Image quality testing based on information metrics

Norman Koren, Imatest, LLC,

Boulder, Colorado 80301, USA

May 16, 2026

A concise guide to camera image quality testing, primarily for machine vision, focusing on metrics derived from information theory, which are more predictive of camera performance than traditional metrics such as sharpness or noise.

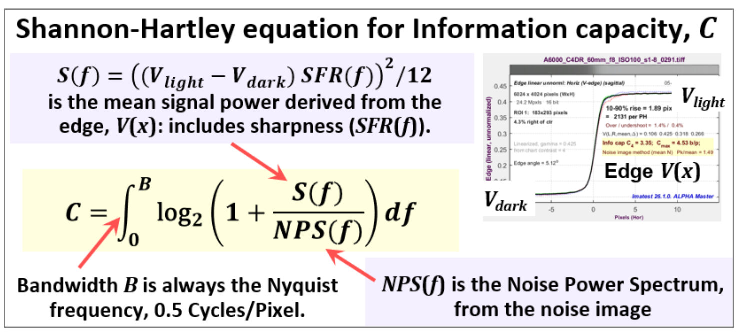

Information Capacity, *C*, is the key performance metric derived from information theory. It is calculated from sharpness, noise, and image amplitude.

C4 is the information capacity measured directly from 4:1 slanted edges. It is the maximum information per pixel for an object with *4*:1 contrast.

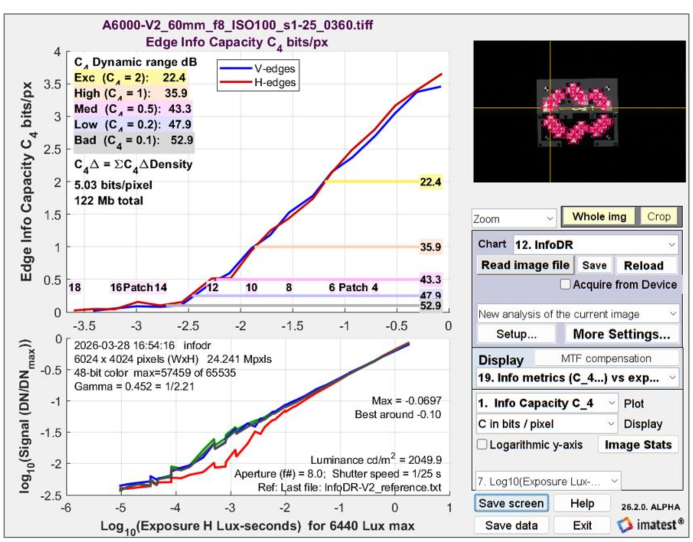

[*SNRi*](#), ideal observer *SNR*, is the key metric for object detectability, calculated from *C4*, object size, and image processing.

This plot, made from the new Information-based Dynamic Range (InfoDR) test chart, shows *C4* for a wide range of illumination. it is a key result of a new approach for measuring performance over a wide range of illumination, including low light and dynamic range, based on information theory.

Table of Contents

Most references are links, which can be opened with control-click. There are a few traditional IEEE/IS&T [*n*] references at the end of the paper. If you are reading this from a hardcopy, most can links be found by searching www.imatest.com/docs/. A few are on Wikipedia, www.wikipedia.org/.

We occasionally use "**green for geeks**" boxes, similar to the ***Imatest*** website to indicate technical material that can be skipped by readers who don't want to dive down deep rabbit holes.

Introduction

Suppose you are an engineer, tasked with selecting a camera to be embedded into a device — ***any*** device — a vehicle, a medical endoscope, a security camera, etc. What test charts should you acquire, what tests should you perform, and how should you interpret the results?

Or suppose you are a reviewer for a publication. You have similar questions, but you might expect the answers to be somewhat different.

This document is designed to help guide you through the process of choosing the right measurements, which can resemble a very deep rabbit hole. It will be mostly high-level, i.e., it will only have a few complex equations, usually in ***green boxes*** that non-technical readers can skip. There will be lots of links to online documents with greater depth. Although this is something of a primer or crash-course on image quality testing, it will include recently-developed measurements based on information theory that are still relatively unfamiliar, but which we believe are superior for evaluating camera performance.

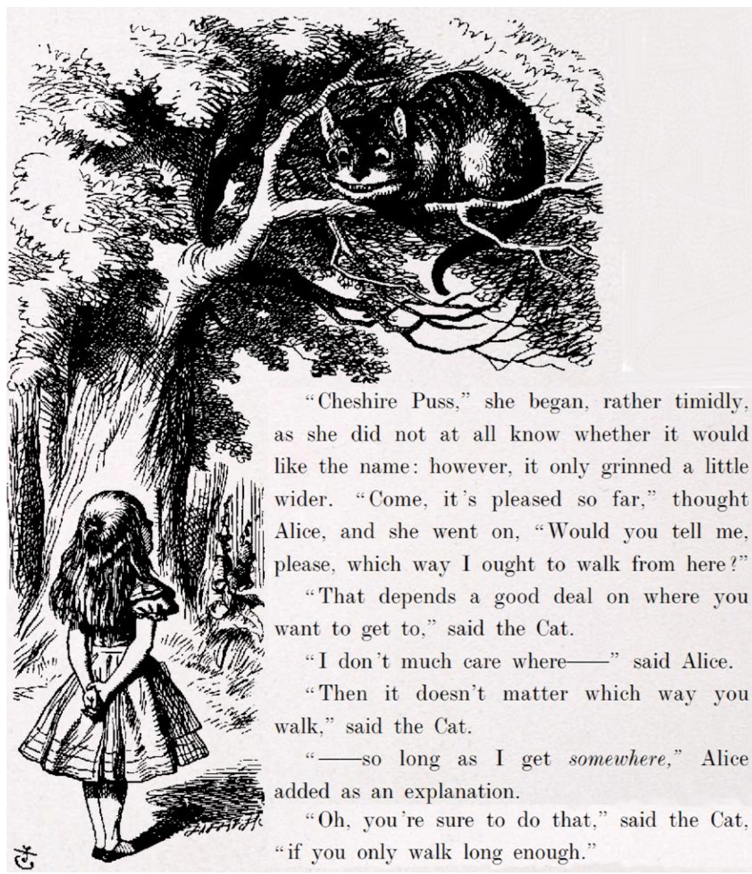

"Cheshire Puss," she began, rather timidly, as she did not at all know whether it would like the name: however, it only grinned a little wider. "Come, it's pleased so far," thought Alice, and she went on, "Would you tell me, please, which way I ought to walk from here?"

"That depends a good deal on where you want to get to," said the Cat.

"I don't much care where——" said Alice.

"Then it doesn't matter which way you walk," said the Cat.

"——so long as I get *somewhere,*" Alice added as an explanation.

"Oh, you're sure to do that," said the Cat, "if you only walk long enough."

Prerequisites — The reader should be familiar with the basic concepts of photography: exposure time, aperture (f-stop), total exposure, noise, etc. If you are new to photography, we recommend Sean McHugh's excellent tutorials, [Cambridgeincolour.com Photography concepts](). If you are an engineer, you should know something about linear systems and what a Fourier transform does (converts signals between spatial and frequency domains).

Note that this document has minimal instructions for running ***Imatest*** modules. More can be found in the **Imatest Documentation** and in the many links.

Basic imaging concepts

Many readers can skip this section and go directly to **Information theory** or **Camera characterization**, which is the heart of this document. We will keep this section as concise as possible.

This section may also be viewed as a guide to the **Imatest Documentation** page, which can be challenging to navigate. It is also as the outline of a hypothetical textbook that we don't plan to write.

Camera

A digital camera is a device that consists of

- a lens that focuses light from a scene (sometimes called an object) on an image sensor (an *m* x *n* array of photosensitive pixels or photosites),
- an image sensor, which converts light on the photosites into electrical signals and also adds noise to the image. The signal is usually digitized on the sensor. The unprocessed digital output of the image sensor is called a ***raw*** image.
- image processing (often abbreviated ***ISP*** for ***Image Signal Processing***), which converts the raw image into an interchangeable image format that can be used and interpreted outside the camera. ISP also enhances the image to be more pleasing for human vision or more useful for machine vision.
- Storage: digital memory cards (photographic film in the old days).

The *m* x *n* pixel size is sometimes referred to as the camera's "***resolution***", but it should not be confused with measured resolution measurements, such as "vanishing resolution." The meaning is usually clear from the context.

Digital image

A digital image is an *m* x *n* x *k* matrix, consisting of *m* rows (height in pixels), *n* columns (width in pixels), and *k* (usually 1 or 3) colors. A ***pixel*** (*picture element*) is a single element of the image, typically consisting of one, two, or more bytes (8, 16, or more ***bits*** (*binary digits*), where one byte is 8 bits). An 8-bit image can have 2^8 = 256 levels (0-255). Similarly, a 16-bit image can have 2^{16} = 65536 levels (0-65535), etc. ***High Dynamic Range*** **(*HDR*)** images usually have bit depths ≥ 16.

Raw image, raw conversion, and demosaicing

Bayer CFA

Raw images directly out of image sensors can have bit depths that are not multiples of 8 (10, 12 and 14 are common). They are generally structured as monochrome images (*k* = 1). For color image sensors, where the photosites are covered by a ***Color Filter Array* (CFA)** (shown on the right for the popular *Bayer CFA*), each pixel represents one color.

Converting a raw image into an interchangeable format (which is associated with a ***color space***) involves several operations.

- ***Demosaicing***: For color image sensors, convert each individual pixel, which represents a single color, into three pixels representing red, green, and blue (R, G, and B) colors (though there are other configurations). Can be quite mathematically sophisticated.
  There are other CFA configurations, for example, RGB-IR (IR shown in **dark gray**), which is normally converted into separate RGB and IR images for analysis.

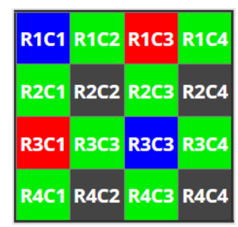

RGB-IR CFA

- Apply a **gamma** curve along with a tonal response to the (usually) linear raw image. $Digital\ Number\ (DN) \mathrel{\hat{=}} illumination^{(\text{encoding gamma})}$.
- **Sharpening**, which is typically visible on edges. It boosts high spatial frequencies.
- **Noise-reduction**, which can be
  - uniform (also called *lowpass filtering*, which degrades sharpness and texture by attenuating high spatial frequencies), or
  - edge-preserving (nonuniform; also called *bilateral filtering*) [6], which lowpass filters relatively smooth regions away from sharp edges, but leaves edges alone. Almost universal in JPEG images from consumer cameras. Makes the image more visually pleasing at the expense of fine texture.
- **White-balance and color correction**. Color correction is often accomplished by applying a 3x3 ***Color Correction Matrix (CCM)*** to each 3-element RGB pixel.

Sharpening and the other image processing operations can also be applied later in the image processing pipeline. See Sharpening and Gamma, Tonal Response, and related concepts.

Sharpening, which boosts medium to high spatial frequencies, is almost always beneficial for human vision (if not excessive), but tends to be less useful for machine vision because human vision has two additional components: the display and the eye, that have losses that need to be compensated by sharpening.

The output of raw conversion is an *interchangeable image file*. ***Imatest*** (MATLAB) supports several types. The most common are

Uncompressed	**BMP, TIF**	May require large amounts of storage
Lossless compression	**TIF, PNG**	Typically, around half the size of uncompressed images
Lossy compression	**JPG, GIF**	

JPEG file compression (quality) is variable, set when saving the file. High-quality JPEGs are suitable for most ***Imatest*** analyses (except for dynamic range), but low-quality JPEGS have blocky artifacts that make them unsuitable. Many of the issues observed with analyzing JPEG files from consumer cameras are caused by strong image processing applied during raw conversion, especially bilateral (nonuniform) filtering [6], ***not*** by JPEG compression itself.

Gamma encoding and Tonal Response Curve

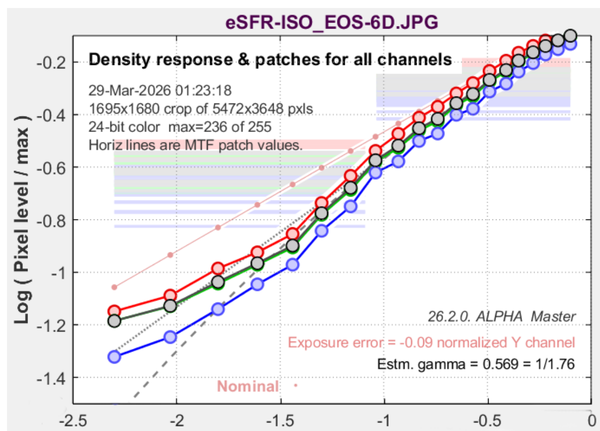

Density (tonal response) plot

Raw images are almost always linear, i.e., their digital numbers (*DNs*) are proportional to the pixel illumination.

