Limits of Resolution. 6. How Many Megapixels?

Phil Service
Flagstaff, Arizona, USA
8 May 2015

Summary

Photosite spacing of 1.5 µm or less is common for smart phone cameras; and 1-inch sensors in cameras such as the the Sony RX100 III have photosite spacing of 2.4 µm. Diffraction-limited line-pair resolution is given for photosite spacing as little as 0.5 µm. Two micron photosite spacing implies APS-C and "full-frame" (FF) sensors with 94 and 216 MP, respectively. In order to approach the theoretical resolution limits of sensors with 2 – 3 µm photosites, it will be necessary to have lenses that perform exceptionally well at apertures f/2.8 – f/4. It is not clear if such lenses can be manufactured at reasonable cost for APS-C and FF sensors. If it is, it may be necessary to sacrifice large maximum apertures, such as f/1.4, in order to make "slower" but sharper lenses. With current technology, 2 – 3 µm photosites will entail a trade-off between resolution and low noise, when compared to current FF and APS-C sensors. For most image uses, 100 MP or greater resolution implies *capture oversampling.* That is, images will be down-sampled for "final" use. It is suggested that such down-sampling may produce a sharper and less noisy final image than could otherwise be obtained by capturing images at lower initial resolution.

Key words: resolution limits, sampling frequency, diffraction, sensor pixel size, photosite size, perfect lens, line-pairs, simulation, oversampling, sharpness, signal-to-noise ratio, down-sampling, image noise, edge acutance

1. Introduction

All other things being equal, closer photosite spacing means higher image resolution (lp/mm).[1] On the other hand, diffraction degrades resolution. If we could pack photosites as closely as we wanted, at what point would diffraction limit resolution? The answer to that question arguably sets the useful upper limit to "megapixels" for a sensor of a given format.[2]

---

[1] I will use the terms "photosite spacing" and "photosite size" interchangeably, with the understanding that the former is more appropriate in the present context, which focuses on sampling frequency. Generally photosite "size" is smaller than "spacing" or "pitch" because gaps must be left between adjacent photosites.

[2] At the risk of being eccentric, or worse, I use the term "photosite" when referring to an individual light receptor on a camera sensor. A "pixel" is the smallest element of a digital image. However, given the apparently universal practice of describing camera sensor resolution in terms of megapixels, it would seem perverse to insist on using "megaphotosites" instead. Hence "megapixels" in the title of this paper. Happlily, however, the usual abbreviation for megapixels, MP, also works for megaphotosites. So, whenever you see the abbreviation MP, feel free to say megaphotosites in your head.

There is more to image sharpness than resolution, *per* se. Most images contain details with a large variety of spatial frequencies. Images appear "sharp" not only when high-frequency detail is visible, but also when lower frequency detail is rendered with crisp edges. Diffraction unavoidably makes edges fuzzy, to a degree dependent on aperture, and sets a minimum width for fuzzy edges. "Sampling" of edges by photosites cannot decrease the width of the fuzzy zone below the minimum set by diffraction. However, low-frequency sampling — by large, widely-spaced photosites — will *increase* the width of the zone. Thus, a second question: at what point does packing photosites more closely together stop yielding useful gains in edge acutance?[3]

The sensor of the iPhone 6 camera has 1.5 µm photosite spacing. Some other smartphone sensors have even smaller spacing, and sensors with photosites smaller than 1 µm are being considered.[4] A "full-frame", 36 x 24 mm, sensor with 1.5 µm photosites would have 384 MP. Whether we will ever see such sensors in consumer products, I cannot say. However, it is probably safe to say that the 1 – 1.5 µm photosite will eventually find its way to sensors larger than those used in current smartphone cameras.

1.1. Terminology

As in the previous papers, it will be useful to define a few terms at the outset. The *lens image* is the image that is formed by the lens. The lens image, or *image field,* is an effectively continuous, analog representation of the external world — the *object field* — in front of the lens. The *sensor* image is the digitized image recorded by the sensor. It is absolutely crucial to understand that this is a *sample* of the image field. Blur affects the lens image, *not* the sampling process. Therefore, there is not a simple relationship between diffraction blur and photosite pitch — a fact that seems generally to be misunderstood. *Resolution* in the present context means line-pairs per millimeter (lp/mm). In general, a *contrast ratio* will be associated with a resolution measure. The contrast ratio is the maximum difference in lightness values of alternating light and dark lines, divided by their sum. A *perfect lens* is a lens with no optical aberrations. In the absence of diffraction, it would produce an image with no blur. A perfect lens is assumed in everything that follows, and diffraction is the only source of lens image blur. *Nyquist rate resolution* is the maximum resolution (lp/mm) achievable with a given sensor. A minimum of two rows or columns of photosites is required to record a line-pair. Thus, if photosite spacing is 4 µm, then 8 µm (= 0.008 mm) are required to sample one line-pair. The Nyquist rate resolution is then 1/0.008 = 125 lp/mm.

---

[3] Acutance is defined here as the rate of change of brightness with distance. If a transition zone (edge) between white and black areas is wide, edge acutance will be relatively lower than if the transition zone is narrower. Resolution and acutance both contribute to the perceived sharpness of an image.

