (other than amplitude scaling by $1/c^4$, where *c* is the crop factor), as shown in Figure 3. This is due to the construction of the target, where the circles are placed in a random fashion spatially; the shape of the curve is therefore still governed by the scale factor *s*. In this case, Equation 5 and Table 1 can still be used, although the relationship between *L* and *N* is modified, with $L = N/(cs)$.

Table 1. PSD extended model coefficients for the dead leaves target.

| L=N/s | ln[A(L)] | B(s) | C(s) |
|---|---|---|---|
| 128 | 12.531 | 2.295 | 0.09991 |
| 256 | 13.985 | 2.400 | 0.12613 |
| 512 | 15.476 | 2.407 | 0.12067 |
| 1024 | 16.690 | 2.718 | 0.19723 |
| 2048 | 17.253 | 3.601 | 0.39951 |

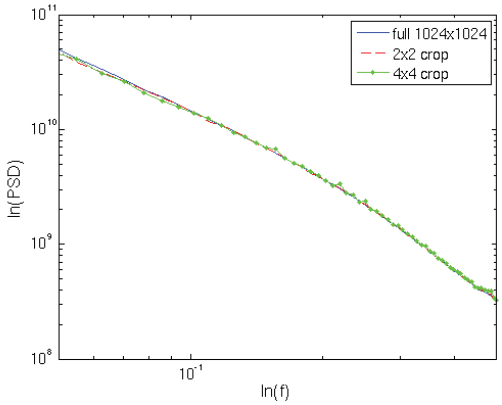


Figure 3. Effect of cropping on the dead leaves PSD curve. Clearly the target is robust to cropping due to its random spatial distribution.

In principle, computation of the spatial frequency response of the dead leaves target is straightforward, and simply involves the calculation of the measured PSD for the particular camera system. Dividing this by the ideal PSD (and taking the square root) results in the spatial frequency response.

However, for real camera systems, sensor noise and JPEG artifacts will corrupt captured images. This is problematic in the context of the dead leaves target, as the target-based PSD falls off very rapidly, and in many cases the noise/artifact PSD will dominate at higher spatial frequencies. Therefore, these artifacts can result in an artificial enhancement of the SFR, and any metric derived from this may be distorted by noise. For this exact reason, the CPIQ technical group has proposed that texture loss should be considered together with noise performance and edge sharpness [6]. The proposal included capturing both a slanted edge target and a dead leaves target simultaneously. The proposed objective metric of texture loss can then be calculated by combining the measured SFR from the slanted edge and the dead leaves target. However, in the present study an alternative method is explored.

If somehow the noise/artifact PSD could be computed for the image, then this quantity could be subtracted from $PSD_{meas}$ to create a more accurate estimation of the SFR,

$$SFR(f) = \sqrt{\frac{PSD_{meas}(f) - PSD_{noise}(f)}{PSD_{ideal}(f)}} \tag{6}$$

Although $PSD_{noise}$ is difficult to extract from the original image, it can be approximated by computing the PSD of a *uniform* patch with a paper-white relative luminance equal to the mean of the dead leaves target (50%). The size of the 50% patch can be equal to the dead leaves target (LxL), or smaller patch can be used that can be subsequently tiled to create an LxL image. Figure 4a shows the dead leaves PSD for an image captured at different light levels, and $PSD_{noise}$ for the corresponding 50% patch using the same camera system and exposure conditions. Clearly, the high frequency PSD of the dead leaves image shows nearly identical behavior to that of the 50% patch PSD for the lowest light level. Figure 4b shows the resulting SFR calculations, and it can be seen that the noise-corrected curve effectively suppresses the effects of the high frequency artifacts introduced by the camera system. For the camera system comparisons that follow, we use Equation 6 to compute the noise-corrected dead leaves SFR.

