ISO-less?

Henry Gordon Dietz* and Paul Selegue Eberhart*
University of Kentucky, Lexington, Kentucky, USA

Abstract. *To determine the proper exposure, cameras generally use the concept of "film speed" – a number representing the film's sensitivity to light. For film, this number was a function of the emulsion and processing, changeable only in batches. However, digital cameras essentially process each shot individually, so most adopted the idea that the film speed of the sensor could be changed for each shot. The catch is that it isn't clear that the sensitivity of a sensor used in a digital camera can be adjusted at all: many digital cameras have been claimed to be "ISO-less," capable of producing similar images for the same exposure independent of the ISO setting used. This paper will present the results of testing the ISO-less behavior of various digital cameras, concluding with a simple proposal for how these results could be used to create a new paradigm for computing exposure and processing parameters.*

1. INTRODUCTION

The sensitivity of photographic emulsions to light is now most commonly specified as the ISO sensitivity or speed. It is a fundamental feature of most digital cameras that the ISO may be changed from one shot to the next. However, it is not clear exactly what is changed to alter the ISO, and thus the impact of a change in ISO is also unclear. A camera is said to be "ISO-less" if reducing the ISO to well below the value needed to make a properly-exposed "raw" capture does not result in a final image of significantly different quality.

For example, Figure 1 shows crops from images captured with a Sony A7. The leftmost image was captured at ISO 100. The middle image was captured at ISO 1600. The rightmost image was captured with the camera set at ISO 100, but was underexposed by 4 stops (as if it had been ISO 1600) and brightened in post-processing. The fact that the middle and rightmost images are essentially identical implies some level of ISO-less behavior.

This paper experimentally explores the phenomenon of ISO-less behavior and proposes a new approach to computing exposure and processing parameters based on the new understanding.

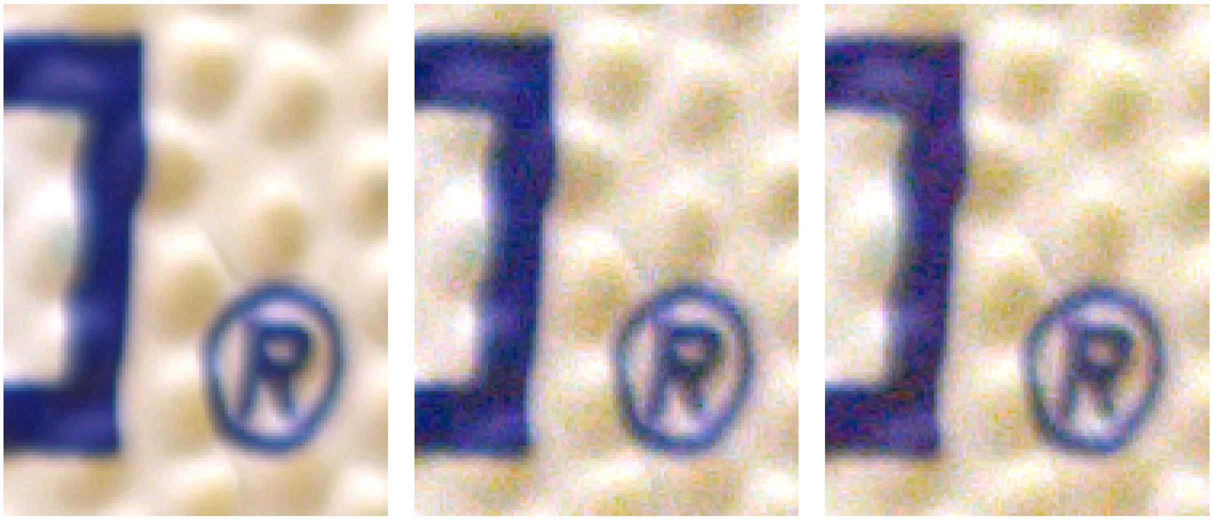

Figure 1. Sony A7 120x160 crops with ISO @100, @1600, and @100 exposed as if 1600

---

* University of Kentucky Department of Electrical and Computer Engineering and the Center for Visualization and Virtual Environments. E-mail: `hankd@engr.uky.edu`

2. WHAT IS ISO?

The concept of specifying the sensitivity of photographic emulsions to light by a single number dates at least to Warnerke's Sensitometer in 1880.[7] In the following decades, a variety of systems relating sensitivity to exposure time were proposed. The 1934 DIN standard 4512 mapped sensitivities into integer numbers in a base 10 log scale and written with a degree mark. The ASA Z38.2.1-1943 standard established an arithmetic scale of film speeds which, through a number of revisions, persists as the linear speed value in the ISO film speed system in common use today. In fact, since 1974, ISO values have officially been pairs of ASA and DIN values, for example, ISO 100/21°. Internal to many digital cameras, ISO film speed is mapped into log-based APEX[11] values to simplify exposure computations.