[4] Agranov, G., *et al.* 2011. Pixel continues to shrink....Small Pixels for Novel CMOS Image Sensors. 2011 International Image Sensor Workshop (IISW), Hokkaido, Japan. http://www.imagesensors.org/Past%20Workshops/2011%20Workshop/2011%20Papers/R01_Agranov_SmallPixel.pdf
Tian, H., *et al.* 2013. Architecture and Development of Next Generation Small BSI Pixels. 2013 International Image Sensor Workshop (IISW), Snowbird, Utah, USA. http://www.imagesensors.org/Past%20Workshops/2013%20Workshop/2013%20Papers/01-4_080-Tian-paper.pdf

2. Methods

The methods are the same as used in a previous paper, and need not be repeated in detail here.[5] As in my previous papers, a "perfect" lens is assumed. Diffraction blur is taken to be 70% of the diameter of the Airy disk for wavelength 550 nm.[6] Complications arising from a color filter array are ignored. Post-capture sharpening, which will increase micro-contrast, is not considered. Finally, I assume that there are no limits to fabricating sensors with ever smaller photosites, and no bandwidth limits with respect to processing sensor data.

3. Results

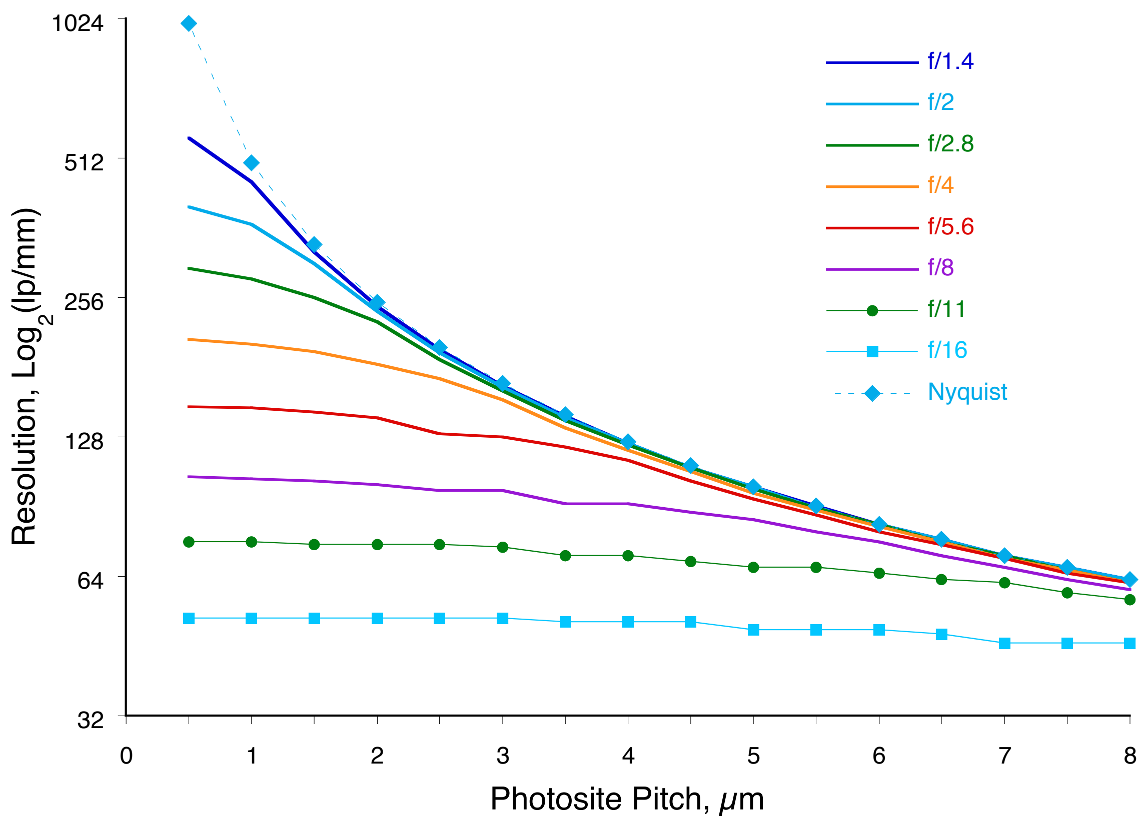

Fig. 1. Line-pair resolution at 50% contrast ratio. "Nyquist" is theoretical maximum resolution achievable with a given photosite pitch. Note the logarithmic scale for resolution.

[5] Service, Phil. 2014. Limits of Resolution. 2. Diffraction.
Service, Phil. 2014. Limits of Resolution. 3. Diffraction and Photosite Size.

[6] Other sources use the diameter of the Airy disk (first minimum of the circular diffraction pattern) as the estimate of diffraction blur. My tests (here and here) suggest that 70% of the Airy disk may be more accurate. For readers who prefer results for the conventional estimate of diffraction blur, an excellent approximation can be obtained from the data presented in this paper. Merely look at the results for the next numerically larger full f-stop. For example, to see what the results for f/4 would have been if I had used the full diameter of the Airy disk to estimate diffraction, look my results for f/5.6.