Most interchangeable images (TIFs, JPEGs, etc.) are encoded with a gamma of around 1/2.2, i.e.,

$$Digital\ Number\ (DN)\ (\text{synonymous with}\ Pixel\ Level) \triangleq illumination^{(encoding\ gamma)}$$

There are several reasons for this. In the old days of Cathode Ray Tubes (CRTs), brightness was proportional to $(grid\ voltage)^{(display\ gamma)}$, where display gamma was between 2 and 2.5, so it made sense to encode the signal with the inverse, around 1/2.2 = 0.454. Also, as explained in Gamma, Tonal Response..., gamma-encoding increases the effective dynamic range, especially for files with bit depth = 8.

The full encoding curve, called the tonal response curve, is often a little more complex than a straight-line gamma curve. The curve shown above has a region of reduced slope in the highlights, called a "shoulder," that helps to control highlight "burnout" in pictorial images.

Images need to be linearized, i.e., gamma-encoding needs to be removed, for *SFR* or *MTF* to be calculated. The linearization can be approximate.

For more information, see **[Gamma, Tonal Response Curve, and related concepts]**.

Color Space

A ***Color space*** is a mapping between an image's *Digital Numbers* (*DNs*), most often in RGB format, and device-independent colors, generally expressed as ***CIELAB (L*a*b*)*** Values. Color spaces also have associated display gammas, most often approximating 2.2, and color gamuts that quantify the range of colors that can be represented.

Raw images don't have a color space, but all interchangeable images have one, either implied or embedded.

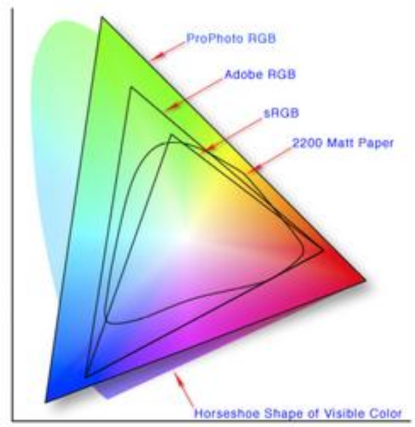

from Jeff Schewe via Wikipedia

Common color spaces

sRGB is the standard color space for still images in Windows and the internet. It is a relatively small color gamut color space with gamma approximating 2.2 (consisting of a small linear region and a region with gamma = 2.4).

Adobe RGB is a medium gamut color space with gamma = 2.2, often used for fine art printing, where a slightly higher color gamut is desirable.

Several other color spaces are commonly used for video and cinema.

Dynamic Range, Stray Light, and Tone Mapping

Dynamic range is the range of illumination over which a camera has good contrast and good Signal-to-Noise Ratio (SNR). It is described in depth in the Dynamic Range page. It can be measured from transmissive grayscale charts, including the three described in Tonal measurements, below: the ***Imatest*** 36-patch Dynamic Range chart, the Contrast Resolution chart, and the new and highly accurate InfoDR chart.

It is important to distinguish between ***camera*** and ***image sensor*** dynamic range. Although some image sensors have extremely high dynamic ranges (as much as 150 dB, measured from flat patches at a succession of light levels), *camera* dynamic range is limited by ***stray light***, also called flare light or veiling glare, that originates with light reflected from lens surfaces or the interior of the lens barrel. Stray light can be reduced by high-quality lens coatings, which reduce the light reflected from a glass-to-air surface from 4% to at best 0.4%, but it can never be eliminated entirely. Read more about stray light here.

Although stray light has little effect on sharpness (*SFR* or *MTF*) or noise, it has a strong effect on C_4 information capacity, as described below.

Because High Dynamic Range (HDR) images cannot be displayed well on standard displays (monitors or prints), they are often [locally tone-mapped] — a highly nonuniform process that maintains local contrast, but lightens dark areas enough to be visible. This can make tonal response and dynamic range measurements from standard grayscale charts quite unreliable. Local tone mapping can often be recognized by very low measured values of gamma (< 0.25; well below the nominal value of 0.45 for most common color spaces). ***Imatest*** designed the Contrast Resolution chart, to produce useful results with locally tone mapped images, but for accurate dynamic range measurements, images with nonlinear or nonuniform processing should be avoided where possible.

Spatial and frequency domains and sharpness: *MTF* & *SFR*

Sharpness, illustrated on the right, is a familiar concept. The upper image is sharp; the lower one is blurred.

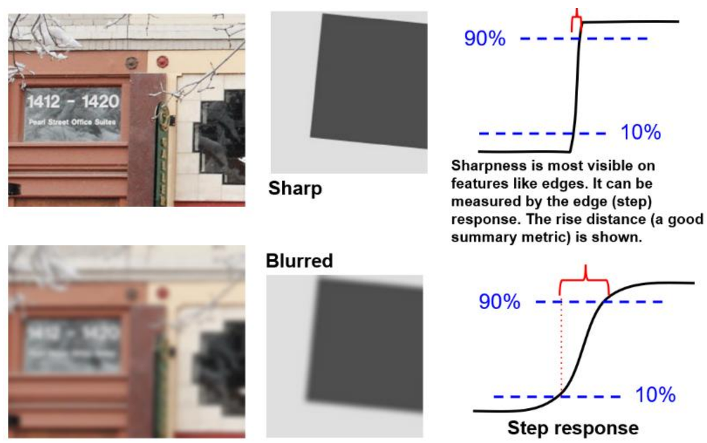

Sharpness and step response

Sharpness can be characterized by a ***step response***. For a camera system that consists of a number of cascaded components, each of which has its own step response, the total system response can be calculated by a process called *convolution*, which is complex, slow, and hard to visualize.

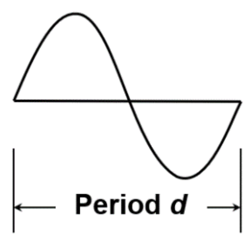

The overall response of a linear system is much easier to calculate in ***frequency domain***, where frequency $f = 1/(\text{Period } d)$. Spatial-domain (real-world) signals are converted to frequency domain for *MTF* analysis by the *Fast Fourier transform*.

The sharpness of a component or a system can be characterized by its ***Spatial Frequency Response (SFR(f))*** or ***Modulation Transfer Function (MTF(f))***. *SFR* and *MTF* are used interchangeably. The *SFR* of an individual component, such as a lens, is the Michelson Contrast of a sine wave of frequency f at its output, $(V_{light} - V_{dark}) \; / \; (V_{light} + V_{dark})$, normalized to 1 at $f = 0$. The total system response can be calculated by multiplying the frequency response of each component.

Frequency response, *SFR*(*f*) or *MTF*(*f*), can be visualized with the sinusoidal pattern below, where the upper half represents the input and the lower half represents the output, whose amplitude decreases with spatial frequency, but they are calculated by much more efficient methods.

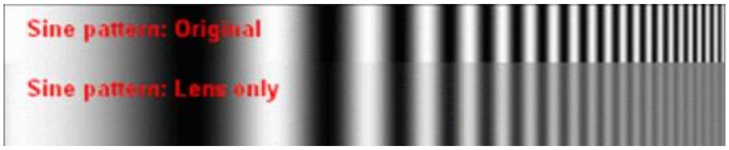

Sine wave response. Upper: input, Lower: blurred output

For the widely-used slanted-edge test pattern [5], *SFR* is the absolute value of the ***Fourier transform*** of the derivative of the step response.

SFR(*f*) and *MTF*(*f*) have native frequency units of *cycles per pixel*, but can be expressed as

- cycles (line widths) per picture height,
- cycles per distance (mm or inches), on the image sensor or scene (object)
- cycles per angle.

MTF and *SFR* are functions of frequency, f, normalized to 1 at $f = 0$. They are plotted along with several summary metrics in the Edge & MTF plot (right). Some commonly used summary metrics are

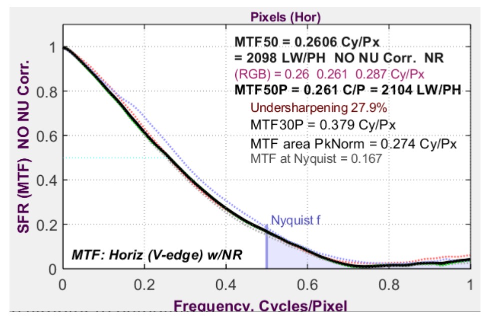

MTF(f), synonymous with SFR(f)

- *MTFnn* (often *MTF50*) — The spatial frequency where *MTF* drops to *nn*% (50%) of its value at $f = 0$. Equivalent to *bandwidth* in electrical engineering. Strongly affected by software sharpening.
- *MTF50P* — The spatial frequency where *MTF* drops to 50% of its *peak* value. Slightly less affected by sharpening than *MTF50*.
- *MTF area (peak normalized)* — The area under the *MTF* curve below the Nyquist frequency, $f_{Nyq} = 0.5$ cycles/pixel (the highest frequency where digital information is correctly conveyed). Relatively insensitive to large amounts of sharpening.

Although more extended *MTF* (higher *MTF50*) implies better sharpness, there is a limit for sampled systems. If there is significant energy above the maximum frequency where the input signal can be reconstructed — the *Nyquist* frequency, $f_{Nyq} = 0.5$ cycles/pixel, an effect called ***aliasing*** causes low frequency artifacts such as Moiré fringing, which can be disturbing for repetitive patterns.

This large subject is covered in **Sharpness** – What is it and how is it measured?

Slanted edge test charts

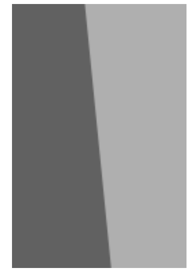

Slanted edge

Slanted edges are ***Imatest's*** most popular pattern for calculating $MTF(f)$ and related metrics, but before we describe them, we must answer two questions.

Why are the edges slanted? They are slanted because if they were perfectly vertical or horizontal, i.e., if they were aligned exactly with the image sensor pixel boundaries, $MTF(f)$ would be highly sensitive to the sampling phase (the position of the edge with respect to the pixel boundaries), resulting in inconsistent measurements. Slanting the edges ensures that they will have multiple sampling phases, resulting in much more stable measurements. The angle is not critical. The standard calls for a 5 or 5.7 degree (arctan(0.1)) tilt, but any tilt between 2 and 8 degrees will work.

Why 4:1 contrast? The original ISO 12233:2000 release called for a contrast ratio between 40:1 (about the maximum matte media can attain) and 80:1. The trouble with these values is that

(1) it was very easy to saturate the image, even with very little exposure error, and saturation severely distorts the $MTF(f)$, making it look *better* than reality,
(2) the value of gamma used to linearize the image had to be estimated accurately,
(3) the high contrast tends to maximize sharpening, especially for bilateral-filtered images [6], which are common in consumer cameras, and
(4) the contrast is much higher than most of the real-world objects that need to be detected.

ISO 12233:2014 (and succeeding versions) corrected this by specifying a 4:1 contrast ratio, which is essentially a *very good* compromise.

(1) it is difficult to saturate, even with moderate exposure error (though we once saw some images that were intentionally overexposed to boost MTF).
(2) The estimate of gamma doesn't have to be very accurate. You can make an excellent estimate from the measured edge contrast if the chart contrast is known.
(3) Signal-to-Noise Ratio, SNR, is decent down to relatively dim light. Although the low-light *SNR* of 2:1 contrast charts is relatively poor, we are exploring potential uses.

Other advantages of slanted edges:

- Compact: small size enables MTF to be mapped over the image surface with good detail. Summary metrics like $MTF50$ or $MTF50P$ are typical displays.

- Fast calculations
- Relatively (though not completely) insensitive to noise

Additional charts for measuring *MTF*, texture, and vanishing resolutions are described **here**.

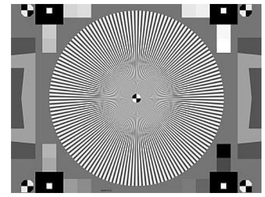

Siemens star

Noise

Noise consists of random perturbations of the image. It is traditionally measured from flat patches, and can be characterized by an amplitude, which is equal to the standard deviation of the signal when illumination is uniform, and a frequency spectrum, also called the ***Noise Power Spectrum*** **(*NPS*)**.

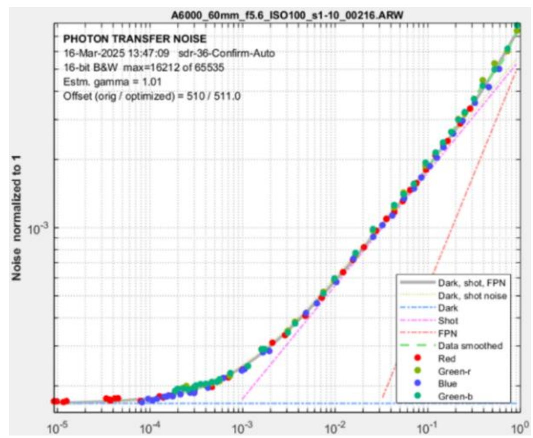

Photon Transfer Curve (PTC), illustrating image sensor noise as a function of signal, obtained from a raw image.