Determination of ISO film speed is a somewhat different process for different types of film (e.g., black-and-white negative, color negative, and color reversal), so it is not surprising that determining speed of sensors used in digital cameras also requires distinct methods, in particular, ISO 12232:20076 *Photography – Digital still cameras – Determination of exposure index, ISO speed ratings, standard output sensitivity, and recommended exposure index*.[10] This document actually provides manufacturers with five different techniques for determining ISO. The three methods from the original 1998 version of the standard compute ISO speed from sensor properties such as sensitivity and noise, while the two 2006 additions compute ISO values based on appearance of the JPEG image produced: SOS and REI. The SOS (Standard Output Sensitivity) is based on matching a reference Y-channel value in the JPEG, and should consistently be about 0.704 times the ISO speed, resulting in somewhat brighter images. The REI (Recommended Exposure Index) is somewhat more arbitrary, intended to approximate the value that would generally give the best metering results for most users and scenes. Clearly, the standard offers manufacturers significant flexibility in computing ISO values.

The question this paper is concerned with is not how the ISO values are computed, but how a camera's *apparent ISO* value, perhaps more properly called EI (exposure index), could be altered and with what impact on image quality. Sensing elements (sensels) used in digital cameras are essentially analog photon counters; thus, it would seem obvious that sensitivity would correspond to the fraction of charge from incident photons recorded – the Quantum Efficiency (QE). However, QE of a sensor is not electrically alterable. Thus, the question becomes what does the ISO setting of a digital a camera really change?

Changing the ISO setting alters how the camera selects the pair of shutter speed and aperture settings to use for exposures, but that tells the user nothing about how the sensitivity change is accomplished. Anecdotally, changing the ISO setting is commonly described as changing the sensor gain, which presumably is implemented by changing the analog amplification before the sensel value is digitized. In early consumer digital cameras, the linear ADC (analog to digital converter) was indeed the weak link in the processing chain with respect to dynamic range. Cameras typically used ADCs providing no more than 8 bits of output with the least-significant 2-3 bits of very poor quality. Thus, although the analog sensor itself might not change sensitivity, changing the ISO setting could be accomplished by adjusting the analog amplification to shift the ADC's usable range.

As camera electronics improved, ADCs have grown to commonly output 12-14 bit digital values – arguably covering the complete dynamic range of the analog sensor without any need to alter analog amplification. This is a basic assumption behind the concept of a camera being "*ISO-less*." It has been observed that some cameras appear to capture just as much scene information no matter how the ISO is set, so that post-processing grossly underexposed low-ISO raw images can produce final images of comparable quality to those with the camera ISO setting boosted to allow correct exposure with the same shutter speed and aperture settings. In various on-line photography forums, this claim is particularly commonly made for Sony cameras – and appears valid, as Figure 1 shows.

However, there are some cameras that clearly are not ISO-less. In on-line forums, Canon EOS cameras are commonly described as benefiting from changing the ISO setting in the camera. In fact, the Magic Lantern (ML) "dual ISO" support[1] uses the strongly ISO-dependent behavior of the sensor to expand dynamic range. In Canon EOS cameras that use two readout channels, the channels correspond to interleaved sets of pixel lines. By setting different ISOs on the control registers for the two separate channels and intelligently combining the results, ML merges the two partially-overlapping dynamic ranges into a larger dynamic range. For example, ISO 100 has a dynamic range of about 11 stops and ISO 1600 has about 10 stops, but these ranges are offset by 4

Camera Model	**DR @ 8MP**	**100**	**200**	**400**	**800**	**1600**	**3200**	**6400**	**12800**	**25600**
Sony A100	11.2	10.9	10.2	*8.9*	*7.7*	*6.9*				
Sony A350	11.5	11.1	*10.3*	*9.2*	*7.6*	7.2	*6.2*			
Sony SLT-A55	12.4	12.0	11.5	11.0	*10.2*	*9.4*	*8.4*	*7.3*	*6.2*	
Sony NEX-5	12.2		11.9	11.6	*10.6*	*9.4*	8.7	*7.5*	*6.4*	
Fuji X10	11.3	11.2	*10.4*	9.8	9.3	*8.2*				
Sony NEX-7	13.4	12.8	*11.9*	*11.1*	*10.1*	*9.1*	*8.2*	7.9	*6.7*	6.4
Canon EOS-M	11.2	10.7	10.7	10.5	10.0	9.3	8.6	*7.4*	*6.5*	
Sony A7	14.2	13.4	*12.5*	11.9	11.1	10.5	*9.3*	*8.3*	*7.3*	6.7

Table 1. DxO measured dynamic ranges in stops at various ISO settings

stops; thus, a combined dynamic range of about 14 stops can be produced using dual ISOs. Both channels use the same exposure parameters other than ISO, so if the camera had been ISO-less, the dynamic range of the lower ISO would have completely covered the dynamic range of the higher ISO, and there would have been no increase in dynamic range.

There does not seem to be a direct correlation between the manufacturer-stated number of bits from the ADC and ISO-less behavior; some Sonys that seem ISO-less have ADCs outputting fewer bits than some Canons that are clearly not ISO-less. Noise issues may still provide a benefit from changing the analog amplification no matter how many bits are delivered. Perhaps ISO-less behavior is really an artifact of sensor/ADC architecture or fabrication technology?

The research reported here does not seek to determine the causes of ISO-less behavior. Rather, the goal is to observe to what extent the phenomenon occurs and to propose ways in which it can be used to improve image quality.

3. TESTING ISO-LESS BEHAVIOR

ISO-less behavior often is claimed for various cameras in on-line photography forums such as `DPReview.Com`, but there is little documentation of it. There also is no standard definition of precisely what requirements must be met to be ISO-less. The ideal definition of ISO-less behavior would be that all final image attributes are indistinguishable between two images exposed identically but with different ISO settings. However, it is reasonable to expect that there will be at least minor differences. This section of the paper presents a simple empirical evaluation of when ISO-less behavior occurs and by which image qualities that ideal behavior is violated.