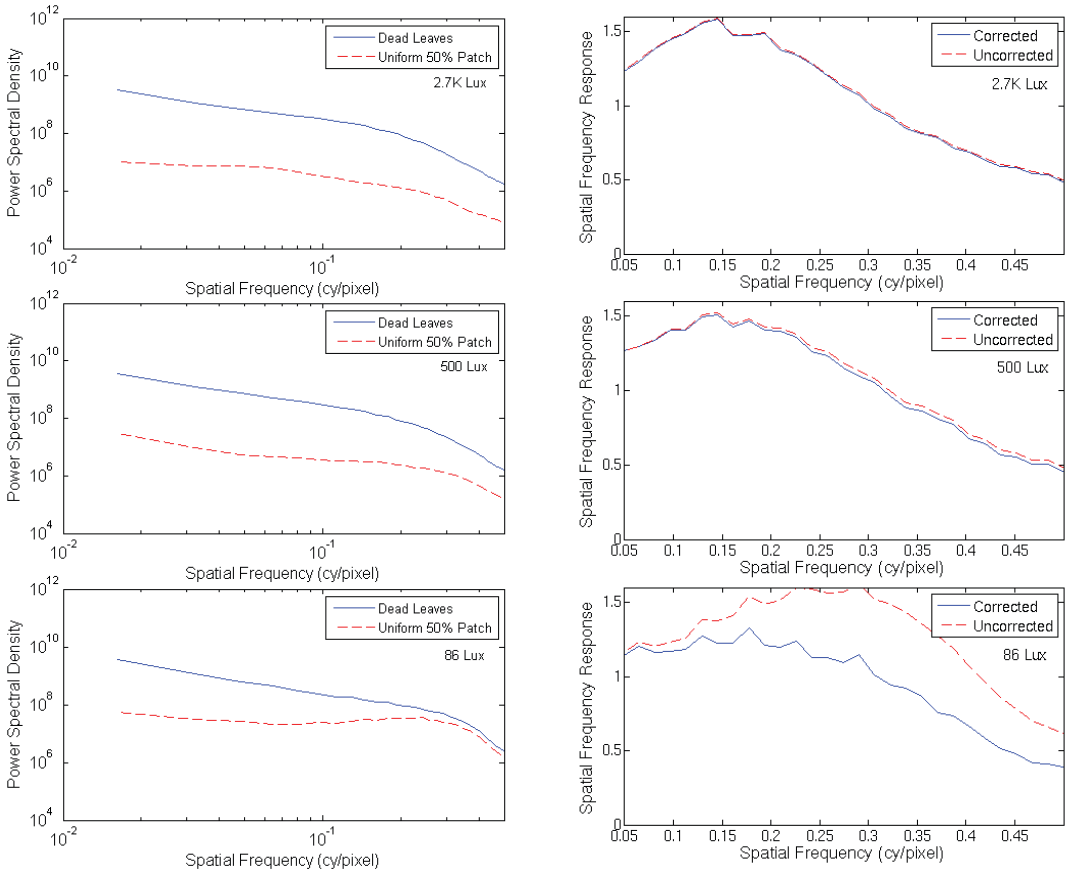


a) b)

Figure 4. a) Dead leaves PSD for a real camera system (solid blue), and the PSD of a 50% uniform patch (dashed red), for different target illumination levels (2.7K lux, 500 lux, 86 lux). For the lowest light level, the dead leaves PSD is dominated by the system noise at higher spatial frequencies; b) computed SFR, with (solid blue) and without (dashed red) noise correction, according to Equation 6.

## 4. CAMERA COMPARISONS USING THE DEAD LEAVES TARGET

### 4.1 Stationary capture

Armed with the formalism of the previous section, we used several camera systems to capture a dead leaves target. The cameras used were as follows: Canon SD1100 IS, Casio Elixim EX-S10, Motorola ZN5, Palm Pre, HTC (Google) G1, Apple iPhone 3GS, and the Samsung Memoir. Most of these systems had auto-focus capability, with the exception of the Palm Pre, which used an extended depth of field (EDOF) approach. In addition to auto focus, the Canon SD1100IS also featured optical image stabilization. Table 2 summarizes the features of each of the camera systems tested.

The dead leaves target was printed on the black channel of a HP Z3100 through a Serendipity Blackmagic RIP, and HP Super Heavyweight Matte paper was used as the substrate. Printer calibration was performed using an X-Rite i1 iSis, and resulted in an output such that (1-$k$) was linear in luminance, where $k$ was the digital level sent to the RIP. Because the system was linearized in luminance relative to paper white, digital values above $k$=0.98 were forced to $D_{max}$ of the printer. However, this clipping of the calibration curve in the shadows was of no consequence to the dead leaves target, since the gray levels of the circles lie in the range [0.25,0.75]. The target itself was 8.1" x 8.1" and positioned over a 45 inch 18% gray circle. A 21-step OECF wedge (linear in luminance relative to paper white) was placed adjacent to the dead leaves target for linearization purposes during analysis. Finally, fiducial markers were placed near the corners of the dead leaves/OECF portion for alignment and automatic detection purposes. To accommodate estimation of the noise PSD, a separate chart was created whereby the dead leaves target was replaced by a 50% patch.