Imatest has developed a sophisticated image sensor noise model, based on the classic *Photon Transfer Curve (PTC)* [14], which displays noise as a function of exposure. The PTC is measured from totally raw (undemosaiced) images, for use in the ***Simatest camera simulator***, which predicts camera performance based on traditional and information metrics. See Image Sensor Noise – Measurement and modeling. As of April 2026, the model does not yet include High Dynamic Range (HDR) sensors, which have multiple operating regions.

Two broad types of noise are

- temporal noise, which is random and different for each exposure. It is gaussian or Poisson-distributed (for photon shot noise with small numbers of photons).
- Fixed pattern noise, which is always the same and often has a fabric-like repetitive pattern (it's not gaussian). There are two types, described in the EMVA 1288 standard: Dark Signal Nonuniformity (DSNU) and Photo Response Nonuniformity (PRNU).

Charts and lighting

Although inkjet chart files can be generated with the ***Imatest*** Test Charts module (this page for SVG charts), for the best results, we recommend the charts in the ***Imatest store***), which has a huge selection, including many charts that must be printed on media other than inkjet, such as film charts for Dynamic Range measurements.

Once you are clear about what tests need to be performed, choosing the right test charts might seem relatively straightforward, but we strongly recommend contacting our sales engineers for assistance. Ordering the wrong chart can result in costly delays. (Our sales engineers frequently contact customers to confirm that they are ordering the charts they really need.)

The two broad categories of test chart are reflective and transmissive. The choice of chart depends on

- the me*asurements needed,*
- the camera resolution (*m x n* pixel size),
- the field of view (i.e., the physical size of the chart),
- the lighting requirements, and
- whether the chart needs to be used for Infrared (IR)).

The operating limits of several test chart media is contained in Test chart suitability for MTF measurements.

- **Reflective charts —** Inkjet charts can be printed very large, but are too coarse for small sizes, though a few specialized small charts (non-inkjet) are available. Maximum optical density, D_{max}, is around 1.6 (40:1 contrast ratio) for matte charts, which don't have specular reflections, and 2.2 for semigloss charts, which can have serious specular reflections, especially for wide angle lenses. (Sometimes lighting can be arranged so specular reflections don't fall on measurement areas.) Finer photographic charts are available.

  Lighting: Several reflective lighting systems are available from the ***Imatest Store***, most notably the excellent Kino Flo lights.

- **Transmissive charts —** can be finer than the finest reflective charts. They include
  - inkjet (coarsest; translucent),
  - photographic film (B&W or color; transparent; higher D_{max}, (3.0 or larger) than reflective charts; for IR, color is transparent but B&W is OK). Often called "LVT" charts because they are made with the "Light Valve Technology" (laser-drum) process,
  - photomask (halftoned Black & White; very fine, visible and IR), or
  - chrome-on-glass (the finest; suitable for microscopic applications; only available in two-tone 10:1 contrast).

  Because film and photomask charts cannot be manufactured with the precise tones or colors, they are supplied with CSV reference files that contain patch densities (log$_{10}$(transmittance τ) for monochrome (B&W) charts or [CIELAB (L*a*b*)](#) values for color charts. The reference files must be entered into ***Imatest*** for reliable results.

Lighting: we recommend [Lightboxes or Light Panels](#) from the ***Imatest Store***.

Information theory and metrics

Image Information Metrics contains a more detailed introduction, with numerous links.

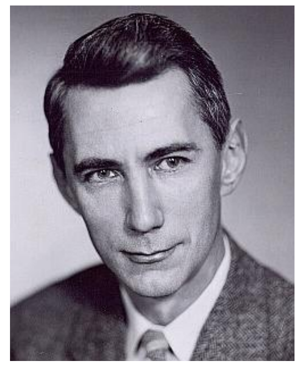

Claude Shannon's ground-breaking work on information theory [1,2] is the basis for calculating information capacity and related metrics from the widely-used slanted-edge test pattern, described in several recent papers from ***Imatest*** [3,4] and summarized in the introductory web page, **Image Information Metrics**.

In electronic communications systems, channel (information) capacity, C, is the maximum rate in bits per second that information can be transmitted through a channel without error. For additive white gaussian noise, it is given by the deceptively simple Shannon-Hartley equation.

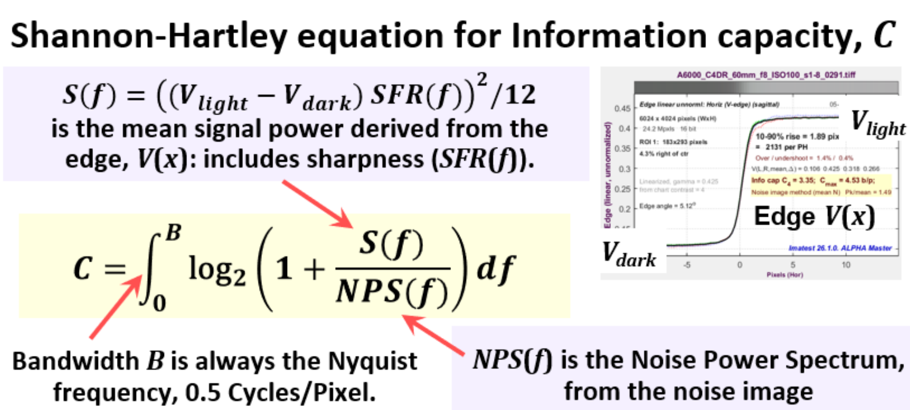

Imaging systems are communication channels where

- Information capacity, C, has native units of ***information bits/pixel*** (though bits/image, bits/distance, bits/angle, etc., can also be calculated). *Information bits* should be not be confused with bits as units of storage (as in image size).
- Because C depends on test chart contrast, we define C_n as the information capacity measured directly from edges with n:1 contrast. C_4, which is measured directly from the ISO 12233-standard 4:1 contrast edges [5] used in most ***Imatest*** slanted-edge

test charts, is the amount of information that can be contained in an object with a 4:1 contrast ratio.

- $S(f) = \left((V_{light} - V_{dark})\ SFR(f)\right)^2/12$ is the mean signal power, where
  - V_{light} - V_{dark} is the signal amplitude, i.e., the difference between the mean Digital Numbers (*DNs*) of the two sides of the slanted edge. It is affected by stray light.
  - $SFR(f)$ or $MTF(f)$ is the spatial frequency response derived from the edge.
  - *T*he denominator 12 scales $S(f)$ to be the mean value of signal power for uniformly distributed amplitudes between V_{light} and V_{dark}, which maximizes C.

Calculating Spatial Frequency Response, $SFR(f)$, from slanted edges has always been straightforward (using the ISO 12233 standard), but combining it with $NPS(f)$, which was traditionally calculated separately in flat regions, was cumbersome and error-prone. The breakthrough that enabled the convenient calculation of C was the discovery of how to measure signal and noise at the *same* location (which was anything but obvious). References [3] and [4] describe the algorithm in detail.

$S(f)$ is a function of image contrast (V_{light} - V_{dark}), that has been generally ignored in image quality calculations. This is a significant omission because (V_{light} – V_{dark}) is degraded by stray light from dust or dirt inside the camera or on the lens. (A good example is dried salt spray, which is common on lenses in climates where roads are salted to melt snow.) Information metrics, especially C_4, rectify this omission.

We are particularly interested in C_4, which is the maximum information capacity that can be conveyed by an object with a 4:1 contrast ratio, which is typical of objects (e.g., neutral-colored cars or clothes on neutral backgrounds, such as concrete) that need to be detected in important applications, such as automotive imaging. C_4 is also the key metric for performance as a function of illumination, measured with the **InfoDR** chart.

Because C_4 is a strong function of signal level, a standard signal level should be used when reporting it. A good choice is to set the exposure so the normalized edge level, $L_{edgeNorm} = DN_{edgeMean}/DN_{max} = 0.18$, where $DN_{edgeMean}$ is the mean linearized Digital Number of both sides of the edge and DN_{max} is the maximum Digital Number, typically $2^{(\text{bit depth})}$-1. 0.18 (18%) is sometimes considered “neutral gray.” We label C_4 measured at this level, $C_4(0.18)$. [We considered 50%, but it was too close to the maximum level, (100%+25%)/2 = 62.5%, for an image where the bright patch is saturated.] $C_4(0.18)$ is illustrated in InfoDR results, below.

The key point is that **information capacity C is a function of three factors**:

(1) ***image contrast***, $V_{light} – V_{dark}$ (or *Michelson contrast*, $(V_{light} - V_{dark}) / (V_{light} + V_{dark})$),
(2) **sharpness**, $SFR(f)$, and
(3) **Noise Power Spectrum**, $NPS(f)$,

making it a ***<u>complete</u>*** image quality metric with units of information bits per pixel, in distinction to the three individual factors, which are *partial* metrics. For comparing cameras, C should be measured at a standard normalized digital level.

[Note that this statement omits other important metrics, such as color and optical distortion, that don’t

References [3-4] describe techniques for calculating $NPS(f)$ and $S(f)$ from the *same* slanted-edge location, making the calculations fast and robust. They also define additional metrics such as *SNRi* (Ideal observer SNR) [4,8-10], which quantifies how well an object of a given size can be detected.

These calculations work best with minimally or uniformly processed images. They are less accurate, though still of interest, for JPEG images from consumer cameras, most of which have been bilateral (nonuniformly) filtered [6].

***The real advantage of information capacity and related metrics* is that they directly answer the question, “How good is this image?”** Sharpness, which is typically measured as *MTF50* (in units of cycles/pixel, cycles/image height, or cycles/distance) doesn’t fully answer the question.

> “Well! I’ve often seen a cat without a grin,” thought Alice; “but a grin without a cat! It’s the most curious thing I ever saw in all my life!”

Significant interpretation is required for MTF measurements, especially since *MTF50* and related measurements are affected by software sharpening, which adds no information to the image. The same holds for Signal-to-Noise Ratio, *SNR*. For a flat patch, SNR = 20 dB (10:1) looks pretty good, but 0 dB (1:1), which is commonly used to define Dynamic Range, looks *terrible*. Visual appearance offers limited clues about how traditional metrics relate to object detection.

Imatest users will need to determine the values of information metrics, usually C_4 and a related metric such as *SNRi*, that need to be specified to meet the performance requirements of their system. Once specified, these values should be more stable and robust than *MTF* or *SNR*. That is why we are confident that even though information capacity is unfamiliar and challenging to visualize, it will become a key image quality metric.

A paper by [Dairmaid Geever et. al.](#) of the University of Limerick, [“Information Capacity as a Predictor of Perception Performance”](#) [13] contains the first results that correlate machine vision performance with information capacity.

Imatest is actively working on **ISO 23654**, “Digital imaging — Image information metrics,” which will describe how information capacity and related metrics, including C_4, are calculated and interpreted. Expected publication date is November 2028. New committee members are welcome.

SNRi — Ideal Observer SNR

Of the metrics derived from information theory, ***Ideal Observer SNR, SNRi***, which quantifies how well an object of a given size can be detected based on Bayesian statistics [4,8-10], is one of the most useful. Unlike C, $SNRi$ is affected by image signal processing (sharpening, etc.).

SNRi is derived from an integral that includes the Fourier transform of the object, $G(\nu_x,\nu_y)$, for spatial frequency, $\nu = \sqrt{\nu_x^2 + \nu_y^2}$.

$$SNRi^2 = \iint \left(\frac{\left|G(\nu_x,\nu_y)\right|^2 SFR^2(\nu_x,\nu_y)}{NPS(\nu_x,\nu_y)} \right) d\nu_x \, d\nu_y$$

SNRI^2 is equivalent to the total (integrated) noise-whitened Signal/Noise energy of the object in the spatial domain.

SNRi is typically calculated from rectangles of size *w* × *kw* pixels, where *k* = 1 for a square or 4 for a 1:4 aspect ratio rectangle, as illustrated on the right (for *w* × 4*w*).

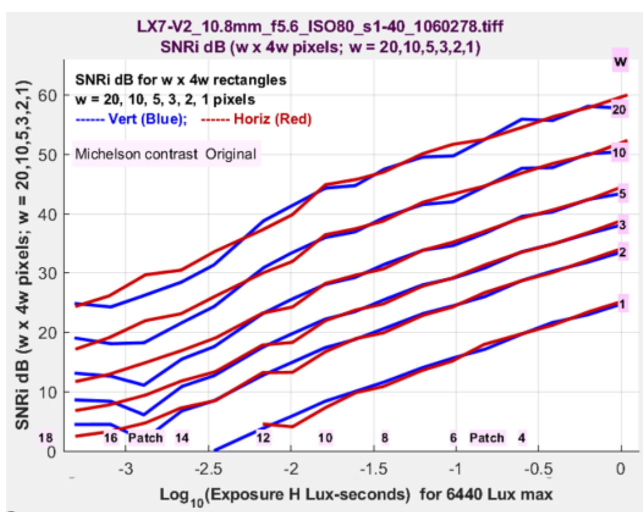

SNRi for w x 4w rectangles with w = 1, 2, 3, 5, 10 & 20 pixels for compact camera.