Three separate mechanisms are used for this evaluation. Section 3.1 uses published dynamic range measurements for 8 cameras to determine under what circumstances ISO-less behavior is possible. To better understand how properties other than dynamic range change, "raw" images for a test scene were evaluated for all full-stop combinations of exposure parameters and ISO settings using 19 cameras. These results are summarized in Section 3.2. Of course, images represented as "raw" sensor data will not exhibit the ISO-less property until they have been processed to create final images for viewing. The processing used to convert "raw" data into final JPEG image files can have a huge impact on the differences between the final images captured with the same exposure settings but different ISOs. Although minimal processing was used in Section 3.2 to avoid biasing the results, the ground truth would be to use the exact same JPEG processing pipeline for both processing of normally-exposed images and those shot with a lower ISO than exposure conditions nominally required. In Section 3.3, results are summarized for using CHDK[3] to reprogram 8 different Canon PowerShot cameras to appropriately adjust the raw data and feed it through precisely the same in-camera JPEG pipeline used for normal exposures.

3.1 Dynamic Range Measurements

One way to determine if cameras are potentially ISO-less is to examine published measurements of dynamic range. To be ISO-less, the same range of scene tones must be able to be recorded. Thus, it is necessary (but not sufficient) that dynamic range increase by approximately one stop for each one stop drop in the ISO.

DxO has published graphs showing thousands of such measurements on their WWW site,[5] and their data has been used for various other analysis, for example at `Sensorgen.Info`.[4] Among that data are dynamic

Figure 2. Test scene for judging ISO-less behavior

range measurements for eight of the consumer cameras used in the research reported here (see Section 3.2). That data is repeated here in Table 1, with font styles used to indicate entries potentially involved in ISO-less behavior. The entries in **bold** are those ISOs which potentially satisfy the constraint that dynamic range increases by approximately the same amount that ISO setting is decreased from the next-highest ISO setting; a slope threshold of 0.8 is used instead of 1.0 to allow for roundoff errors in reported values of dynamic range and differences between ISO settings and measured ISOs. If all potentially ISO-less cases are actually ISO-less, the entries marked in *red italics* would never be beneficial. For example, on a Sony NEX-7, ISO 100 could be harmlessly substituted for any ISO from 200 to 3200.

Table 1 shows significant ranges of potentially ISO-less behavior for each camera, but *none of these cameras could be completely ISO-less.* The positioning of the runs seems to follow a rough pattern, with a run starting near the minimum ISO setting and a second run starting around ISO 3200 and ending near the maximum ISO setting. Near the minimum ISO setting is often a special case in that it is where dynamic range of the sensor is most likely to exceed that of the analog/digital processing pipeline, resulting in clipping of the dynamic range. The EOS-M data is consistent with the common notion that the processing pipeline in Canon's EOS-family cameras is unable to handle as large a dynamic range as some other cameras, thus limiting performance at lower ISOs despite fairly good performance from the sensor at higher ISOs. The gap in ISO-less potential commonly seen around ISO 1600 might be explained by processing pipelines enabling noise reduction logic above that ISO. As will be discussed in Section 4.4, dynamic range measurements are based on measuring signal/noise ratio. Despite making the image blurry, aggressive noise reduction can improve the measured signal/noise ratio, resulting in inflated dynamic range measurements for high ISOs and a discontinuity around the ISO where the reduction becomes enabled.

3.2 Test Scene "Raw" Sensor Data

While covering the dynamic range is a fundamental requirement for ISO-less behavior, it is not the only requirement. All image properties should be similar across ISOs. Significant properties include dynamic range, the noise

Figure 3. Cameras used in the experiments

Camera Model	Year	Prog	Sensor	MP	Min ISO	Max ISO	BPP
Canon G1	2000		CCD	3	50	400	10
Sony F828	2004		CCD	8	64	800	14
Canon A620	2005	CHDK	CCD	7	50	400	10
Canon A640	2006	CHDK	CCD	10	80	800	10
Sony A100	2006		CCD	10	100	1600	12
Canon A590	2008	CHDK	CCD	8	80	1600	10
Canon SD770	2008	CHDK	CCD	10	80	1600	12
Sony A350	2008		CCD	14	100	3200	12
Canon A480	2009	CHDK	CCD	10	80	1600	12
Sony SLT-A55	2010		CMOS	16	100	12800	12
Sony NEX-5	2010		CMOS	14	200	12800	12
Fuji X10	2011		CMOS	12	100	3200	12
Sony NEX-7	2011		CMOS	24	100	16000	12
Canon A4000	2012	CHDK	CCD	16	100	1600	12
Canon EOS-M	2012	ML	CMOS	18	100	12800	14
Canon ELPH115	2013	CHDK	CCD	16	100	1600	12
Canon N	2013	CHDK	CMOS	12	80	6400	12
Sony A7	2013		CMOS	24	50	25600	14
Sony A7 II	2014		CMOS	24	50	25600	14

Table 2. Some properties of the cameras used in the experiments

Figure 4. Crops from Canon A4000, Sony NEX-7, and Canon EOS-M

distribution, the retention of texture detail and sharpness, and handling of color information. Color is important because each color channel typically has somewhat different sensitivity, making most image characteristics color dependent. Some of these image properties are difficult to quantify, so our approach was based on informally comparing images of a test scene constructed specifically to reveal the kinds of differences expected. This scene is shown in Figure 2. Using each camera, the scene was shot at all full-stop ISO settings with exposures ranging from being appropriate for the highest ISO to matching the exposure needed for the lowest ISO.