Table 2. Cameras compared and features.

| Camera | Class | Sensor Size | Pixel Size | f/# | Optical Features |
|---|---|---|---|---|---|
| Canon SD1100 IS | DSC | 8 MP | 1.75 µm | 2.8-4.9 | 3x zoom, auto focus, image stabilization |
| Casio Elixim EX-S10 | DSC | 10 MP | 1.67 µm | 2.8-5.3 | 3x zoom, auto focus |
| Motorola ZN5 | Mobile | 5 MP | 2.2 µm | 2.8 | auto focus |
| Palm Pre | Mobile | 3.1 MP | 1.75 µm | 2.4 | EDOF |
| HTC G1 | Mobile | 3.1 MP | 1.75 µm | 2.8 | auto focus |
| Apple iPhone 3GS | Mobile | 3.1 MP | 1.75 µm | 2.8 | auto focus |
| Samsung Memoir | Mobile | 8 MP | 1.75 µm | 2.8 | auto focus |

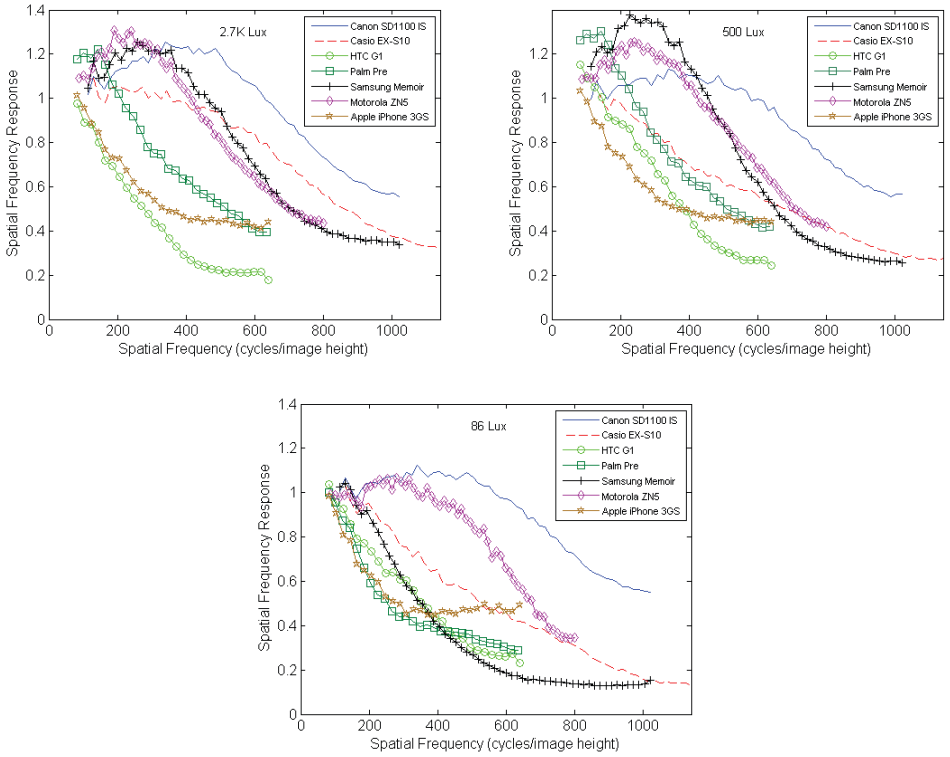


Figure 5. Spatial frequency response of the dead leaves target for the various camera systems measured, with the cameras mounted on a tripod: 2.7K lux (top left); 500 lux (top right); 86 lux (bottom).