It’s important to note that *SNRi* for *w* × *kw* pixels only indicates the *presence* of an object (detection). To *identify* it, i.e., to distinguish a car from a boat from a person, much smaller sizes should be used: roughly 1/10 the total object size. (This needs more research.) *SNRi* can be plotted for several values of *w* (1, 2, 3, 5, 10, and

20 pixels) for a wide range of exposures and object contrasts (based on Michelson contrast) using the **InfoDR chart**, as shown here.

Performance specifications — *more questions than answers*

The introduction of image information capacity and related metrics, especially *SNRi*, is a paradigm shift that offers strong potential benefits but comes with a number of challenges. A major potential benefit is improved test results, with fewer false positives and false negatives, saving time and money. But this requires determining suitable pass/fail thresholds, i.e., setting the specifications ("specs"), which requires effort and will need to be communicated to customers.

Better specifications, i.e., improved pass/fail criteria based on information metrics such as *SNRi*, produce more reliable test results with fewer false positives or false negatives, leading to *significant savings in time and money*.

Determining pass/fail thresholds for traditional metrics, such as sharpness and noise (or *SNR*) may well be ***more*** complex than for a single information metric such as *SNRi*, where the selected object size (a fraction of the *total* size) and the minimum specified value depends on the task to be performed.

Typical tasks can be described in order of increasing difficulty as

- Detection ("Something is present"),
- Recognition ("It's a car, not a tractor, bicycle, or boat."),
- Identification ("It's Tom Magliozzi's 1963 Dodge Dart").

The size used for identification is significantly smaller than the total object size because relatively small features (a person's hand, face, shoe, etc.) enable objects to be recognized. How much smaller will require research.

An additional task (useful for determining motion in automotive imaging), is to find the *position* or *location* of the object. [It's not quite ready for prime time.]

Some older summary metrics used for pass/fail, such as *MTF50*, may be familiar, but can be *highly* misleading because they are strongly affected by sharpening and bilateral filtering, which can make sharpness numbers arbitrarily high without improving object detectability. They can even degrade performance by boosting artifacts like noise. See the paper, Correcting Misleading Image Quality Measurements [15].

Avoid setting pass/fail thresholds based on "vanishing resolution," which can be measured from bar patterns such as wedges. Vanishing resolution was a valid measurement for optical systems, but is unrelated to the information content of imaging systems. To paraphrase Winnie-the-Pooh, it's "where the woozle wasn't." Such measurements include *MTF10*, *MTF10*, and *MTF20*, which can be boosted by noise and by demosaicing algorithms for edges (though less for sinusoidal patterns). *MTF50*, *MTF50P*, and related metrics are better because they correspond to the system bandwidth and are less sensitive to noise and artifacts. And information-based metrics, to be described below, are ***much*** better.

Error probability (Bit Error Rate)

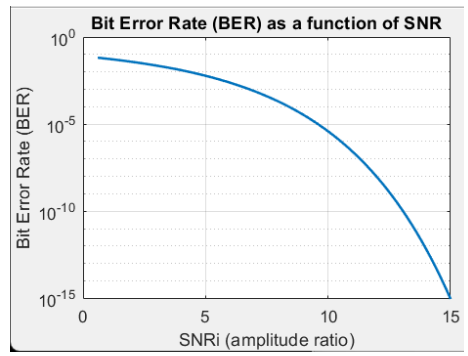

BER = 1/2 erfc(SNRi)

A key claim of papers that describe *SNRi* [7-10] is that it can be used to derive Bayesian statistics for object detectability. But the equations for doing so are missing from the papers. So, in May 2026, we took the audacious step of deriving them in hopes of finding better ways of setting measurement specifications. Starting with *SNRi* and following the steps described in the **[Appendix]** and in [16], we calculated the key metric of detectability for rectangular objects of various sizes: the *Error probability,* also called *Bit Error Rate*.

$$BER = \frac{1}{2}\operatorname{erfc}\left(\frac{SNRi}{2\sqrt{2}}\right)$$

Then we added plots of $\log_{10}(Bit\ Error\ Rate) = \log_{10}(BER)$ for *w* x *w* square or *w* x *4*w rectangular objects, where *w* = 20, 10, 5, 3, 2, and 1 pixel. Log_{10} was used because of the extreme range of error rates. Since low contrast objects may be important, we added the option to display $\log_{10}(BER)$ for several object contrasts, mostly lower than the default 4:1 chart contrast, noting that *SNRi* is proportional to Michelson contrast $C_{Mich} = (C_R-1)/(C_R+1)$ for contrast ratio C_R:1.

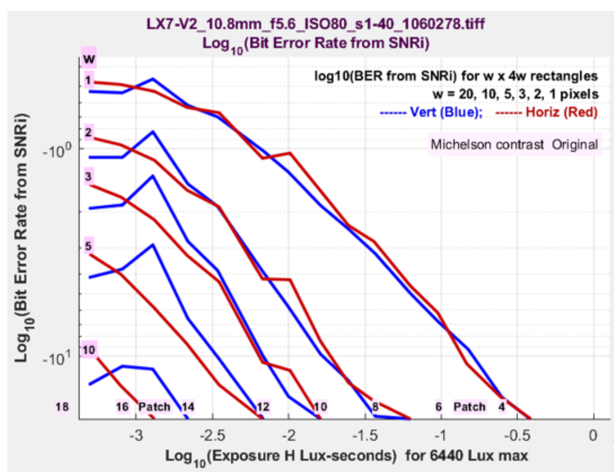

Log_{10}(Bit Error Rate, BER)

What makes *BER* uniquely useful is that it can be compared *directly* with actual measured machine vision/AI performance, i.e., the object recognition error rate. Research is needed to correlate *BER* to measured error rates, which may be dominated by *SNRi* for rectangle sizes smaller than the total object size (for example, hands or faces in a person, which are around 1/10 the person's size). Contrast ratios can also be lower than the ISO 12233 standard C_R = 4:1 (C_{Mich} = 0.6). We will try to find the relative sizes and contrasts that best correlate with measured *BER*, and if things go really well, we may be able to define a figure of merit for machine vision/AI performance (perhaps for specific applications). For now (mid-2026), *BER* is new and still somewhat experimental.

Camera characterization

To fully characterize camera performance, we recommend the **"Three Pillars of image quality measurement."**

I. **Spatial measurements over the image field,** consisting of sharpness, information capacity and related metrics, and often lateral chromatic aberration, each measured at multiple locations in the image, usually at a single illumination level. Spatial measurements may include optical distortion, uniformity, and texture (which is only an issue for nonuniformly processed, i.e., bilateral-filtered images). A spatial measurement at a single location is *never* sufficient to characterize a camera's performance.

II. **Tonal measurements**, usually made near the center of the image. These require a grayscale test chart — typically a transmissive chart with a large density range. Measurements include tonal response and noise, Signal-to-Noise Ratio (SNR), and Information capacity, C_4 (measured from the newly designed [InfoDR chart]), as a function of illumination. Dynamic Range is derived from these measurements.

III. **Other measurements.** These include color response, (flat field) uniformity, and defect analysis (also from a flat field image), none of which are directly related to information metrics.

Believe it or not, that pretty much covers the key image quality measurements. Each measurement can, of course, be made in a number of different ways, depending on the application, field of view, lighting requirements, and camera resolution. Many test charts enable multiple measurements from a single image.

Spatial measurements over the image field

These involve measuring sharpness, information capacity, and often lateral chromatic aberration, measured over the image field, usually at a single illumination level. This is typically done with a test chart that contains multiple slanted edges, which make efficient use of space for measuring *MTF*. Although Distortion and Uniformity (light falloff) can also be considered spatial measurements, we will limit our discussion of them because they don't directly contribute to information capacity.

We recommend one of the four multi-ROI charts supported by ***Imatest Rescharts***. These charts are summarized by a table in the Imatest Documentation. All can measure sharpness for all edges and Lateral Chromatic Aberration (LCA) in the outer edges. They are designed to fill the frame, individually (SFRplus, eSFR ISO, or Checkerboard) or as a group (SFRreg).

Table of spatial resolution charts

SFRplus — Excellent spatial resolution: many slanted edges), Color and grayscale patches are present, but too small for good noise statistics. Good optical distortion measurements (though less detailed than Checkerboard). Must be framed with white space above and below the top and bottom bars.	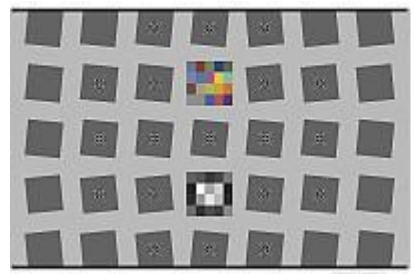
eSFR ISO — An enhanced version of the edge SFR chart illustrated in the ISO 12233 standard [5], with additional slanted edges, wedges, and colors. Good spatial resolution (though not as good as SFRplus or Checkerboard). Very limited optical distortion measurements.	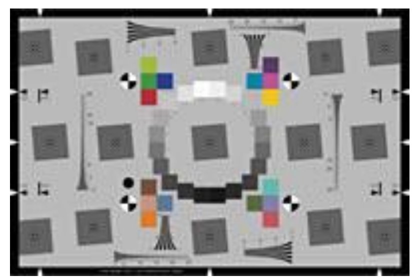
Checkerboard — High spatial resolution. Most detailed optical distortion measurements. No grayscale or colors. Works over a wide range of magnifications, limited only by chart quality and ROI size.	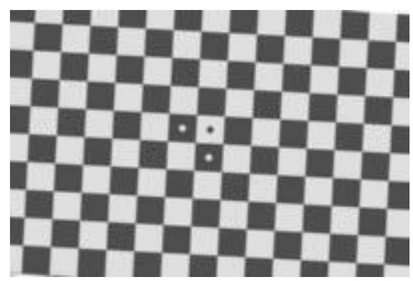
SFRreg — Several charts are needed to cover the field of view. Not recommended for applications where any of the above three charts do the job. Useful for long distances or extreme wide-angle lenses. No distortion, color, or grayscale measurements, though color and/or grayscale charts can be added between the SFRreg charts.	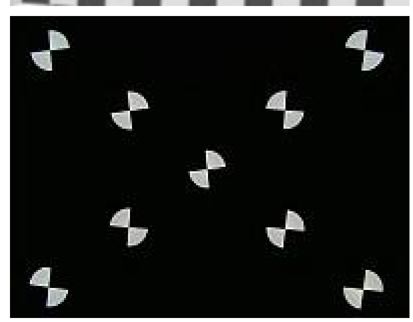

These charts come in a great many sizes and media, allowing you to choose the one that best suits your needs.

The test chart must be large and fine enough to produce reliable sharpness results, i.e., if the chart is magnified (shrunk) to the same size it would occupy on the image sensor, it should be significantly sharper than the lens. Test chart suitability is discussed in two ***Imatest*** web pages, [Test chart suitability for MTF measurements](), and (for charts operating close to their limits), [Compensating camera MTF measurements for chart and sensor MTF]().

Spatial (slanted-edge) results

The results shown below are from [eSFR ISO](), but can be generated by any of the slanted-edge charts listed above.

Edge & MTF plot

The Edge & MTF plot is one of the original plots included in ***Imatest***. The most recent enhancement is the addition of Information capacities (C_4 and C_{max}). We show three plots.

1. From a raw image converted from LibRaw with minimal processing (no sharpening or noise reduction). The complete Rescharts window is shown.
2. Same as 1 with **Plot large ROI…** checked: allows you to examine the ROI closely.
3. From a JPEG image, which has been sharpened and bilateral filtered.

Note that the shapes of the edge and MTF in the raw-converted image (1 & 2) are characteristic of unsharpened images. The peaks in the JPEG results (3) are characteristic of strongly sharpened images. They can be thought of as a "signature" of sharpening.

MTF50 and *MTF50P* are greatly boosted for the JPEG, which is strongly sharpened, but C_4 information capacity is nearly unchanged.