3.1. Resolution – 50% Contrast Ratio

Line-pair resolution as a function of aperture and photosite pitch is shown in Fig. 1. This figure illustrates several important points. (1) In general, and as expected, resolution increases with decreasing photosite pitch. For a perfect lens at f/1.4, resolution continues to increase all the way down to 0.5 µm. On the other hand, over the range of spacing used here — 0.5 – 8 µm — the effect of photosite spacing on resolution is almost negligible at f/16 and f/11. (2) In order to fully exploit the resolution potential of small photosites, it is necessary to use large apertures. For example, with 3 µm photosites, significant gains in resolution are obtained by opening the aperture to f/4. With 2 µm photosites, f/2.8 produces substantially higher resolution than f/4. At the same time, for any given photosite pitch there is a point beyond which further increases in aperture fail to yield significant increases in resolution. For example, with 5 µm photosites, there is little to be gained — in terms of line-pair resolution with 50% contrast ratio — by using apertures larger than f/5.6. (Note that I am NOT saying that there could be NO improvements to

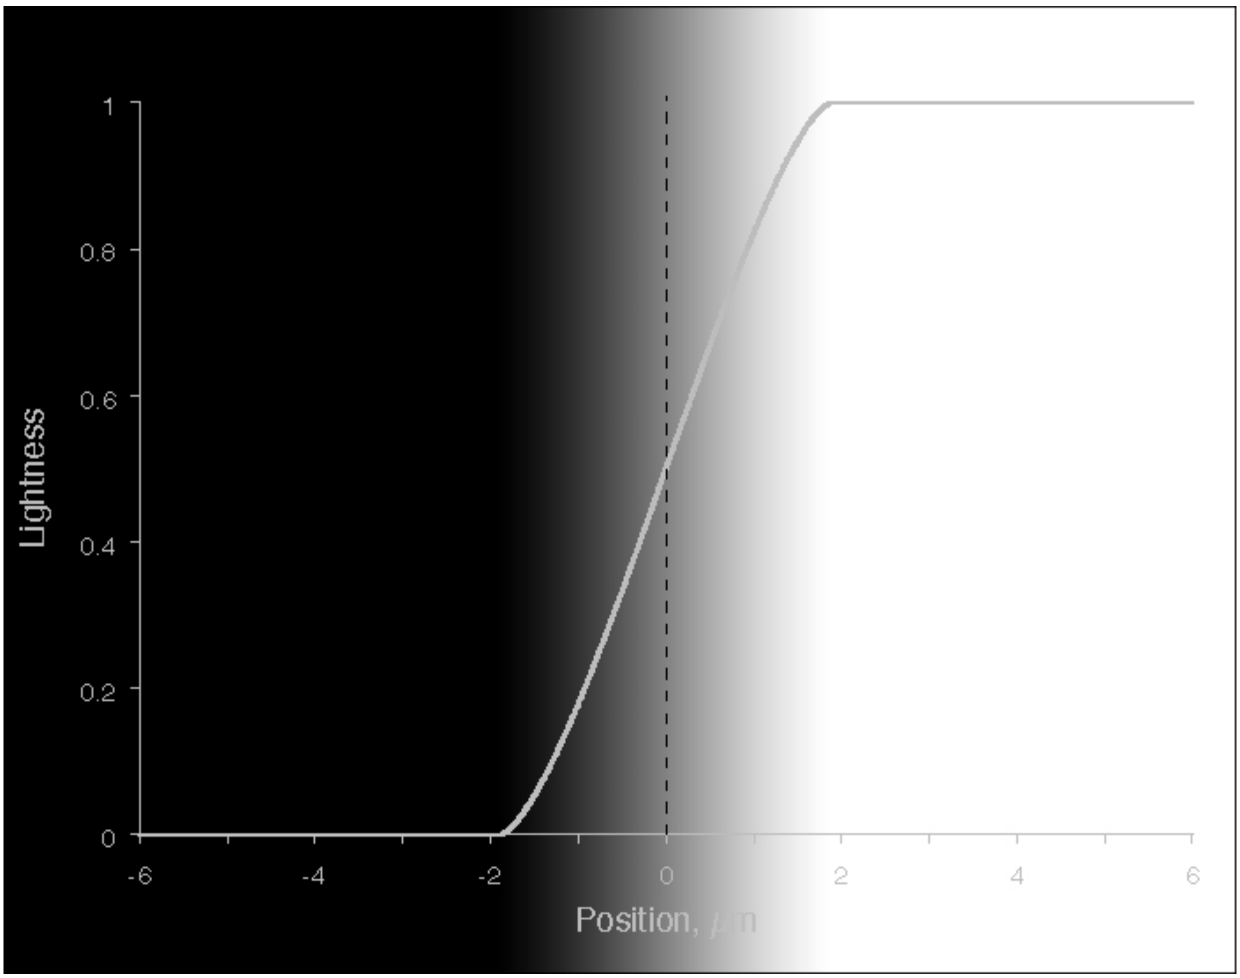

Fig. 2. Simulation and graphical representation of a blurred border between black and white areas. The un-blurred border would be at position 0 μm. The width of the blurred zone is approximately 3.76 μm, which corresponds to the diffraction blur circle diameter for f/4.