The test chart was illuminated using a pair of Kinoflo Image 85 (5500 K) lights at three different levels (2700 lux, 500 lux, 86 lux). During the capture process, each camera was positioned such that the radius of the background circle on the chart intersected the corners of the sensor. In this way, the width of the dead leaves target in the final capture was equal to approximately 0.36 image height, or 22.5% of the horizontal field of view (for 4:3 aspect ratio sensors). As the HFOV was approximately equal for the cameras tested, the distance from the cameras to the test chart was approximately 90 cm. It is important to note that the flash was disabled for those cameras equipped with that feature.

The captured JPEG images were then analyzed, first by detecting the fiducial markers. The OECF patches were extracted from the image and measured to create an 8-bit lookup table relating the captured digital level to a luminance level in the range of 0-255. The dead leaves portion of the image was cropped down to the closest $2^k$ x $2^k$ sub-image. The OECF lookup table was then applied to the dead leaves sub-image, and the DFT was taken (using the FFTW [10] library). From the DFT, the PSD was computed, and subsequently the averaged PSD vs. spatial frequency $f$ in the range [0.05, 0.5]. Likewise, the same procedure was performed on the 50% gray "noise estimation" patch. The spatial frequency response was then computed by dividing the measured (noise-corrected) PSD by the ideal PSD (using Equation 6 and table 2 for $L=2^k$), and taking the square root at each spatial frequency.

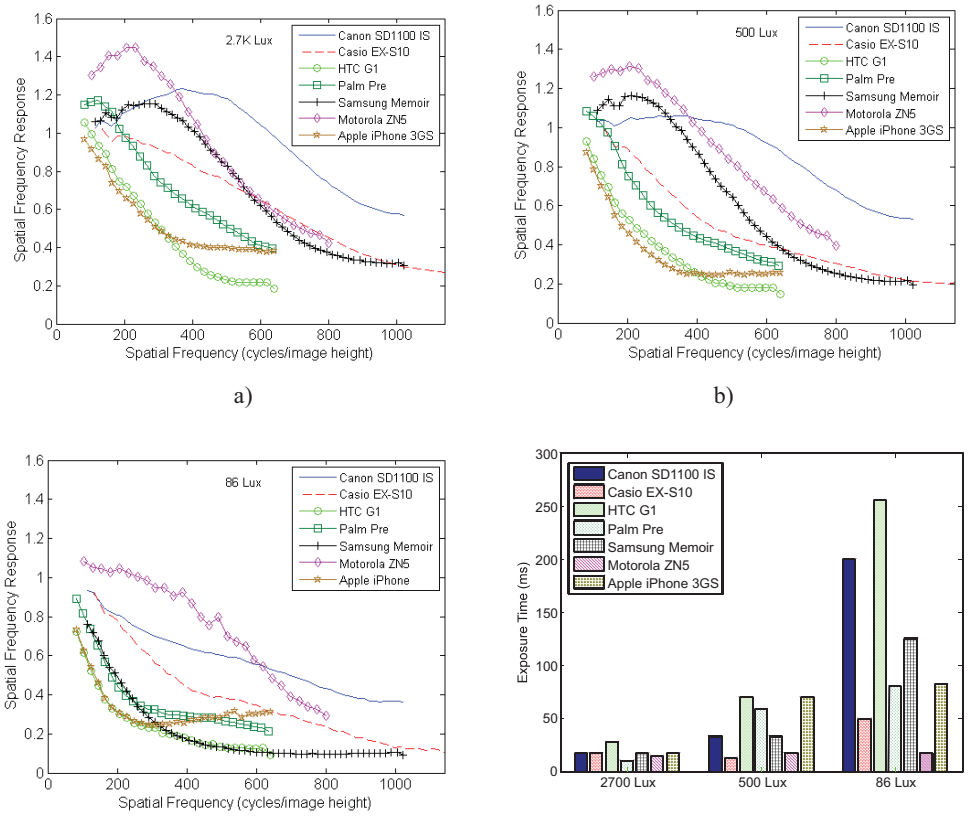


Figure 6. Spatial frequency response of the dead leaves target for the various camera systems measured, including the effects of motion blur for: a) 2.7K lux; b) 500 lux; c) 86 lux. Clearly, the systems with longer exposure times perform significantly worse than when mounted on a tripod, with the exception of the image stabilization-equipped Canon SD1100.