In order to perform this testing, a group of 19 cameras, all supporting capture of "raw" sensor data, were collected. The cameras are shown in Figure 3 and some of their properties are described in Table 2. The cameras span nearly a decade and a half, including mirror-less, SLT, DSLR, and compact formats. Many of these particular cameras were selected because they are programmable – a feature used in Section 3.3. Both CCD and CMOS sensor technologies are represented, with a wide range of pixel resolutions (and hence of pixel sizes) and bits-per-pixel (BPP) depths in the digital portion of their raw processing pipelines.

Although thousands of "raw" test exposures were made, the findings can be summarized concisely.

As expected, completely ISO-less behavior was not seen from any of the cameras discussed in Section 3.1. However, it was seen from *all* the Canon PowerShot compact cameras (but neither the G1 nor EOS-M). The raw files produced very similar final images even when the lowest ISO setting was used with exposure parameters matching the highest available ISO. Surprisingly, even a grossly underexposed JPEG capture from these cameras could be brightened in post-processing to produce an acceptable image – perhaps because the camera's dynamic range does not dramatically exceed that of JPEG encoding. The top line of images in Figure 4 shows three crops from an A4000. The first two are minimally-processed raw images with ISO set to 1600 and 100, but exposure parameters set for ISO 1600. The third image is an ISO 100 JPEG captured with the same exposure parameters, and then brightened in post processing; the image looks surprisingly good, except in that dark texture information is missing.

Cameras that are often described as being ISO-less, such as most cameras using Sony sensors, showed many ISO-less properties across ranges of ISOs, but were not completely ISO-less even within a run. There were easily visible variations in rendering of textural details, colors, and even noise. The variations were not always favoring the ISO by which the exposure was set. For example, the middle row of Figure 4 shows crops from three captures of our test scene with a NEX-7 using exposure parameters set as though ISO 16000 was used (again, post-processing of the raw images was kept to a minimum e.g., no noise reduction was applied). The image captured at ISO 16000 is grainy, but not bad for such a high ISO. The image captured at ISO 800 has slightly more noise, but it also has noticeably more textural detail in the cloth – perhaps Sony is applying some noise reduction processing at higher ISO settings? The ISO 100 image suffers additional noise and a significant reduction in dynamic range, but the image is still recoverable and textural detail may again be superior to that of the native ISO 16000 capture. The decade-old Sony F828 showed similar properties, but image quality was overall *consistently better* digitally brightening raw images (or even JPEGs) captured at lower ISOs than the exposure parameters would suggest, perhaps in part because it used a 14-bit image pipeline that clearly exceeded the dynamic range capabilities of the sensor.

Canon's EOS cameras, which the EOS-M represents in our tests, are often described as being very different from those using Sony sensors, and certainly not ISO-less. The data presented in Table 1 makes it clear that the EOS-M cannot be ISO-less. However, using the ISO 800 setting with exposure parameters for ISO 12800 actually compromises dynamic range by only 1/2 stop. The result, shown in the bottom row of Figure 4, is that a boosted ISO 800 image is not entirely worse than one shot at the native ISO 12800 setting. Noise is higher, but again texture detail is preserved a little better. Unfortunately, the ISO 100 crop clearly suffers from too much loss of dynamic range, and is probably beyond recovery. The ancient Canon G1 proved to be even further from ISO-less.

In summary, the raw test scene images suggest cameras fall into two categories. The first group includes the Canon A620, Canon A640, Canon A590, Canon SD770, Canon A480, Canon A4000, Canon ELPH115, and Canon N. These (typically low-end cameras) are essentially ISO-less, and there is little reason not to always shoot at the lowest available ISO. The second group includes the Canon G1, Sony F828, Sony A100, Sony A350, Sony SLT-A55, Sony NEX-5, Fuji X10, Sony NEX-7, Canon EOS-M, Sony A7, and Sony A7 II. These cameras

Figure 5. JPEG Crops from Canon A4000 at ISO 1600, 16x ISO 100, and 16x ISO 100 filtered

can make usable images at a variety of ISOs, but with visible differences in specific image quality metrics. These quality metrics do not always favor the ISO setting normally associated with proper exposure. Thus, there is the potential to improve image quality by characterizing image quality properties of different ISO settings and carefully selecting which to use for each capture (as discussed in Section 4.3).

3.3 Test Scene JPEG Pipeline Data

As compelling as the above raw image tests are, most photographers shoot JPEG images most of the time. Raw post processing software has access to computational resources far beyond those found in typical cameras. It is also possible for a skilled user to control the processing in sophisticated ways that could bias comparisons of post-processed raw captures. A more fair test would directly employ the camera's JPEG processing pipeline to create the final images for comparison.