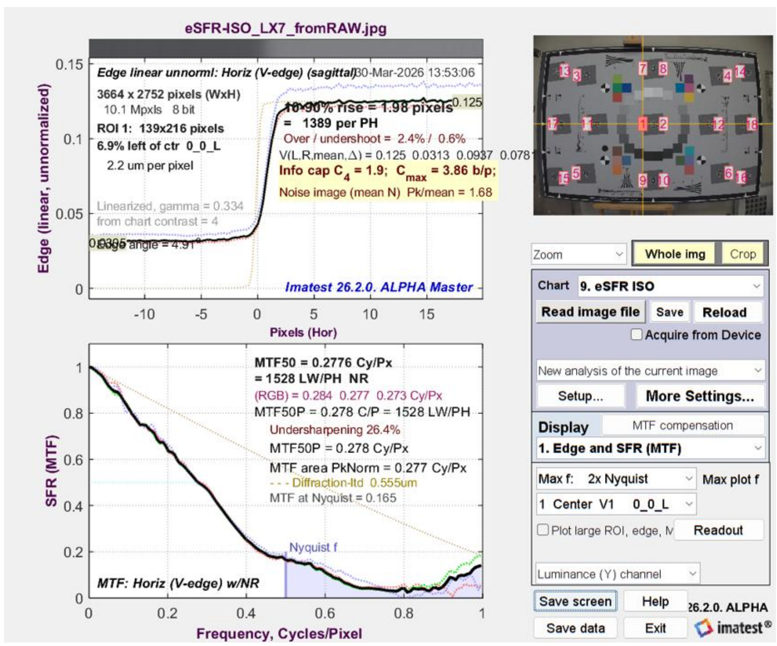

1. Edge & MTF plot from raw-converted image with minimal processing

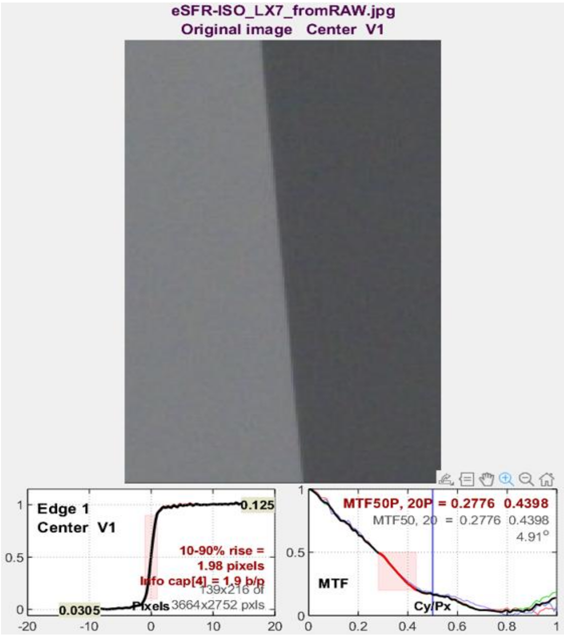

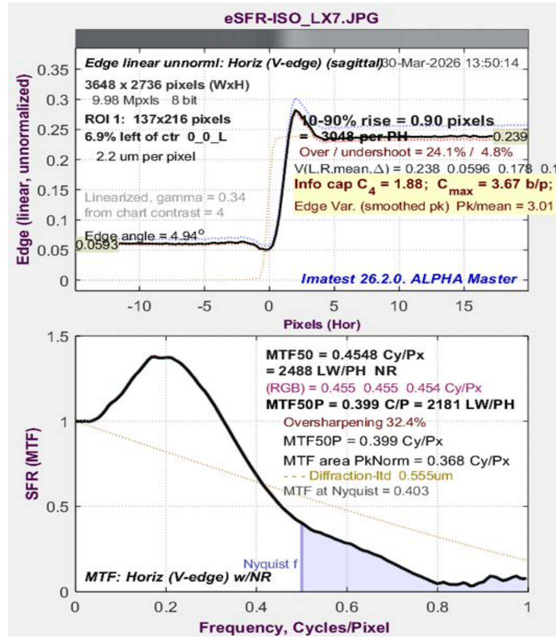

2. Same as 1 with Plot large ROI... Checked	3. Edge & MTF from JPEG

Edge & Spatial noise plot

The spatial noise plots, shown below, are quite new. Nothing like them was available before information capacity was added to ***Imatest***. The plots are for (1) a raw-converted (unsharpened) image and (2) a JPEG (bilateral filtered and strongly sharpened) image.

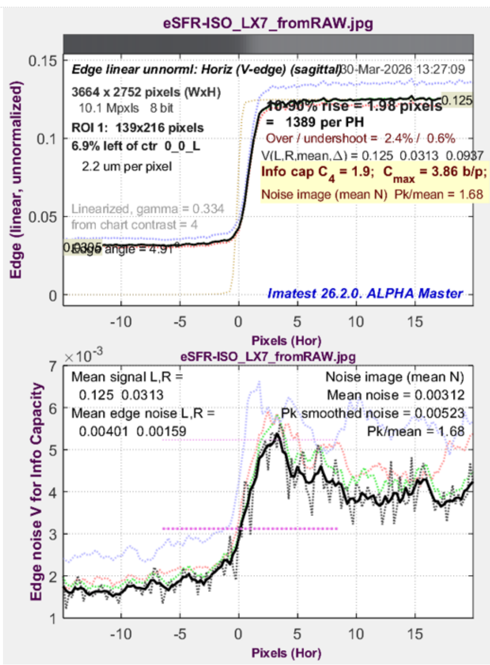

1. Edge & Spatial noise plot from raw-converted image with minimal processing

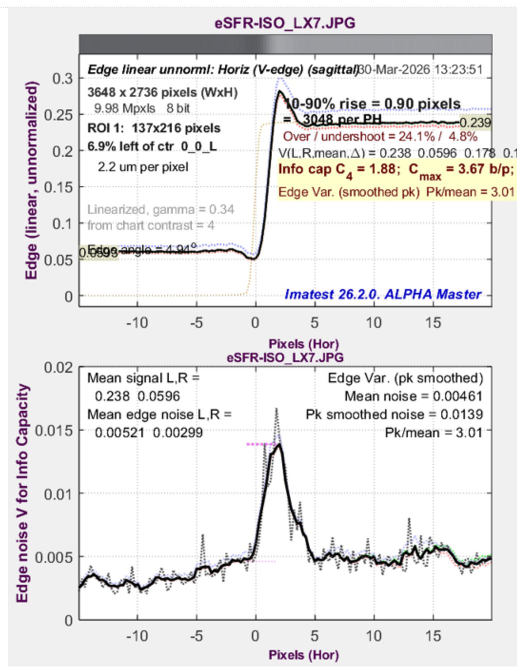

2. Edge & Spatial noise from JPEG

The large spatial noise peak for the JPEG (lower plot, bottom-right) is striking and important. As we've indicated in earlier documents introducing the information metrics, it indicates the presence of bilateral filtering [6], which can cause texture loss, but it's not 100% reliable. Sometimes peaks appear in uniformly processed images, but they are generally small, like the peak on the left, below. A peak may be absent If the bilateral-filtered image is out of focus or the lens is defective. On one occasion we found an unexpected strong peak in a "raw" image from a premium camera phone that didn't seem to affect the performance.

If a strong peak indicating bilateral filtering is present, it may be worthwhile to analyze the camera with a [Dead Leaves (Spilled Coins)]() or [Log Frequency-Contrast]() chart to determine the amount of the texture loss.

3D plot (summary results plotted over the image surface)

3D plots, which include weighted summaries of results, are available for a large number of summary metrics (*MTF50*, *MTF50P*, *MTF Area*, Information capacities C_4 and C_{max}, etc.) We show results for $MTF50P$ — the spatial frequency where *MTF* falls to half its peak value — which is relatively insensitive to exposure. *MTF50P* is identical to *MTF50* and very close to *MTF area (Peak normalized)* for images that are unsharpened or have low sharpening (no response peaks).

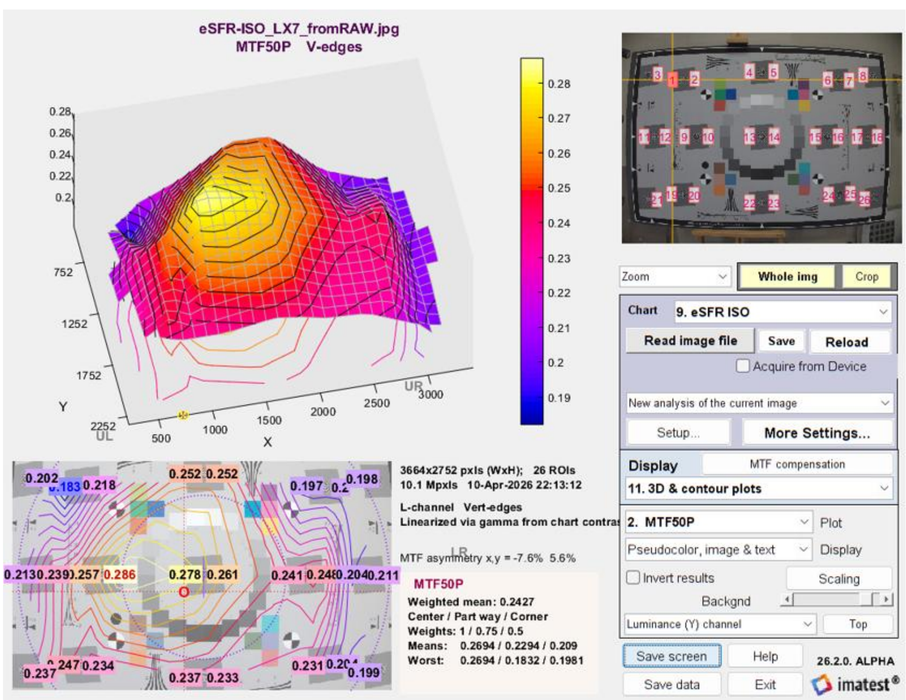

3D plot of MtF50P over the image surface.

C_4 information capacity closely tracks *MTF50P*, but it is highly sensitive to exposure, hence it is useful primarily for its *relative* values, unless exposure is carefully controlled. If you plan to use C_4, we recommend aiming for an average edge reflectance (the mean of the light and dark portions of the linearized ROI) of 0.18, which is close to "neutral gray." Measurements of C_4 as a function of exposure are described in the section on the [InfoDR]() chart below.

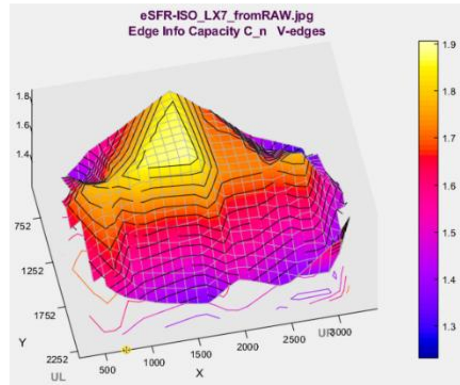

3D plot of C_4 information capacity

Additional spatial resolution charts

Several additional charts, shown below, are available in ***Imatest*** for *SFR* (*MTF*) measurements. They are all less efficient in their use of space than the slanted edge, hence less appropriate for mapping sharpness over the image field, but they have several specialized uses, such as measuring texture.

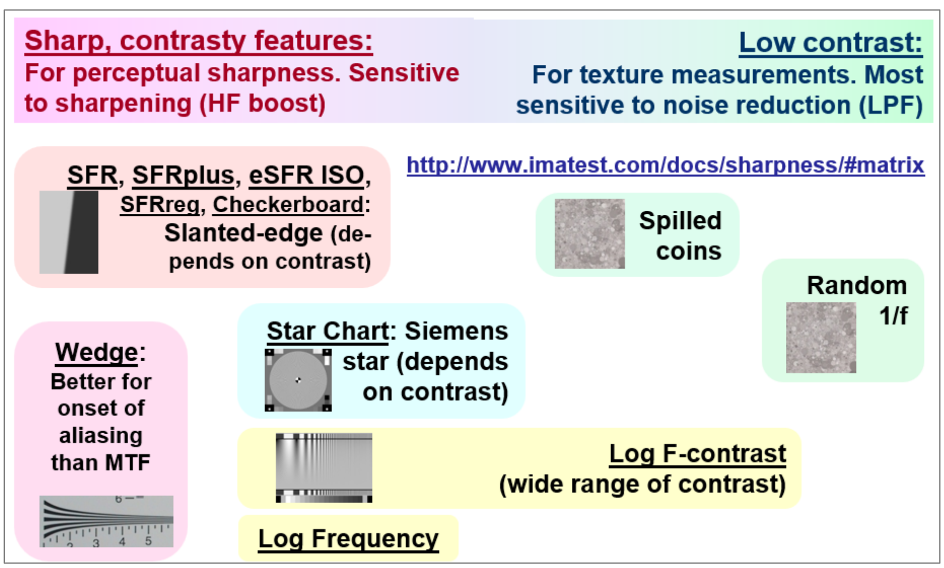

MTF test patterns, ordered by contrast ranges: high on left; low on right

Several additional charts, shown above, are available in ***Imatest*** for *SFR* (*MTF*) measurements. They are less efficient in their use of space than the slanted edge, hence less appropriate for mapping sharpness over the image field, but they have their uses, such as measuring texture. They are included here primarily for reference.

Wedge — ***Imatest's*** logarithmic wedges have a much better frequency distribution than traditional hyperbolic wedges, resembling Bode or frequency response plots. Primarily designed for visual analysis, but ***Imatest*** can analyze them for the onset of aliasing (closely related to vanishing resolution). There is significant demand from industry, but we don't consider wedge measurements to be vital for assessing camera performance.