Shown in Figure 5 are the dead leaves SFR results for the various tripod-mounted cameras, at 3 different illumination levels. Note that we chose to represent spatial frequencies in cycles/image height, due to the wide range of pixel counts among the cameras. At 2700 lux, the Motorola ZN5, Canon SD1100 IS, and the Samsung Memoir all show a significant amount of sharpening that pushes their response above unity. The Casio Elixim, HTC G1, Palm Pre, and the iPhone 3G all demonstrate little if any sharpening in their image processing. At 500 lux, sharpening appears to be reduced for the Canon, but maintained for the ZN5 and Memoir. At the lowest light level (86 lux), the Memoir appears to completely disable sharpening, whereas the sharpening for the ZN5 and the Canon remains significant.

It is interesting to note that the SFR for these systems does not completely fall to zero as would be expected in the case of the knife-edge SFR. This characteristic exists for all of the cameras - even at the brightest scene illumination levels, where noise contributes an insignificant amount to the PSD. Although the cause has not been fully characterized, it is believed that this effect is created by edge enhancements associated with sharpening, JPEG compression, and possibly pixel cross-talk. This may be fundamental to the construction of the dead leaves target, where a large percentage of the image consists of edges. Considering only sharpening, its contribution is partially supported by the Samsung Memoir SFR curves in Figure 5 - for the lower light levels where sharpening is reduced or disabled, the high frequency offset is reduced.

### 4.2 Handshake blur measurements

For consumer-level digital still cameras, and to a greater degree mobile cameras, it is very rare that images are captured on a stable surface such as a tripod. Instead most images are captured with the user holding the camera, and quite often there is some amount of hand motion during the exposure time. The characteristics of this motion depend on many design factors associated with the camera, including orientation, hand position(s), location of the capture button, and the

"click-to-capture" delay. A common rule of thumb for the amplitude of this motion is 5 degrees/sec [11,12,13], although the motion can be translational or rotational (or a combination thereof) depending on design. In the final consumer image, the sharpness quality will be determined by both the system optics/image processing *and* the motion during capture. For most systems, exposure time factors very strongly in this domain - a short exposure time will minimize motion blur artifacts whereas longer exposure times can significantly increase the image blur. However, in lower lighting the short exposure times can lead to objectionable levels of noise, and can require image filtering (added blur) to reduce this noise. These system-level SFR contributions are discussed in detail by Xiao, et al [11]. Furthermore, for elevated noise levels, the color saturation embedded in the system color correction matrix (CCM) must be reduced due to the generally high condition number for these matrices.

In order to assess the effects of handshake motion blur, the same camera systems were used to capture a modified test chart. The dead leaves target was bordered on the left and right with facial images, and were placed in order to activate the optical image stabilization for cameras with such capabilities. For all of the cameras, seven different individuals captured 4 images of the test chart, for each of the three illumination levels. The dead leaves SFR was measured in the same way as described in the previous section, except that a single OECF lookup table was used for all images with the same camera and illumination level. For each camera and illumination level, the *SFR* vs. *f* result was obtained by averaging over 4x7=28 curves.

Figure 6 shows the results of these SFR calculations, and it is clear that the cameras behave differently as a function of light level. To provide more clarity on these differences, we also measured the exposure times for each of the cameras at the different light levels, shown in Figure 6d. At the very brightest light level, each of the cameras chooses a relatively short exposure time, so the motion blur artifacts are minimal (compare to the curves in Figure 5). However, as the illumination level is reduced, significant differences are observed. With the exception of the ZN5 and the Casio Elixim, all cameras choose longer exposure times. Of these cameras, all exhibit significant degradation in the dead leaves SFR, although the Canon (equipped with optical image stabilization) shows the least amount of performance reduction. Due to its apparent 16 ms exposure time ceiling, the ZN5 shows the least amount of motion blur degradation at low light levels. Although, at these light levels, the short exposure times require a very high ISO speed (gain) resulting in objectionably noisy images. Because of the slightly longer exposure time at low light levels, the Casio shows greater motion blur degradation relative to the ZN5, although the noise levels are not as high.

These results highlight important system trades from the standpoint of camera design. In the context of sharpness/texture blur, a camera with very good combination of optics and sensor can actually perform poorly in the presence of handshake motion if longer exposure times are chosen. Conversely, forcing a short exposure time may result in a superior texture-based SFR, but the noise levels may be excessive. The camera system image processing pipeline may be configured for noise reduction in these cases, but that would necessarily result in a reduced response to the textured dead leaves target. From the user preference standpoint, there needs to be a fundamental understanding of these trades as they pertain to the market segment for which the camera is intended.