resolution or edge sharpness by using apertures larger than f/5.6 on 5 µm photosites,as will be discussed below.) (3) The converse of (2) is that small apertures seriously degrade resolution when photosites are small. Consider, for example, the effect of using f/11 with 4 µm photosites — relative to f/5.6 or f/4.[7]

3.2. Edge Acutance

Edge acutance is the second aspect of image detail. Well-defined edges of recognizable elements in an image contribute to the overall impression of sharpness. Edges are sharp if the boundaries between elements are narrow. Diffraction blurs edges. In a simple simulation of a straight edge, where there are no interactions with other nearby edges, the width of the blurred border is the width of the diffraction blur circle. Fig. 2 is a graphical depiction of blur at the border between a black area on the left side (lightness = 0) and a white area (lightness = 1). This particular graph is for f/4. Graphs for other apertures would be similar, the only difference being the width and slope of the transition zone.

3.2.1. Matters of Scale — Maintaining Perspective. In order to illustrate the effects of diffraction and photosite size on edge acutance it is necessary to use greatly magnified figures. Under such conditions, differences engendered by different apertures, for example, may seem quite noticeable — compare, for instance, the effects of diffraction at f/2.8 and f/8 in Fig. 3.[8] It is worth remembering, however, that most images are presented at *reduced* size on computer, tablet, or smart phone displays. The effects illustrated in this paper may not be evident under "normal" viewing conditions. And even if they are, whether they are important is a matter of individual taste and preference.

Fig. 3a represents a portion of a high contrast edge in an un-blurred *image field* (i.e., without diffraction effects).[9] Figs. 3b and 3c simulate the appearance of the edge with blur corresponding to f/2.8 and f/8, respectively. For concreteness, I have included a scale of measurement. The "features" of this image edge — projections and indentations — are on the order of 4 – 11 microns (µm). With a 50 mm lens, 4 – 11 µm in the image field corresponds to about 0.4 – 1.1 mm in the object field at a distance of 5 meters (16.4 ft). For comparison, the diameter of a human hair is about 0.1 mm.

In Fig. 4, I superimpose 2 µm and 4 µm photosites on the image field. The customary way of describing magnification when viewing digital images is the (linear) ratio of display pixels to image pixels (x 100%). By that metric, image pixels corresponding to the 4 µm photosites in Fig. 4b have been magnified by a factor of 8000% (each photosite is represented by an 80 x 80 pixel square). Similarly, the image pixels corrresponding to the 2 µm photosites (Fig. 4c) have been magnified by a factor of 4000%. That is not to say that the effects that I

---

7 Lest anyone think this example extreme, it is worth noting that the photosite pitch of the 16 MP m4/3 sensor in the Olympus OM-D E-M1 is approximately 3.7 *µ*m. With a reasonably good lens and suitably detailed subject matter, the degradation in image sharpness for f/11 compared to f/5.6 would be obvious.

8 Figs. 3 – 5 are contained in separate documents that can be accessed by clicking on the embedded links.

9 Whether one views this as a silhouette of a dark edge against a light background, as I tend to, or vice versa, is immaterial.

demonstrate would not be visible in images viewed at 100% magnification — only that they would be much less obvious.

3.2.2. Diffraction. Simulations of the effect diffraction blur on edge acutance are shown for f/2.8 and f/8 (Figs.3b and 3c). The diameter of the diffraction blur circle is about 2.6 µm for f/2.8 and about 7.5 µm for f/8. None of the major "features" of the edge are completely obscured by diffraction at f/8, although the edge is very "soft" compared to the f/2.8 example. Diffraction blur at f/16 (and possibly f/11) would most probably leave the narrow black lobe in the bottom third of the image unresolved.

3.2.3. Photosite Size (Sampling Rate). Four- and two-micron photosites are superimposed on the image field for the f/2.8 diffraction blur case (Figs. 4b and 4c). A given photosite can encode only one color, or shade of gray — obtained by averaging within the borders of each simulated photosite shown in Figs. 4b and 4c. The resulting simulated sensor images for the 4 µm and 2 µm photosites are shown in Fig. 5. It is clear that, for this particular example at least, 2 µm photosites do a much better job of reproducing the edge. It's not that the 4 µm photosites do not capture the major "features" of the edge, it's just that they do so very crudely. Note particularly that in the 4 µm case, contrast is reduced between the interdigitating black and white lobes.