Directly using the standard in-camera processing to encode final JPEG images from underexposed raw data would result in dark JPEGs. However, CHDK (the Canon Hack Development Kit)[3] allows arbitrary user-written code to run inside 8 of the Canon PowerShot models listed in Table 2. The user code to run inside one of these cameras can be written in interpreted BASIC or Lua or, with some additional complications, C compiled into fast native ARM32 code. Thus, all that is needed is a way to boost the raw pixel values before feeding the raw data into the camera's JPEG pipeline.

Re-scaling the data in the camera's raw memory buffer to compensate for underexposure could be done by compiled C code, but CHDK now provides "raw development" support functions that can do the boosting and JPEG encoding directly from a Lua script. Fundamentally, for each stop of underexposure due to a lower ISO setting, the raw pixel values need to be doubled. CHDK's raw development support includes a function that can add or average a series of images, so multiplying an image's pixels by two is accomplished by a Lua script that adds the image to itself. Multiplication by larger numbers is not as tedious as one might expect; a sequence like `i=i+i; i=i+i; i=i+i; i=i+i;` efficiently implements multiplication of `i` by 16, which would appropriately brighten an ISO 100 image shot captured using an exposure suitable for ISO 1600. This programming is simple enough that this task was used as an assignment in the EE599/EE699 "Cameras as Computing Systems" course that Dietz taught in Fall 2013 and again in Fall 2014.

Figure 5 shows the quality of JPEGs generated in-camera from ISO-less exposures. The leftmost image crop is from the native JPEG generated by a Canon A4000 using a setting of ISO 1600 with suitable exposure parameters. The center crop is taken from a 16x ISO 100 JPEG given the same exposure. It was created using Lua code to direct the camera to make a raw capture, multiply the pixel values by 16, and then render a JPEG.

This crop appears dramatically sharper than the "properly exposed" version. Unfortunately, it also is much noisier, making it difficult to judge which is better. The rightmost image resolves that question by showing that the noise level of the ISO 100 image can be reduced to be comparable to that of the native ISO 1600 shot, yet still preserve more detail. The denoising of the last image was done outside of the camera, but the algorithm used easily could be implemented as compiled C code under CHDK.

It is worth noting that compiled C code takes seconds to scan the 16MP raw buffer of the Canon A4000, so the boosting operation would be annoyingly slow implemented in this way. However, the boosting algorithm is highly suitable for integration with the existing hardware enhancements in the JPEG pipeline.

4. ISO-LESS EXPOSURE

The concept of ISO film speed was primarily about simplifying exposure computations, but with powerful computers in every camera, this simplification is no longer necessary. As discussed above, no digital camera tested is literally ISO-less in the sense that there are variations in image quality dependent on ISO setting. In fact, significant image qualities depend on the mix of analog and digital processing used to achieve a particular effective ISO. The inescapable conclusion is that using ISO to determine exposure parameters is a simplification that, by ignoring such variations, may compromise image quality.

Thus, the preferred notion of "ISO-less" cameras in the current work is not that the user ISO setting does not matter, but that the computation of exposure parameters is fundamentally based on analysis rather than an ISO-based concept of exposure equivalence. To clarify the difference, it is first useful to summarize how ISO is traditionally used in APEX exposure computations. The following subsections present several alternative ISO-less methods for computing exposure parameters with the goal of achieving superior image quality.

4.1 APEX

The precise method for computing exposure settings inside most digital cameras is not documented. However, some of the EXIF data recorded in image files suggests APEX (Additive System of Photographic Exposure)[11] computation. CHDK exposes the fact that Canon PowerShot cameras internally employ a variant of APEX known as APEX96.

Sensitivity to changes in exposure factors in general is logarithmic, so APEX maps exposure parameters into a base 2 log system in which a unit change in any parameter corresponds to doubling or halving that factor – a one stop change. This is the math behind the common photographic knowledge of how aperture and shutter speed settings trade off.

The basic APEX formulation is:

$$Av + Tv = Bv + Sv = Ev$$

In which:

Av Aperture value represents the rate of light transmission by the lens such that adding 1 to the Av value represents delivery of light at half the rate. Using a perfect lens, Av is determined solely by the aperture f/number, which is simply the ratio of the lens focal length divided by the diameter of its circular aperture. However, for real lenses, reflections and other imperfections reduce the light transmitted by a small amount, so it would be more correct to say that Av is determined by the transmission-corrected effective f/number, or T/number. If f/1.0 is assigned Av=0, then f/1.4 is 1, f/2.0 is 2, etc.

Tv Time value represents the exposure integration period, commonly known as shutter speed even for systems that lack a mechanical shutter. An increase of Tv by 1 represents halving the integration period. If 1s is assigned Tv=0, then 1/2s is 1, 1/4s is 2, etc. Although the progression is in powers of 2, the markings on most shutter controls round values to multiples of 5 or 10; thus, 1/15s is really 1/16s, 1/125s is really 1/128s, etc. Film typically suffers *reciprocity failure* at long shutter speeds which would break the log-linearity of Tv values, but that effect is usually negligible at shutter speeds less than 1s and does not exist per se for electronic sensors.