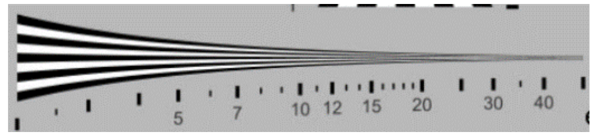

Siemens star — widely-used, but less space-efficient than the slanted edge. Slightly less sensitive to bilateral (nonuniform) filtering. Low-contrast versions may be useful for measuring texture. It can be used to measure information capacity, where its continuous-tone design makes it valuable for measuring the effects of data compression.

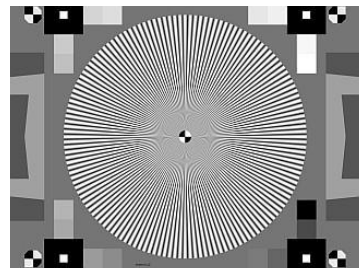

Dead Leaves or Spilled coins (***Imatest's*** implementation of Dead Leaves) — Widely used for measuring texture, but results can be confusing because the threshold for sharpening/noise reduction can be lower than the maximum contrast (3:1) of the chart. This causes some edges to be sharpened and others to be blurred, muddling the results. Its random design makes it relatively immune to AI "cheating" (recognizing the chart design, then creating a perfect replica to pass tests flawlessly).

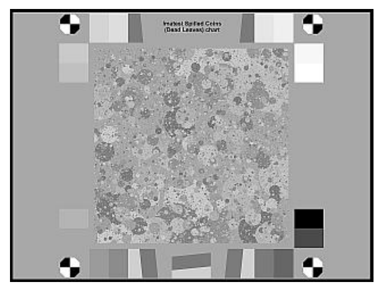

Log-F Contrast — a sinusoidal chart whose spatial frequency increases along the x-axis and contrast^2 (i.e., modulation^2) increases along the y-axis. Useful for displaying the effects of bilateral filtering [6] on texture when a noise peak in the Edge & Spatial noise plot indicates its presence.

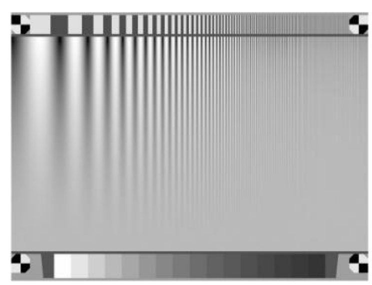

Because Log-F Contrast is particularly good for illustrating the texture loss from bilateral filtering, we illustrate a key result. The high contrast sinusoids near the top of the image (modulation > 0.2) have a significant sharpening boost that disappears at low modulation (< 0.1), where lowpass filtering (noise reduction), which blurs out fine texture, is increased. The contour boundaries are mostly vertical when uniform processing is applied.

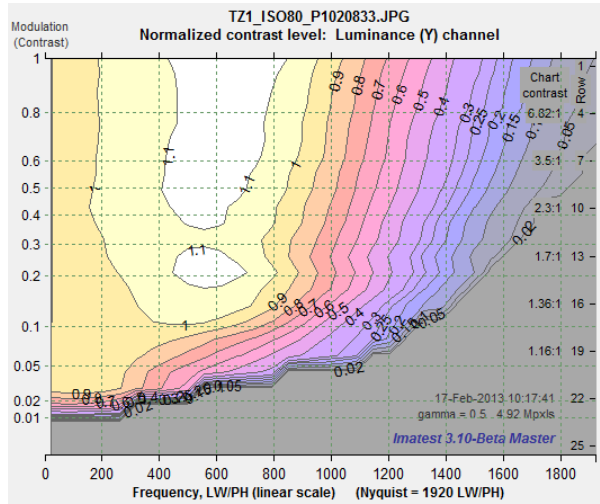

Log F-Contrast results

Tonal measurements

Tonal measurements are complimentary to the spatial measurements described in the previous section. They are made from transmissive grayscale charts, which can have a larger tonal range (maximum density, D_{max}) than reflective charts. For medium to high resolution cameras (> 2 megapixels) the test chart images should occupy the central portion of the image, i.e., the chart image is not designed to fill the frame.

Although we describe three charts for tonal measurements, we will emphasize the InfoDR (Information-based Dynamic Range) chart, which conveniently measures C_4 information capacity over a wide range of illumination from a single exposure. *Strongly recommended* for new work.

Photographing and analyzing tonal charts — Charts should be backlit with a lightbox or light panel and photographed in a dark environment, with care taken to minimize reflections back to the chart. They normally occupy the central portion of the image. If you can control the exposure,

- For standard Dynamic Range (DR) measurements from flat patches (DR36 charts, etc.), expose so a maximum of one patch is saturated.
- For **InfoDR** DR measurements, which are derived from slanted edges, i.e., pairs of patches, expose so the maximum mean density is a little *under* saturation. This reduces a calculation inconsistency when the lighter patch nears saturation.

The use of transmissive charts for measuring dynamic range is described in Dynamic Range. There are several flavors. All except DR_{sensor} are camera dynamic range.

Name	Chart	Description
$DR_{SNR\text{-}36}$	DR36	Widely-used traditional *DR* measurements, from the scene-referenced SNR of flat patches on two different test charts. A minimum slope is required. Quality levels from "high" (20 dB) to "low" (0 dB).
$DR_{SNR\text{-}info}$	InfoDR	
DR_{CR}	Contrast Resolution	Based on SNR_{CR} = (light signal-dark signal)/(mid noise), where "mid" refers to the large surrounding patch (middle density).
DR_{C4}	InfoDR	Measured from C_4 information capacity. Quality levels from "Excellent" (C_4 = 2) to "Bad" (C_4 = 0.1) are not exactly the same as the SNR-based measurements. Best performance predictor.
DR_{sensor}	Flat field image	Image sensor *DR*, from a succession of differently-exposed flat images. From EMVA 1288. Generally larger than camera *DR* because there is no stray light.

36-patch Dynamic Range (DR36) charts

36-patch Dynamic Range charts are widely used for measuring tonal response and dynamic range, $DR_{SNR\text{-}36}$. They come in one-, two-, and three-layer versions with maximum (patch – base) densities, ($D_{max} - D_{base}$), of 2.5, 5, or 7.5, equivalent to 50, 100, and 150 dB. The added layers cover the bottom three rows of the active chart pattern.

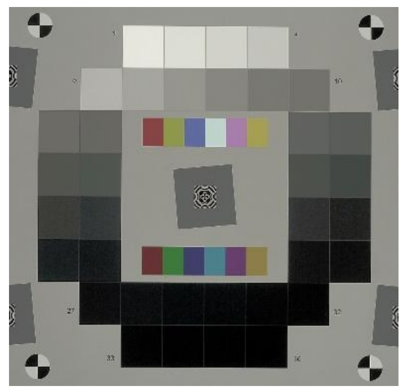

DR36 chart

Dynamic range $DR_{SNR\text{-}36}$ is measured from scene-referenced SNR, which can be simplistically defined as the noise in each patch divided by the derivative of the patch's digital number with respect to illumination.

The figure below shows the primary output of a DR36 run for a 10-megapixel compact camera: the Noise analysis, displaying scene-referenced SNR and Dynamic Range for several quality levels.

The upper plot is the tonal response (also called OECF). Measured gamma = 0.440 is almost exactly the ideal value of 0.454 = 1/2.2 for standard color spaces. The "plateaus" in the dark region (Exposure < −90 dB) are caused by stray light ("ghost images") reflected from the lighter upper rows of patches to the darker lower rows. This has been slightly improved in the new V2 version, where the lighter patches are in the central rows. The dynamic range is shown in the colored horizontal bars in the lower plot: 31.7 dB for "High" (*SNR* = 20 dB = 10); 59.8 dB for "Low" (*SNR* = 0dB = 1), where the quality is so poor that it would be difficult to distinguish any detail in moderate-contrast objects. There is no simple way to relate SNR-based Dynamic Range, $DR_{SNR\text{-}36}$, to C_4 information capacity.

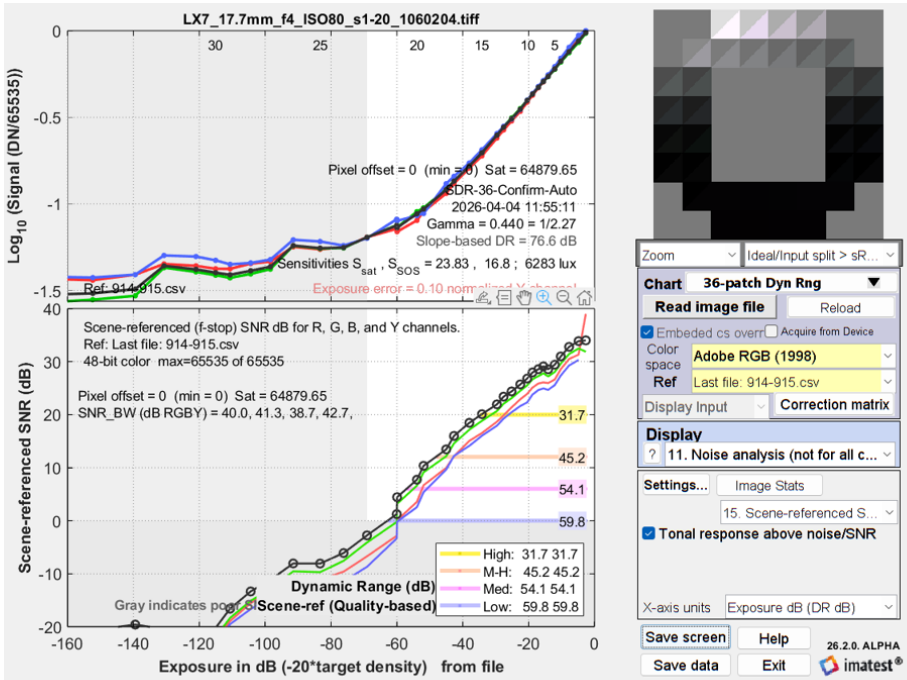

DR36 results: Noise analysis showing Scene-referenced SNR & Dynamic Range, *DR*$_{SNR\text{-}36}$

Contrast Resolution chart

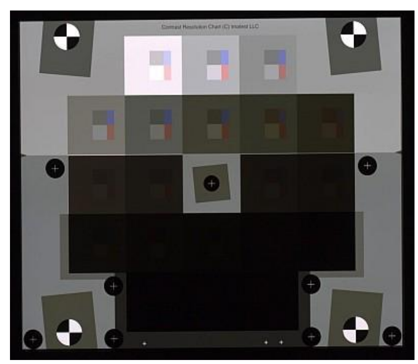

Contrast Resolution chart

The [Contrast Resolution chart](#) was designed to measure the visibility of low contrast objects over a wide range of tones.

It consists of twenty large patches that cover a 95 dB tonal range, each of which contains four smaller patches. The small light and dark gray patches have a 2:1 (6 dB) contrast ratio (Michelson contrast = $(\tau_n - \tau_{n-1})/(\tau_n + \tau_{n-1}) = 1/3$) with the same mean density as the surrounding large patch.

The difference between them defines the signal for the [Contrast Resolution Signal-to-Noise Ratio (SNR_{CR})](#) measurement, where noise is measured in the larger gray patch, which has better noise statistics. The difference signal responds correctly to flare light or uncorrected black level off-set, resulting in good measurements in the

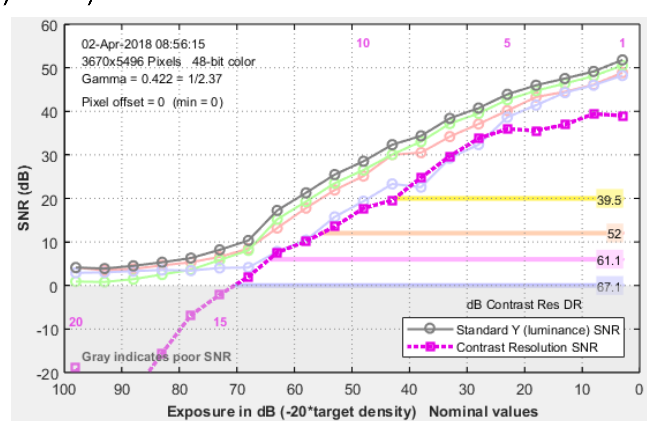

Contrast Resolution Dynamic Range, DR$_{CR}$

presence of [tone mapping](). SNR_{CR} is a good indicator of the visibility of small objects over a wide exposure range. The red and blue patches are for visual analysis-only.

Information-based Dynamic Range (InfoDR) chart

The Information-based (InfoDR) Dynamic Range chart, introduced in early 2026, is a major addition to ***Imatest's*** menagerie of measurements. Because its approach to measuring C_4 information capacity over a wide range of illumination is a major improvement over previous Dynamic Range and low-light measurements, we describe it in detail.

The transmissive InfoDR chart consists of 6 groups of patches: 3 lighter on top and 3 darker on the bottom. Within each group, all boundaries between all adjacent patches consist of slanted edges with a 4:1 contrast ratio (Δ*D* = 0.6). The groups are offset from their neighbors by Δ*D* = 0.2.