Table 1. Resolution (millions of photosites) by Sensor Size and Photosite Pitch

Sensor Format	Photosite Pitch, µm			
	1	2	3	4
One-inch (13.2 x 8.8 mm)	116 MP (13,200 x 8,800)	29 MP (6,600 x 4,400)	13 MP (4,400 x 2,933)	7 MP (3,300 x 2,200)
Micro 4/3 (17.3 x 13 mm)	225 MP (17,300 x 13,000)	56 MP (8,650 x 6,500)	25 MP (5,767 x 4,333)	14 MP (4,325 x 3,250)
APS-C (23.7 x 15.8 mm)	374 MP (23,700 x 15,800)	94 MP (11,850 x 7,900)	42 MP (7,900 x 5,267)	23 MP (5,925 x 3,950)
"Full-frame" (36 x 24 mm)	864 MP (36,000 x 24,000)	216 MP (18,000 x 12,000)	96 MP (12,000 x 8,000)	54 MP (9,000 x 6,000)

4. Discussion

The illustrations of the effect of diffraction blur (Figs. 3b and 3c) show that diffraction sets an upper limit to realized edge acutance. Edge width (in the absence of sharpening) cannot be less than set by diffraction. Large photosites may increase the width of an edge, and therefore decrease acutance. On the other hand, very small photosites may sample an edge with relatively high fidelity. But if the edge is wide — perhaps because a small aperture has introduced substantial diffraction blur — small photosites will do nothing to enhance edge acutance — a detailed sample of a blur is still a blur. The unsurprising conclusion is that for maximum resolution and acutance *in the plane of focus*, use the largest apertures possible and sensors with the smallest photosites. Remember that we are assuming perfect lenses, and that depth of field is not a consideration. As already mentioned, some smart phone cameras have sensors with 1.5 µm or smaller photosites. The 20 MP 1" sensor used in the Sony RX100 III, for example, has 2.4 µm photosite spacing. A 24 MP APS-C sensor has 4 µm photosites. Table 1 shows the resolutions of larger format sensors for various photosites pitches. I have no idea if we will ever see 864 or 216 MP full-frame camera sensors. However, 96 MP does not seem out of reach in the reasonably near future.

4.1. Rule-of-thumb Approximations: Matching Image-field "Feature" Size, Diffraction Blur, and Photosite Size.

The smaller black lobe in (the bottom half of) Fig. 3a is about 5.5 x 4.5 µm. The f/2.8 diffraction blur circle is about 2.6 µm, and the lobe is well-depicted in the image field. At f/8, the diffraction blur circle is about about 7.5 µm, and the lobe is in danger of being lost to diffraction blur. Without considering photosite size, these results suggest that the smallest aperture consistent with a high likelihood of resolving a feature *x* microns in size in the image field, is the aperture associated with a diffraction blur circle *x* microns in diameter. For the black lobe in question, that would be about f/5.[10]

Sampling by photosites comes after image formation, and therefore after the "addition" of diffraction blur. Figs. 4 and 5 indicate that reasonably faithful sampling of the smaller black lobe blurred by diffraction at f/2.8 requires photosite spacing of 2 µm, or possibly less. That is, photosite spacing should be substantially less than the diffraction blur diameter — here 2 µm vs. 2.6 µm — in order to record all features that are present in the blurred image.

Note that I am not proposing hard-and-fast rules. Given that we are trying to record a feature of a particular size in the image field, if we choose to minimize diffraction blur, the constraints on photosite size may be eased somewhat. Alternatively, if diffraction blur pushes the feature to the limits of resolution, then it will be necessary to use very small photosites in order to sample the feature well enough to prevent it being lost from the sensor image.

4.2. Noise

The most commonly voiced objection to small photosites is that, all else being equal, they are noisier than larger photosites. However, the story of imaging sensors over the last 15 years, or so, is that "all else" is seldom equal. In particular, while there has been a progression to higher megapixel counts, and therefore smaller photosites, sensors have simultaneously become less noisy. There is some suggestion in recent data, however, that a plateau may have been reached.

Table 2 shows specifications for a number of camera sensors together with data on signal-to-noise ratios (SNR), and low-light ISO (SNR and ISO data published by DxO Mark). SNR results are for illumination equivalent to 18% and 1% gray-scale, with the camera set at ISO 200 (manufacturers' setting). 1% gray-scale illumination is approximately 6 2/3 EV below 100% gray, which, I believe, corresponds to sensor saturation. Thus, the 1% gray SNR (Table 2,

[10] Recall that the actual, object-field size of the *x* µm feature in the image field depends upon image magnification (*i.e.*, lens focal length and object distance).

Table 2. Signal-to-Noise Ratio (SNR) of Selected Camera Sensors (2004 – 2015)*

Model A	Launch Date (yyyy-mm) B	Total sensor Mpx C	Pixel pitch (µm) D	Nominal Pixel area ($µm^2$) E	18% Gray SNR @ ISO 200 (linear) (DxOMark) F	1% Gray SNR @ ISO 200 (linear) (DxOMark) G	1% Gray SNR / $µm^2$ (linear) H	Low-light ISO (DxOMark) I	Low-light ISO/$µm^2$ J
Full-frame (FX)									
D3	2007-08	12.20	8.40	70.48	118.85	20.42	0.29	2,290	32.49
D3x	2008-12	24.59	5.90	34.86	109.65	14.29	0.41	1,992	57.14
D3s	2009-10	12.20	8.40	70.48	134.90	21.38	0.30	3,253	46.15
D4	2012-01	16.43	7.21	52.01	138.04	22.39	0.43	2,965	57.01
D810	2014-06	36.37	4.86	23.66	130.32	14.96	0.63	2,853	120.57
A7S	2014-04	12.21	8.30	68.93	134.90	26.92	0.39	3,702	53.71
APS-C (DX)									
D70	2004-01	6.12	7.89	62.33	49.55	10.12	0.16	529	8.49
D50	2005-04	6.12	7.80	60.78	55.59	12.45	0.20	560	9.21
D300	2007-08	12.48	5.42	29.41	69.18	12.30	0.42	679	23.09
D90	2008-08	12.36	5.48	29.98	83.18	14.79	0.49	977	32.59
D7000	2010-09	16.37	4.73	22.36	81.28	13.34	0.60	1,167	52.19
D5200	2012-11	24.26	3.91	15.29	94.41	13.18	0.86	1,284	83.99
D7200	2015-03	24.16	3.91	15.26	95.50	13.34	0.87	1,333	87.36
One-inch									
RX100 III	2014-05	20.18	2.40	5.77	57.54	8.13	1.41	495	85.81