Algorithm 1 example P mode APEX exposure computation

```
if $Ev_{target}$ - $Av_{min}$ < $Tv_{blur}$ then $Av$ = $Av_{min}$ and $Tv$ = $Ev_{target}$ - $Av_{min}$
else if $Ev_{target}$ - $Tv_{blur}$ < $Av_{limit}$ then $Tv$ = $Tv_{blur}$ and $Av$ = $Ev_{target}$ - $Tv_{blur}$
else $Av$ = $Av_{limit}$ and $Tv$ = $Ev_{target}$ - $Av_{limit}$
```

Bv Brightness value represents the metered scene luminance such that a unit increase in Bv represents doubling the brightness. If 1 foot-lambert (fL) is Bv=0, 2 fL is 1, 4 fL is 2, etc. Luminance is often expressed in candela per square meter (cd/m^2); 3.4262591 cd/m^2 is 1 fL.

Sv Speed value represents the light sensitivity of the film or sensor – the ISO. A unit increase in Sv means the image can be formed with half as much light. ISO 100 is Sv=5, which means that Sv=0 is around ISO 3.125.

Ev Exposure value represents the total amount of image-forming light, but is most often used as a shorthand to express any equivalent combination of $Av + Tv$ or Bv + Sv. In other words, two exposures are expected to produce "equivalent" images as long as Ev is the same. For example, $f/1.4$ for 1s is Ev=1 (1+0), but so is $f/1.0$ for 1/2s (0+1) or f/2.0 for 2s (2+-1).

While these values are real numbers, in practice it is rare that settings for aperture or shutter speed are calibrated with more than 2 or 3 markings per stop. Thus, it is convenient to implement the computation using scaled integers. In APEX96, the APEX values are multiplied by 96, which probably provides higher accuracy than can be achieved by the mechanical systems controlling most apertures and shutters.

For a camera using APEX with a user-specified ISO, automating exposure setting is simply a matter of determining $Ev_{\text{target}} = Bv + Sv_{\text{set}}$ and then selecting Av and Tv such that $Av + Tv = Ev_{\text{target}}$. In `A` (aperture priority) mode, Av_{set} is manually set and $Tv = Ev_{\text{target}}$ - Av$_{\text{set}}$. The `S` (shutter priority) mode instead has the user manually select Tv_{set} and computes $Av = Ev_{\text{target}}$ - Tv_{set}. The fully-automated `P` (program) mode picks both the Av and Tv values for a a given Ev_{target}, most often based on identifying Tv_{blur} and Av_{limit} as respectively the slowest Tv for which hand-held motion blur is expected to be acceptable (often, based on the 1/focal length rule) and the highest Av such that image sharpness is not degraded by diffraction. An example of such a control algorithm is given in Algorithm 1. If the user specifies "auto ISO," generally similar logic is used such that Sv is increased to avoid Tv dropping below Tv_{blur}. Many low-end cameras lack an aperture control mechanism and some also lack a method for controlling shutter speed. In such cases, automatic ISO control is the only settable variable: $Sv = (Av_{\text{fixed}} + Tv_{\text{fixed}})$ - Bv.

Fundamentally, APEX exposure is based on a concept of equivalent exposures. However, a scene consists of not one, but many different light levels, so an equation specifying that the imaged intensity of one particular scene brightness is preserved is not really providing equivalent exposures. The following subsections discuss ways in which exposure computations can be made to consider more image qualities.

4.2 APEX Shift Exposure

Sv is really a compound value, with components determined primarily by analog gain and subsequent digital processing. Minimally, that justifies extending APEX with an additional equation:

$$Sv = Sv_{\text{analog}} + Sv_{\text{digital}}$$

The simplest way to implement ISO-less exposure control would be to apply the usual APEX exposure formulas, but with an optimal decomposition of Sv into Sv_{analog} and Sv_{digital} . Where there are clear image quality differences (as discussed in Section 3.2), Sv_{analog} should be selected as the value that generally will provide the optimum quality. In cases where the camera is truly ISO-less, Sv_{analog}should be Sv_{min}in order to minimize the probability of highlight clipping. The exposure would be made using Sv in the standard APEX formula, but the raw would be marked as needing a digital boost of Sv - Sv_{analog} and that same boost would

Algorithm 2 ISO-less exposure to maximize information content

```
capture a test image, $i$, with exposure parameters $p_{test}$
for ($p_j$ is each possible exposure parameter set) {
  $c_j$ = 0
  for ($v_{ref}$ is each pixel value in $i$) {
    $v_j$ = estimate_value($v_{ref}$, $p_{test}$, $p_j$)
    $c_j$ = $c_j$ + information_content($v_j$, $p_j$)
  }
}
$p_{best}$ = $p_j$ with maximum $c_j$
```

be applied in-camera before JPEG conversion. In fact, this boosting is precisely what we implemented using CHDK in Canon PowerShots – with great success, as described in Section 3.3.

While this approach easily extends APEX, it is possible to completely reformulate the exposure computation so that neither ISO nor Sv is treated as an input. The methods proposed in the following sections treat Sv, Sv_{analog} , and Sv_{digital} as outputs.

4.3 Maximizing Information Content

ISO sensitivity is measured by observing signal characteristics in an image captured using a known exposure. According to DxO,[6] the most common way to compute ISO is based on the exposure required to reach saturation – the maximum representable pixel value. This measurement method is closely related to the photographic concept of *exposing to the right* (ETTR)[12] to maximize image information content.[13]

Both CCD and CMOS sensors generally have the property that their output is approximately linearly proportional to the number of photons that deposited charge, not the log scale that human perception uses. Half the value range of the sensor is used to represent the brightest stop. For example, a 14-bit linear raw digital value's range is from 0 to 2^{14}-1, but values from 8192 to 16383 are all distinguishing brightness levels within a single stop. Very few levels within a darker stop are distinguished, thus potentially leading to visible posterization effects. Much finer tonal gradations can be recorded, and signal to noise ratio can be maximized, by setting the exposure such that the brightest pixels have the maximum value. The ETTR name comes from the fact that a histogram of pixel values in the image is thus skewed to the right (bright) side, rather than a centered distribution.