## 5. CONCLUSIONS

Use of the dead leaves target for spatial frequency response estimation was analyzed. Key details associated with generation of the list of circles, as well as printing considerations were discussed. By analyzing several dead leaves distributions, it was found that the ideal PSD does not exhibit a simple power law dependence on spatial frequency as previously reported. Instead, modeling ln(PSD) as a quadratic function of ln(f) resulted in a much better agreement across the entire spatial frequency range. Furthermore, it was shown that the PSD model coefficients are only loosely scale-invariant, so that they must be empirically determined for a given scale $s$ and capture size $L$. Because the dead leaves target is spatially distributed, it exhibits a fundamental sensitivity to noise and other artifacts that are introduced by the camera system. The PSD of these artifacts can dominate the measured PSD at high spatial frequencies, and result in artificial enhancement of the computed spatial frequency response. To circumvent this problem, additional measurement of a 50% uniform patch was proposed, and the subsequent computation of the spectral density was used to compensate for the artifact-related portion of the measured PSD. With this modification, it was shown that the resulting SFR curves are well-behaved and consistent with expectations.

The dead leaves target was also used to measure the effects of motion blur, and it was particularly well suited for this endeavor due to its spatially distributed nature. Camera systems with long exposure times may perform reasonably well in a statically mounted (tripod) configuration. However, in the presence of real handshake motion, these cameras

perform substantially worse, ultimately resulting in captured images with an excessive amount of blur. This issue becomes particularly problematic in low light situations, and from the system level the choice must be made between blur artifacts caused by 1) motion (for longer exposure times), 2) spatial averaging (for short exposure times), or 3) higher noise levels if spatial averaging is not applied in conjunction with short exposure times. However, such artifacts are mitigated in cameras with optical image stabilization, since much longer exposure times can be afforded without a large penalty of handshake blur artifacts.

## REFERENCES


[1] ISO 12233, Photography - Electronic still picture cameras - Resolution and spatial frequency response measurements, 2000.

[2] Loebich, C., Wueller, D., Klingen, B., and Jaeger, A., "Digital camera resolution measurement using sinusoidal Siemens star," Proc. SPIE, Vol. 6502, 65020N (2007).

[3] http://www.imatest.com/docs/log_f.html.

[4] Cao, F., Guichard, F., Hornung, H., "Measuring texture sharpen of a digital camera," Proc. SPIE, Vol. 7250, 72500H (2009).

[5] Lee, A.B., Mumford, D., Huang, J., "Occlusion Models for Natural Images: A Statistical Study of a Scale-Invariant Dead Leaves Model," IJCV 41 (1/2), 35 (2001).

[6] Phillips, J. B. and Christoffel, D. W., "Validating a texture metric for camera phone images using a texture-based softcopy quality ruler," Proc. SPIE, Vol. 7529, (2010).

[7] Cao, F., Guichard, F., and Hornung, H., "Dead leaves model for measuring texture quality on a digital camera", Proc. SPIE, Vol. 7537, (2010).

[8] Phillips, J., Jin, E., Chen, Y., Clark, J., "Correlating Objective and Subjective Evaluation of Texture Appearance with Applications to Camera Phone Imaging," Proc. SPIE, Vol. 7242, 724207 (2009).

[9] Jin, E. W., Keelan, B. W., Chen, J., Phillips, J. B., and Chen, Y., "Softcopy quality ruler method: Implementation and validation," Proc. SPIE, Vol. 7242, 724206 (2009)

[10] http://www.fftw.org/.

[11] Xiao, F., Farrell, J.E., Catrysse, P.B., Wandell, B., "Mobile Imaging: The Big Challenge of the Small Pixel", Proc. SPIE Vol. 7250, 72500K (2009).

[12] Xiao, F., Silverstein, A. and Farrell, J., "Camera motion and effective spatial resolution," ICIS'06, 33 (2006).

[13] Xiao, F., Pincenti, J., John, G. and Johnson, K., "Camera-motion and mobile imaging," Proc. SPIE Vol. 6502, 650204 (2007).