* SNR and Low-light ISO data from DxOMark. Original SNR data was in dB. I have transformed it to linear scale. Low-light ISO (column I) is defined as "the highest ISO setting for a camera that allows it to achieve an SNR of 30dB while keeping a good dynamic range of 9 EVs and a color depth of 18bits." [30 dB = 31.6 SNR (linear)]. Note that dynamic range "corresponds to the ratio between the highest brightness a camera can capture (saturation) and the lowest brightness it can capture (typically when noise becomes more important than the signal, i.e., a signal-to-noise ratio below 0 dB)." It should be pointed out that this is a very generous definition of dynamic range — 0 dB, or a linear SNR of 1, means that noise is equal to signal. ISO 200 refers to the manufacturers' on-camera ISO setting.

column G) gives us information about shadow noise. Low-light ISO (Table 2, column I; also referred to as the "Sports" score) is a measure of the usefulness of the camera in poorly lit situations. The 1% gray SNR and low-light ISO are given in absolute terms (columns G and I); and also "standardized" for difference in nominal photosite area (columns H and J).

With respect to APS-C sensors, manufacturers appear to have been unwilling to sacrifice shadow SNR (at low ISO) for increased resolution. Thus, the 1% gray SNR (column G) has remained quite steady over time in absolute terms (actually increasing slightly), despite the fact that pixel area has decreased by about 75% (compare the D5200 with the D70). This reflects real advances in sensor technology — essentially a five-fold improvement in shadow SNR on an area-adjusted basis (column H). At the same time, low-light ISO of APS-C sensors has improved about 2 1/2 times in absolute terms (column I), and about ten-fold on an area-adjusted basis (column J). The suggestion that sensor development has plateaued comes from the fact the APS-C sensor performance has remained quite static since 24 MP versions were introduced in 2012 (*e.g.,* Nikon D3200/D5200).

Unlike APS-C sensor development, in which resolution has always increased with time, full-frame sensors have carved out multiple niches. In particular, relatively low-resolution sensors with large photosites (*e.g.,* Nikon D3S and Nikon D4) coexist with sensors with more and smaller photosites (*e.g.,* Nikon D3x and D810). Not surprisingly, then, among the full-frame sensors there is a clear trade-off between photosite size and shadow SNR — compare, for example, the D810 and the D4 in Table 2, column G. Note also that there is much less proportional difference in 18% gray SNR between those two sensors (column F). In general, the latest full-frame sensors do not perform as well as the latest APS-C sensors on an area-adjusted basis, the one notable exception being the area-adjusted low-light ISO of the Nikon D810.

In order to get some idea of what it might mean in terms of noise performance if we were to have APS-C or FF sensors with 2 - 3 μm photosites, I have also included data for the Sony RX100 III in Table 2. Not surprisingly, in absolute terms, the sensor in the RX100 III is no match for the latest APS-C sensors. However, it is quite good on an area-adjusted basis — in fact it has by far the best area-adjusted 1% gray SNR among all the cameras shown. That said, given current technology, much higher resolution APC-C and FF sensors will be “noisier” than current sensors of the same format. It appears that we can choose to minimize noise by using relatively large photosites, or we can maximize resolution by using small photosites. But the days of increasing both ISO and resolution simultaneously are over, at least for now. That is perhaps an argument for modular cameras with interchangeable sensors, or “backs”.

4.3. Speculations about Lens Design for Maximum Image Sharpness

Throughout this series of papers, I have assumed perfect lenses and, therefore, that diffraction is the only source of blur. With perfect lenses, resolution and edge acutance always increase with larger apertures (provided that photosite size is not limiting). In real life, however, most photographic lenses perform best in the range of f/4 – f/8, where the total blur from diffraction and uncorrected lens aberrations is minimized. The clear message of this analysis is that maximizing image sharpness requires maximizing lens performance *at the largest practicable aperture.* Several factors seem likely to influence “largest practicable aperture”. These include: (1) lens maximum aperture; (2) sensor size and therefore image field size; (3) lens size, weight, and cost. I have no training in optics and lens design. Therefore what follows must be considered purely speculative.