However, there are problems with ETTR. An obvious problem is that middle tones in the scene are likely to be recorded overly bright, requiring adjustment in later processing. Arguably, such adjustment does not work after the images have been recorded as JPEGs, because the data has suffered lossy compression and log encoding. Later adjustments generally reveal artifacts that would have been invisible if the JPEG were viewed without further processing.

A more fundamental, but less commonly recognized, problem is that *ETTR does not maximize the total information content in the image.* Suppose that the scene contains a tiny, but very bright, spot – for example, a reflection of the sun off a shiny surface near a window in an otherwise poorly lit room. By placing the value of that spot at saturation, the vast majority of the pixel values are pushed into the encoding of lower stops, and the total information content of the image is actually much less than if the spot had been allowed to clip. One can argue about the importance of letting some pixels clip, but in fact this importance can be approximately known by the camera prior to capture, so it can be used to intelligently decide how much clipping is beneficial.

Given an image, the total information content can be computed as the sum over all pixels of the information content of each pixel. Others have suggested vaguely similar analysis to select ISOs for multi-exposure HDR.[8] The information content of a pixel can be reasonably approximated by the signal/noise ratio of its value. The signal/noise ratio is not just a function of the value, but of the particular combination of analog and digital manipulation used to capture the image. That analog and digital manipulation is not merely adjusting ISO, but properly includes any processing that might be applied to improve image quality, such as noise reduction algorithms. Clipped (also called "blown") pixels carry no information, apparently contributing nothing to the total information content. However, there are post processing highlight reconstruction methods that can effectively recover the values of some clipped pixels, in which case the information content of a clipped pixel would really be an expected value for the accuracy of recovery. Alternatively, clipped pixels could be disproportionately discouraged by assigning them negative values.

The above definition of total information content can be directly used to drive an ISO-less exposure computation, outlined as Algorithm 2. A test image is captured using a reasonable guess at the exposure parameters, p_{test}, which could be initially determined using the usual APEX formulation. This image need not be a high-quality capture, but trivially can be a lower-resolution sampling from a live view. Note that in many cameras the live view image stream is coming from the JPEG processing pipeline, so some care must be taken to compensate for any transformations applied in that pipeline to obtain reasonably accurate v_{ref} pixel values. Then, for each possibly viable exposure parameter set, the algorithm computes the approximate total information content that the image would contain if captured using those parameters. This is done by estimating what the value of each pixel would be if the exposure parameters were to be changed from p_{test} to p_j, and then using pre-computed characterizations of the signal/noise ratio for each pixel level in an image captured with parameter set p_j. The exposure parameter set p_j that produced the highest total information content, c_j, is p_{best} – the best to use for the exposure. Of course, if the user does not trigger an exposure, this computation can repeated for the next live view image, in which case it would be natural to set p_{test} to produce an image slightly darker than p_{best} so that any clipped pixels might come back into the exposure computation with more meaningful values.

4.4 Maximizing Pixels Of Sufficient Quality

The previous section presented an ISO-less exposure scheme to maximize image information content. However, it may be more appropriate to *maximize the number of pixels of acceptable tonal quality.*

In Table 1, the dynamic range values listed were taken from the extensive collection of measurements published by DxO.[5] However, DxO is not the only organization conducting careful measurements of dynamic range, and the numbers obtained by various methods often differ by a stop or more. Saturation means that there is little disagreement about the top of the tonal scale, but how does one determine the darkest tone? Defining the darkest tone as that when adjacent tones can no longer be distinguished seems intuitive, but that does not yield a particularly useful measure. Instead, analysis software, such as Imatest,[9] bases dynamic range on measurement of noise. An appropriate threshold for noise could be argued to be anywhere between 0.1 and 1.0 stops, but exposure settings are usually made with precisions of 1/2 or 1/3 stop, so approximately that level of noise may be generally acceptable.

`Imaging-Resource.Com` reports Imatest scores for a variety of cameras, including a number of those listed here. For example, DxO quotes a 12.4-stop dynamic range for the Sony SLT-A55, which really represents a 1-stop noise threshold measured on a raw image scaled to a "print resolution" of 8MP. With a native 16MP raw, DxO and Imaging Resource agree on 11.1 stops with the same noise threshold, but Imaging Resource notes that dynamic range decreases to 10.9 stops with a 1/2-stop threshold, 10.5 stops with a 1/4-stop threshold, and 9.16 stops with a 1/10-stop threshold. DxO's data shows how dynamic range decreases as ISO is increased; the 11.1 stop range at ISO 100 has decreased to just over 6 stops at ISO 3200... and, of course, that value would go lower still with a smaller noise threshold.

Given this strong dependence on noise threshold, why not allow the acceptable noise threshold to be given directly as a user-settable parameter in determining exposure? The ISO-less exposure computation described in Section 4.3 and Algorithm 2 can be adapted for this with almost no change. It would be sufficient to modify the `information_content()` function to return 1 when the expected signal/noise ratio for the argument pixel value is above the user-specified threshold and 0 otherwise.