The overall chart has 27 distinct values of *D* with Δ*D* = 0.2, for a patch tonal range of 104 dB or an edge tonal range of 92 dB

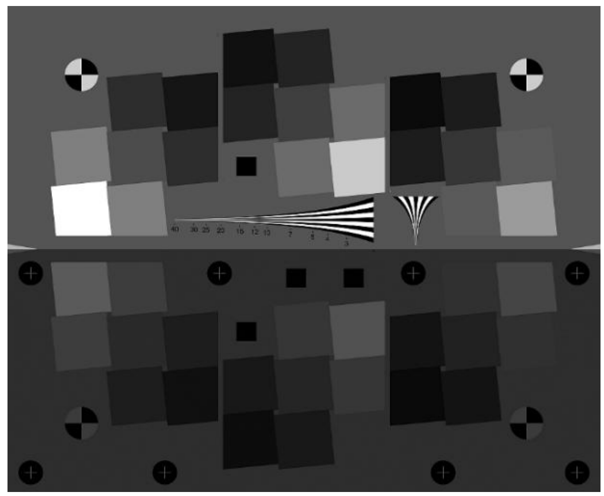

2-layer transmissive InfoDR chart

InfoDR chart design

The InfoDR chart is designed to

- measure C_4 from 4:1 contrast edges (Δ*D* = 0.6)
- over a wide range of illumination in a compact area where sharpness (*SFR*) is reasonably consistent,
- with smaller density steps (Δ*D* = 0.2) for the overall chart to achieve good tonal resolution.

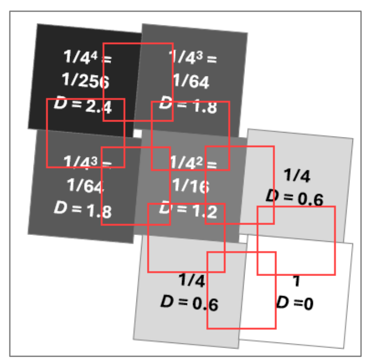

Building blocks of the InfoDR test chart

Because 4:1 (Δ*D* = 0.602) edge density steps are coarser than desired, we designed a two-layer transmissive chart, with 3 groups of 7 patches on the top and 3 on the bottom that mirror the top, where each group is offset by Δ*D* =0.2 from its neighbors. The bottom half of the chart is covered with a sheet with *D* = 2.4.

The design was accomplished with the building blocks shown above. The lightest patch (transmittance $\tau = 1$; density $D = 0$ relative to the base density) is shown on the lower right. The two neighboring patches (to the left and above) have $\tau = 1/4$ ($D = 0.602$). The next single patch, adjacent to the previous two, has $\tau = 1/16 = 1/4^2$ ($D = 1.204$). This progression {1, 2, 1, 2... patches} continues as needed. All of the boundaries between adjacent patches are slanted edges with a 4:1 contrast ratio, equivalent to density step = $\Delta D = 0.602$ or Michelson contrast = $(\tau_n - \tau_{n-1})/(\tau_n + \tau_{n-1}) = 0.6$.

Property	Patches	Edges
Total number	42	48 total: 24 nr-H, 25 nr-V
Distinct density (D) values	27	
Total density range OD (dB)	5.2 (104 dB) in 26 steps	4.6 (92 dB) in 23 steps
Density step D_{step} (total chart)	0.2 (4 dB)	

The **InfoDR** chart's edge density range of 92 dB is sufficient for the great majority of cameras, including High Dynamic Range (HDR) cameras, which can have up to 150 dB *sensor* dynamic range. But practical *camera* dynamic range is limited to around 100 dB by stray light from lens surfaces and interior reflections.

Working with the InfoDR chart

The most important considerations when photographing the InfoDR chart are

- The active area of the chart should fill only the central portion of the image, about 600 to 1000 pixels (vertically). Fewer pixels may reduce measurement consistency; more is unnecessary and increases the likelihood that the outer edges may be far enough from the image center to have reduced *SFR*.
- The chart should be back-Illuminated with a lightbox or light panel in a dark environment, taking care to minimize the light reflect back to the chart.
- For best accuracy, exposure should be set so the brightest patch is just below saturation.
- **The chart must be accurately focused to obtain correct values of C_4.** This is not required for traditional SNR-based Dynamic Range measurements, where some misfocus can be tolerated.
- To display C_4 results as a function of *absolute* illumination, lightbox luminance must be measured. Otherwise, the x-axis will have *relative* units — Log_{10} exposure (−Density), Exposure (dB), or F-stops (EV) (all based on chart density).
- As of April 2026, the InfoDR chart is only available in 7.75x9.25-inch (197x235 mm) LVT color film. It should be easy to make smaller (but not microscopic) versions. We are working on a larger VisNIR photomask version.

Measuring luminance — Because all illumination comes from behind the chart, a luminance (reflected light) meter is required for absolute light measurements. Luminance meters have limited fields of view for measuring the source (lightbox) luminance, L_{source}, or patch luminance, L_{patch}. The relationship between the two is $L_{source} = L_{patch}\ 10^{D_{patch}}$, where D_{patch} is the patch density obtained from the density reference file. The meter can be held very close to the chart (even in contact) because shading isn't an issue.

Edge contrast adjustment

The nominal edge contrast for measuring C_4 is 4:1 (ΔD = 0.602). However, actual chart densities, and hence edge contrasts, vary because transmissive charts cannot be manufactured with perfect consistency, which is why they are supplied with individually-measured density files.

The actual density increment of edge i is the difference between the adjacent measured patch densities, i.e., $\Delta D_i = D_j - D_k$. for adjacent patches j and k. To correctly represent the edge signal for the C_4 calculation, replace nominal $\Delta V_i = (V_{light} - V_{dark})$ with

$$\Delta V_{i-corrected} = \Delta V_i\ 10^{(mean(\Delta D) - \Delta D_i)}$$

Where $\mathrm{mean}(\Delta D) \cong 0.602$ is the *mean* measured patch density increment.

InfoDR Results

Tests were performed on several consumer cameras that had raw output, which could be challenging to obtain with some development systems.

The figures on this and the next page contain the most significant results from the InfoDR chart. The upper plot of each figure contains C_4 as a function of exposure and the Information-based Dynamic Range, DR_{C4}. The lower plot displays the logarithm of the normalized digital number, $\log_{10}(DN/DN_{max})$ for each patch, i.e., the density response or OECF (Opto-Electric Conversion Function), which is equivalent to a classic film characteristic curve.

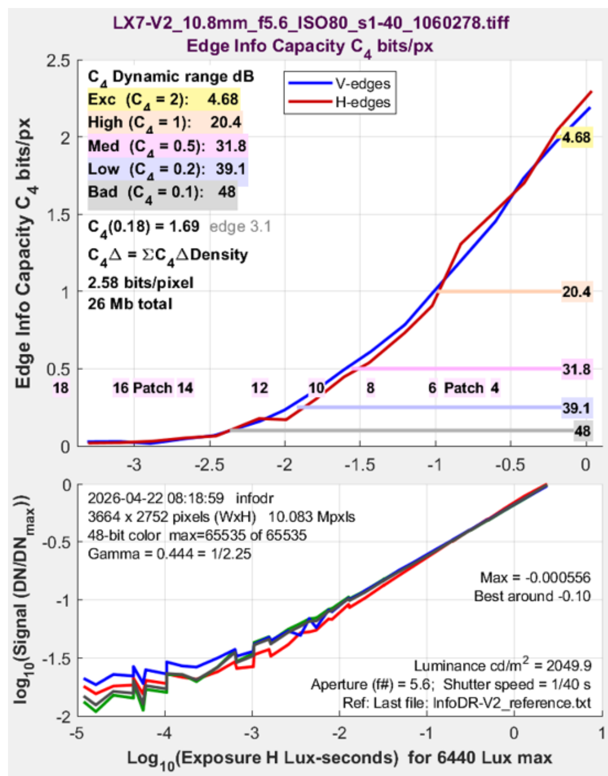

Edge information capacity C_4 and tonal response for a 10-megapixel compact consumer camera with 2.14 µm pixel size and a premium zoom lens.

The bumps in the lower tonal response plots are caused by reflections from the light to dark patches. They are more regular but clearly worse for the DR36 chart, shown above, where the minimum value of $\log_{10}(DN/DN_{max})$ only reaches −1.5.

The x-axis, Log_{10}(Exposure H in Lux-seconds) of the figures, is the approximate exposure at the focal plane for each patch, derived from ISO standard 12232:2019, Annex B [12].

$$H \cong \frac{0.65\,L\,t}{A^2}$$

where A is the aperture (the lens f-number), t is exposure time in seconds, and L is the patch luminance in candelas per meter^2 (cd/m^2).

A and t are often available from EXIF metadata. The equation involves several approximations, most notably, lens transmission factor T = 0.9. T is easy to find for cinema lenses, but is rarely available and difficult to measure for still camera lenses. Since it can vary from about 0.85 to 0.95, depending on the number of lens surfaces and the quality of the coatings, the 0.9 approximation should be adequate for most applications. The ISO 12232 standard has a more precise equation for close distances (image distance < 10∗lens focal length) or when T is known.

$C_4(0.18)$ is C_4 measured at the mean linearized & normalized edge Digital Number (DN) of 0.18.

$C_4\Delta$, shown on the left of the upper plots, is a preliminary heuristic figure of merit that combines C_4 and dynamic range. Its value is close to the maximum information capacity, C_{max}, and it correlates well with the perceptual quality of cameras we've tested, which range from the compact camera above to the high-end heavyweight below.

$C_4\Delta$ is calculated from a simple summation,

$$C_4\Delta = \sum C_4(x)\Delta x$$

where $x = \log_{10}$(Exposure H in Lux-seconds) and Δx is the x-axis increment = 0.2 OD (Optical Density units). $C_4\Delta$ is the area under the C_4 curve in the upper plot.

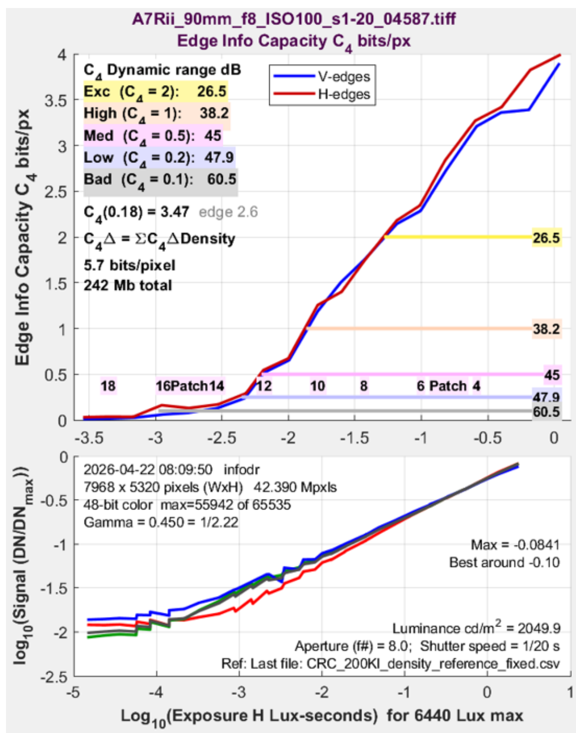

Edge information capacity C_4 and tonal response for a 42-megapixel professional-grade camera with 4.51 μm pixel size BSI sensor and an excellent 90mm macro lens.

[Note that C_{max}, which was described in earlier papers on information metrics, has been de-emphasized because there are several pitfalls in calculating it: it can be inaccurate for HDR image sensors, where good noise models would be complex (and different for each sensor) if they were available, or if the maximum Digital Number is less than $2^{(bit\ depth)}-1$.]

Note that even though the two figures look similar, the numeric results and x and y-axes of the C_4 plots are *very* different. $C_4\Delta$ for the 42-megapixel full-frame camera on the right is 9x higher than for the compact 10-megapixel camera above: not unexpected given the difference in the price and weight of the cameras. (Sometimes, you get what you pay for.)

Although $C_4\Delta$ is of interest for consumers and engineers tasked with selecting cameras, the added detail in the plots of C_4 as a function of Exposure H should be especially useful for camera designers concerned with low light performance.

If $C_4\Delta$ turns out to be a durable metric for predicting the performance of a large variety of cameras, it will need a better name.

SNRi

InfoDR can calculate the object detectability metric, ***SNRi***, for multiple object sizes (values of *w* for *w* x 4*w* pixel rectangles) at multiple illumination levels. This plot can be produced after applying Image Signal Processing (ISP) with the **Simatest** simulator. It is a particularly good indicator of low light performance. You can select a Michelson contrast for this display; You are not restricted to the chart's contrast (0.6 for a 4:1 contrast ratio).