4.3.1. Re-thinking the Need for Very Fast Lenses. Very fast (f/1.4 or faster) “normal” focal length prime lenses seem to occupy a special place in the minds of many photographers. The very best are very expensive, very large, very heavy, and very good by most measures. My feeling is that the preoccupation with fast lenses is largely baggage that we are still carrying from the film era, when film speeds were often less than ISO 100. The best lenses in this class usually show their highest resolution at about f/4 – f/5.6 — apparently, again, the apertures which minimize total blur from uncorrected lens aberrations and diffraction. My speculation is that maximizing lens speed conflicts with maximizing image sharpness at smaller apertures, say f/2.5 – f/3.5. If that is true, then it should be possible to design a lens with a maximum aperture of, say, f/2.8 that would be sharper than a much faster lens, *when both are used at f/2.8.* Our hypothetical, highly-corrected f/2.8 lens would most probably be lighter and more compact than its f/1.4 competition, and possibly less expensive, although not cheap. A *perfect* f/2.8 lens could make good use of photosites as small as 1 µm (Fig. 1); whereas a faster lens stopped down to f/4 (or f/5.6) — and performing perfectly — would be close to its 50% contrast resolution limit with 2 (or 3) µm photosites (Fig. 1). An obvious objection to this lens strategy, assuming it is technically feasible, is that the market for such lenses would be too small. That is, most photographers would not be willing to invest in relatively expensive, relatively "slow" lenses. Additionally, if small photosites mean low maximum ISO, as would appear to be the case, slow lenses would be adding insult to injury. Nevertheless, to take maximum advantage of very high resolution sensors I suggest that is will be necessary to re-think lens design, or at least to re-think the trade-off between lens speed and resolution.

4.3.2. Matching Lens Speed, Sensor Size, and Photosite Size. My impression is that it is more difficult to design fast lenses for larger sensors than for smaller ones. If that is true, then it is another reason to abandon the attachment to f/1.4 lenses for full-frame sensors. Leave the f/1.4 lenses for smaller sensors, and scale maximum aperture accordingly. For example, a perfect f/1.4 lens in front of a 1" sensor with 1 µm photosites would have a 50% contrast resolution of about 4,000 line-pairs per picture height (LP/PH). A perfect f/4 lens in front of a full-frame sensor sensor with 3 µm photosites could capture about 3,700 LP/PH — or about the same total resolution.[11]

The "cameras" described in the preceding paragraph are approximately equivalent: the "crop factor" for a 1" sensor is 2.73, or about 3. The photosite sizes also differed by a factor of 3, and the numerical f-values — 4 vs 1.4 — differed by a factor of 2.9. The example is meant to illustrate how aperture, sensor size, and photosite size can be scaled to yield images of similar resolution. However, even if we accept that f/4 is a desirable maximum aperture for lenses designed for full-frame sensors, there is no necessary reason to limit photosite size to 3 µm. For a perfect f/4 lens, 50% contrast resolution increases from 154 to 184 lp/mm when photosite size is reduced from 3 µm to 2 µm (Fig. 1). A further reduction in photosite size to 1 µm would yield a more modest additional improvement in resolution, to 203 lp/mm, indicating that increases in resolution due to higher sampling rate are being opposed by f/4 diffraction blur. On the other hand, if we could produce a perfect f/2.8 lens, 1 µm photosites would yield 281 lp/mm, and 2 µm photosites would give 227 lp/mm with 50% contrast. Both are substantial improvements over the f/4 case, although for a full frame sensor, 2 µm photosites would mean 216 MP (Table 1). If we decide that the smallest practicable photosites for a full-frame sensor are 3 µm, then there is negligible resolution benefit to increasing aperture to f/2.8 from f/4 — 161 vs 154 lp/mm with 50% contrast.

---

[11] These calculations are based on resolutions (lp/mm) presented in Table 2 of Limits of Resolution. 3. Diffraction and Photosite Size.

Whether or not a full-frame sensor, for example, can make maximal use of 3 or 2 µm photosites will depend on reasonably-priced lenses that are up to the task. As I write this (July, 2015) the highest resolution full-frame sensor currently available is the 50 MP sensor in the Canon 5DS and 5DS R, with 4.1 µm photosite spacing. Presumably, some careful testing will tell us if any current lenses can extract the full potential of that sensor. The Nyquist limit resolution of the Canon sensor is about 120 lp/mm or 2,880 LP/PH — theoretically achievable with 50% contrast by a perfect lens at f/4.

4.4. Oversampling — Why Sensors with Hundreds of Megapixels May be a Good Thing

The clear message of these simulations is that if diffraction and photosite size are the only factors that limit resolution, then photosites of 1 – 2 µm will yield useful resolution gains, provided that apertures are relatively large — ≥ f/4. But, *even without resolution gains* — as might be the case with less than perfect lenses — I suggest that *oversampling,* in its various guises, is the primary reason why sensors with 100 or more megapixels will be a good thing. Oversampling may be useful in at least three ways: (1) for producing sharper *final* images; (2) for producing less noisy *final* images; and (3) for making higher quality prints.