4.5 Zone System

Although the previous section takes a very different approach to exposure than is commonly discussed, the idea that exposure is really about obtaining acceptable tonal quality for the key parts of the image is nearly as old as photography itself. Perhaps the best-known codification of a method for preserving specific tonal qualities of the scene is Ansel Adams' *Zone System*.[2] Why not directly base exposure parameters on Zone System concepts?

Perhaps the most daunting aspect of the Zone System is that it is not so much about preserving tonal properties of the scene as it is about the visualized result:

> "The Zone System allows us to relate various luminances of a subject with the gray values from black to white that we *visualize* to represent each one in the final image."[2]

However, much of what the Zone System is doing is actually fairly mechanical. Fundamentally, there are 11 zones described, roughly corresponding to an 11-stop subject luminance dynamic range. The darkest (0) and lightest (X) zones serve as limits only. Zones I and IX might be distinguished by slight difference in tonality from 0 and X, but show no perceptible texture. Texture begins to become visible in zones II and VIII, and all the most important parts of the image will fall between these zones... but not all zones look right for all types of image content.

Dark foliage and landscape shadows are expected to fall in zone IV. A clear blue sky is in zone V. Snow ranges from zone VII to VIII with texture, and is in zone IX in flat sunlight. However, the most important content in many images is human skin, not only because humans are important subjects, but also because abnormal toning of them is more noticeable. Adams identifies Caucasian skin as normally falling in zone VI and never above zone VII; shadows on skin normally land in zone IV and highlights in zone VIII. Dark skin is placed in zone V. Exposing to ensure that skin tones are given consistently natural shades, free of noise artifacts, is a common theme in many exposure guides, such as the Wolfcrow System.[14]

The interesting fact is that most current digital cameras incorporate software that can recognize human faces in the live view feed. The faces are used to select focus point and aperture so that depth of field will cover the faces present, delay capture of an image until faces detected are all smiling, crop the image to a more interesting composition, and to bias the exposure. In fact, fields in the EXIF data indicate the bounding boxes for each of the faces found in each photo. It is a minor step to use the detected faces to select regions of interest that should be mapped into specific zones. Algorithm 2 can thus implement this exposure control by simply modifying the `information_content()` function to return 0 for non-face pixels and otherwise return a quality metric based on the probability that the pixel is assigned an appropriate tone. Of course, similar logic or user input (e.g., using a live-view touch screen) could identify target zones for other scene content, in effect providing exposure control via an ISO-less version of the Zone System.

5. CONCLUSION

This paper has presented an overview and preliminary investigation of the issues surrounding so-called "ISO-less" behavior. Some aspects of this behavior were found in every one of the 19 cameras tested. Despite that, the occurrence of literally ISO-less cameras is rare. The more common circumstance is that a set of multiple underexposing ISOs can all produce viable final images, but with slightly different qualities.

This realization, combined with the concept of an ISO consisting of separable analog and digital components, is then used to define several potential methods for computing exposure parameters in which ISO is not an input parameter: ISO-less exposure computations. The evaluation of these new exposure paradigms is future work.

REFERENCES

1. a1ex. *Dynamic range improvement for some Canon dSLRs by alternating ISO during sensor readout.* Magic Lantern Forum http://acoutts.com/a1ex/dual_iso.pdf, July 2013.
2. A. Adams. *The Negative.* Number v. 2 in The New Ansel Adams Photography Series. New York Graphic Society, 1980.
3. *Canon Hack Development Kit (CHDK).* http://chdk.wikia.com/wiki/CHDK.
4. Roger Clark. Sensorgen.info. *http://www.sensorgen.info*, 17 2014.
5. Camera Sensor Ratings by DxOMark. *http://www.dxomark.com/Cameras/Ratings*, 12 2013.

6. DxOMark measurements for lenses and camera sensors. *http://www.dxomark.com/About/In-depth-measurements/Measurements/ISO-sensitivity*, 1 2015.
7. J. Hannavy. *Encyclopedia of Nineteenth-Century Photography.* Taylor & Francis, 2013.
8. S.W. Hasinoff, F. Durand, and W.T. Freeman. Noise-optimal capture for high dynamic range photography. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 553–560, June 2010.
9. Imatest home page. *http://www.imatest.com/*, 1 2015.
10. *ISO 12232:2006 Photography - Digital still cameras - Determination of exposure index, ISO speed ratings, standard output sensitivity, and recommended exposure index.* International Organization for Standardization, 2011.
11. Douglas A Kerr. Apex-additive system of photographic exposure. *Issue*, 7(2007.08):04, 2007.
12. Michael Reichmann. Expose (to the) Right: Maximizing S/N Ratio in Digital Photography. *http://www.luminous-landscape.com/tutorials/expose-right.shtml*, 7 2003.
13. Michael Reichmann. Optimizing Exposure: Why Do Camera Makers Give Us 19th Century Exposures With Our 21st Century Cameras? *http://www.luminous-landscape.com/tutorials/optimizing_ exposure.shtml*, 8 2011.
14. Sareesh Sudhakaran. How to Expose and Grade S-Log2 for the Sony A7s. *http://wolfcrow.com/blog/how-to-expose-and-grade-s-log2-for-the-sony-a7s/*, 1 2015.