4.4.1. Oversampling for more sharpness and less noise. The vast majority of digital images are downsized for viewing. Whenever that is the case, image capture entails *oversampling* because the initial pixel count is greater than the pixel count of the final, displayed image. This sort of oversampling seldom seems to be a conscious strategy — who goes about capturing images with the intention of "throwing away" pixels? — yet that is what is done in most cases. If we have cameras with 100 or more megapixels, practically all images will be down-sampled for viewing, and so virtually all images will be oversampled at capture. The question is: can *capture* oversampling — and the consequent down-sampling of the final image — be a useful strategy for improving image quality? There are two reasons to believe that it might be: down-sampling improves sharpness and reduces noise. The crucial issue to be investigated, however, is whether a down-sampled image will be sharper and less noisy than an alternative image captured at lower resolution, but not subsequently down-sampled. To be concrete, consider the following two APS-C format sensors. Sensor X has 6 MP (3,000 x 2,000) and sensor Y has 24 MP (6,000 x 4,000). Both sensors are to be used with the same lens and exposure settings, say f/5.6, 1/250 sec, ISO 200. The final image size will be 3,000 x 2,000 pixels, a size that will fit on a 4K display. What we want to know is: will the down-sampled image from the 24 MP sensor be sharper and less noisy than the un-resampled image from the 6 MP sensor?

The answer will certainly depend on the method used for down-sampling. With regard to sharpness, the problem is illustrated by Figs. 4 and 5. Because we are using the same lens at the same aperture, the image fields will be identical. Our 6 MP sensor X is analogous to Fig. 4b, and the 24 MP sensor Y is analogous to Fig. 4c. The corresponding sensor images are shown in Figs. 5b and c. What we want to know is, if Fig. 5c were down-sampled to 80 pixels (simulated) from its native 320 pixels (simulated), would the final result be sharper than the sensor image depicted in Fig. 5b? My guess is that in many cases, the answer would be "yes", if for no other reason than sharpening seems to be "built-in" to many down-sampling algorithms — for example, Photoshop Bicubic Sharper.

The one source of noise that is directly related to photosite size is the randomness of photon arrival at photosites. That is sometimes referred to as "photon shot noise". It is unavoidable, but usually not noticed in bright areas of an image, where the signal-to-noise ratio is high. For our hypothetical 6 and 24 MP sensors, the pixel-level signal-to-shot-noise ratio of the 24 MP sensor will be one-half that of the 6 MP sensor (all else being equal).[12] In other words, although the area of sensor Y photosites is only one-fourth that of sensor X photosites, the signal-to-shot -noise ratio is reduced by only one-half. That is reason to think that "intelligent" down-sampling could produce a 6 MP final image from sensor Y that would be less noisy than the 6 MP image from sensor X.

Note that this argument that oversampling may reduce noise in the final image does not contradict the earlier assertion that sensors with smaller photosites will most likely be "noisier" than sensors with larger photosites. The argument about possible lower noise in down-sampled images applies when sensors are used at, or close to, base ISO — and provided that downsampling is done "intelligently".

4.4.2. Oversampling for printing. The Epson Stylus Pro 3880 is an excellent printer capable of making 22 x 17 inch prints. I have argued that in order to extract the maximum possible resolution from this printer *and from images*, it is necessary to print at 720 ppi.[13] Therefore, to make a moderately large print that is 18 x 12" in printed area, we need an image that is 12,960 x 8,640 = 112 MP. This is oversampling because the Epson 3880 has a second, lower native resolution of 360 ppi. In other words, we "need" only a 6,480 x 4,320 = 28 MP image to make out 18 x 12" print. In fact, 720 ppi corresponds to about 14.2 lp/mm, which is beyond the limits of human visual resolution — although not beyond the printing capabilities of the 3880 — and for carefully prepared images the difference between the two printing resolutions may not be easily perceived. But, and this is the key point, the oversampling at the stage of image capture means that we can print at the higher native printer resolution without resampling the image; and *the detail which we can see in the print* may be rendered with greater acutance than if the image had been captured at the lower resolution.[14] The problem with typical arguments about how many image pixels are "needed" for printing at a given size is that they ignore the possibility that capture oversampling may improve the quality of the captured and, therefore, printed images.

4.5. Conclusion

Given that 2.4 µm photosite spacing is already being used with 1" sensors, it seems reasonable to assume that 2 - 3 µm photosites will eventually be used for APS-C and "full-

[12] This is a direct consequence of the fact that the number of photons arriving at each photosite is Poisson distributed. For more details see this page at ClarkVision.

[13] Service, Phil. 2015. Limits of Resolution. 4. Image Capture for Maximum Detail Printing.

[14] Assuming, of course, that image blur is controlled well-enough to take advantage of the "oversampling" sensor.

frame" sensors, implying total resolutions >100 MP. In order for such sensors to perform close to their theoretical resolution limits, it may be necessary to sacrifice lens "speed" in order to optimize sharpness at f/2.8 - f/4. Even if such sensors do not achieve their theoretical resolution limits — perhaps because the necessary lenses are not available — they may still provide advantages over current 24 – 50 MP sensors. In particular, most >100 MP images will be down-sampled for "end use". I suggest that such down-sampling may result in a sharper and less "noisy" final image than would otherwise be obtained with lower resolution sensors.