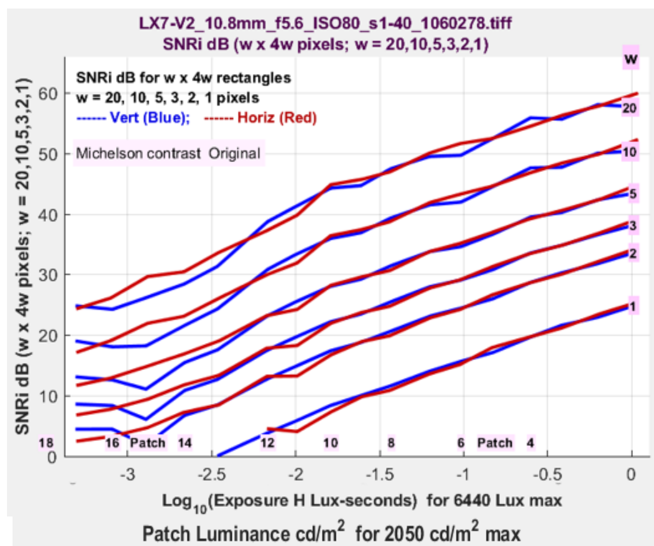

SNRi for multiple rectangle sizes over a range of illumination

Deeper exploration

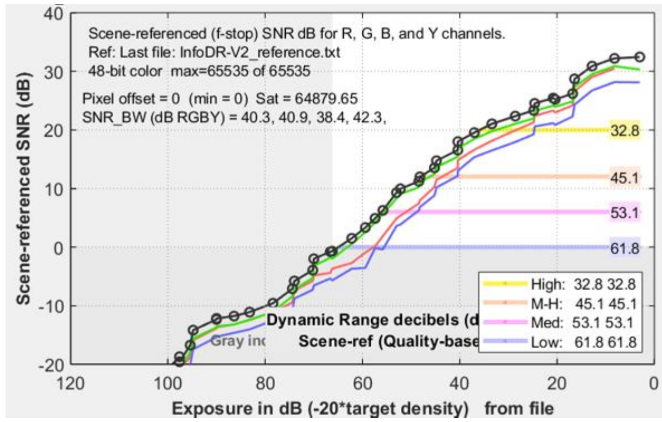

SNR-based Dynamic Range $DR_{SNR\text{-}info}$ for the same 10 megapixel image capture used for DR_{C4}. $DR_{SNR\text{-}info}$ is nearly identical to the $DR_{SNR\text{-}36}$ measured on the 36-patch DR36 chart, above.

In addition to information-based Dynamic Range (*DR_{C4}*), the new InfoDR chart can also be used for all standard slanted-edge measurements as well as traditional *SNR*-based *$DR_{SNR\text{-}info}$*, which doesn't fully characterize performance.

For the 10-megapixel camera, *DR_{C4}* is slightly lower than *$DR_{SNR\text{-}info}$*, in part because the *DR_{C4}* quality levels are new, somewhat arbitrary, and don't correspond exactly to the classic *$DR_{SNR\text{-}info}$* quality levels. But they better represent camera performance.

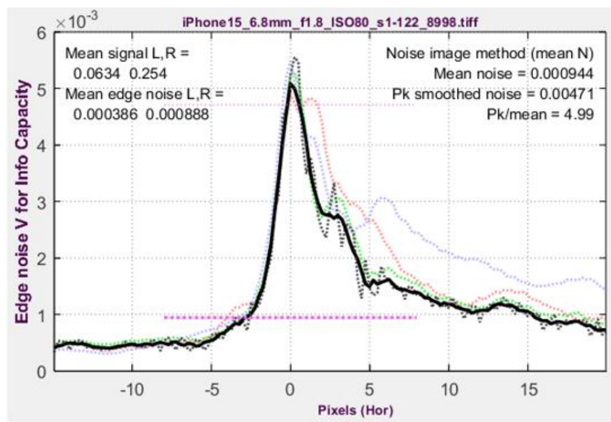

Spatial noise, showing unexpected peak

Results can sometimes be surprising. For example, a peak in the spatial noise can indicate the presence of bilateral filtering [6], which is a form of edge-preserving noise reduction that degrades texture response. Its presence indicates that it might be worthwhile to measure Log F-Contrast or Spilled Coins to better understand the image processing.

Although it is almost universal in JPEG files from cameras, we've only seen it once in a raw image — for a premium camera phone, that appeared to have noise reduction but no sharpening. When it is present, information capacity is estimated using the amplitude of the smoothed peak noise instead of the mean noise. Even with noise reduction in the "raw" image, it had impressive performance.

Photon Transfer Curve and Simatest

A particularly useful result from any of the three transmissive tonal charts is the **Photon Transfer Curve (PTC)** [14], which can be calculated from pure raw (undemosaiced) test chart images, as a result of a special property: the noise in each patch is a function of the mean Digital Number (*DN*) of the patch, *independent* of color.

The PTC, which is a plot of noise as a function of exposure, characterizes an image sensor's noise behavior. It is the heart of the ***image sensor noise model***, which consists of dark noise, photon shot noise, and PRNU (Photo Response Nonuniformity) noise, each of which responds differently to light. For linear sensors, noise can be characterized by just three coefficients: $\sigma_N = \sqrt{k_{Ndark}^2 + k_{Nshot}^2 V + k_{PRNU}^2 V^2}$ for normalized Digital Number, V.

Imatest calculates the noise coefficients using nonlinear optimization. [As of April 2026, **Imatest** does not yet support High Dynamic Range (HDR) sensors, which have multiple operating regions.]

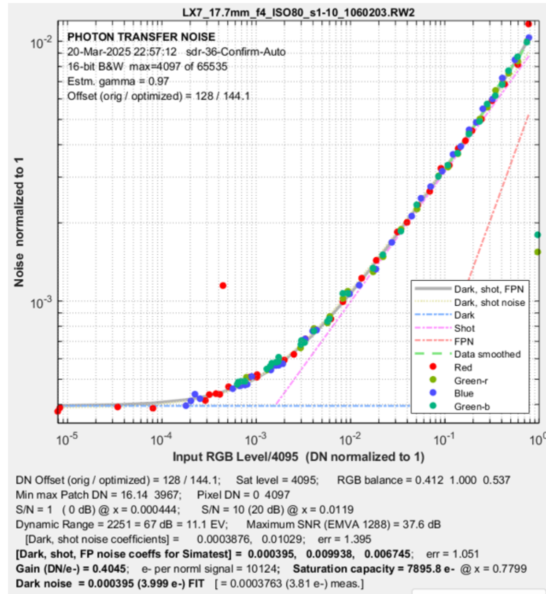

Photon Transfer Curve (PTC)

The image sensor noise parameters can be entered into the **Simatest camera simulator**, along with simulated (ideal) images of any test chart described in this document, to predict the performance of prototype cameras under a wide variety of lighting conditions and Image Signal Processing (ISP).

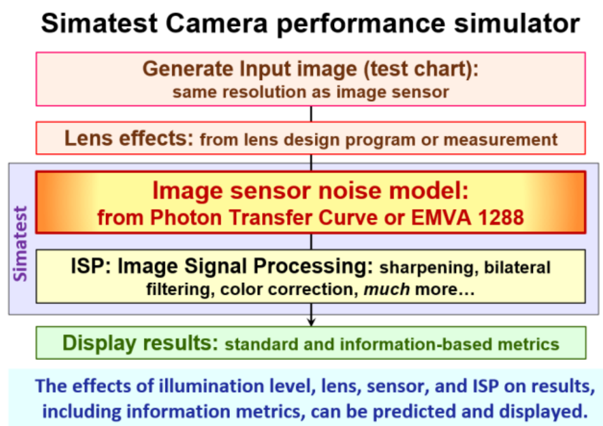

Simatest: Simplified block diagram.

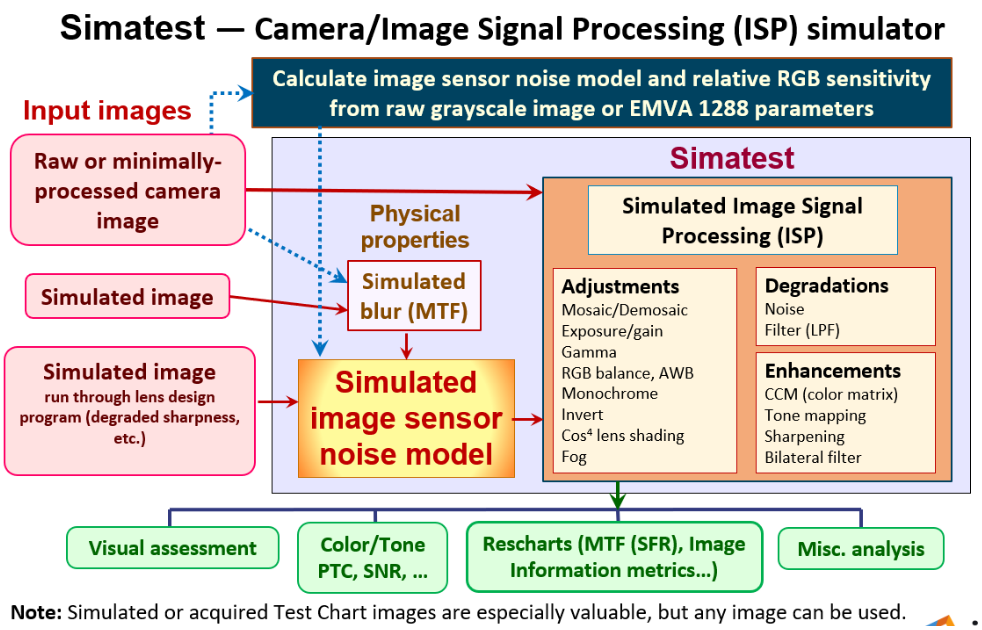

Simatest: Detailed block diagram showing input, ISP, and output options.

Any of the results shown in this document can be calculated from Simatest output. Details are on **Simatest: Overview**, **Simatest: instructions and reference**, and related pages.

Other measurements

The spatial and tonal measurements presented so far — the first two pillars of image quality measurement — are the key measurements for characterizing camera performance, but more may be needed to complete the picture. Since none of these are directly related to information capacity, we'll keep the descriptions brief.

Color

Color can be measured with two of the spatial charts: eSFR ISO or SFRplus, in Rescharts (interactive) mode or from their fixed, batch-capable modules. The best results are obtained for eSFR ISO when color patches have been measured and a reference file is available. Split colors (reference/input) and a*b* color error (on the right) can be displayed.

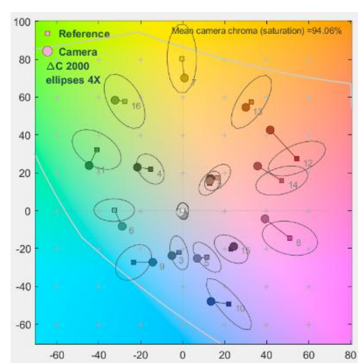

a*b* color error

Most color charts are designed be analyzed with Color/Tone, which can provide more detail than the spatial chart modules (above) and can be run interactively or as a batch-capable fixed module. The most popular color charts are

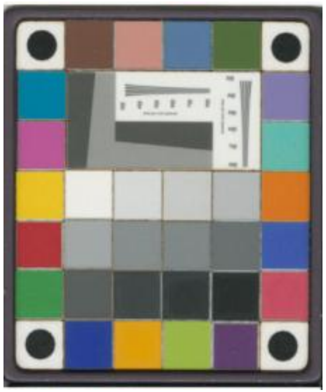

Rezchecker

- the familiar 24-patch **[Colorchecker Classic](#)**,
- **[Colorchecker SG](#)**,
- **[Colorgauge](#)** and **[Rezchecker](#)**: both tiny charts, well-suited for endoscopes, available in several sizes, whose color patches are made from the same pigment-based material as the Colorcheckers.

Colorchecker Classic

Uniformity (Light Falloff) and Defects (Blemishes)

Uniformity and defects (blemishes) can both be measured from the [Flatfield module](#) [(Using Flatfield, Part 1](#), [Using Flatfield, Part 2](#), [Using Flatfield-Interactive](#), [Using Flatfield Blemish Detect)](#), which runs in both fixed and interactive modes, by directly photographing a light panel or lightbox, which doesn't have to be in perfect focus. No chart is needed.

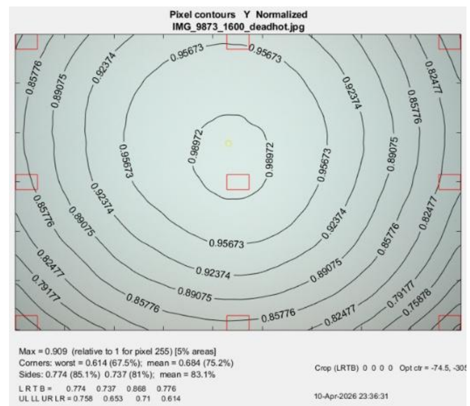

Flatfield contours

Limited uniformity results (but not blemishes) are available from SFRplus (from the light regions between the slanted squares) and Checkerboard (from the light squares).

Problematic images and misleading results

Monkey wrenches, spanners, and wooden shoes

imply situations that can ruin your work. In the US we say, "throw a monkey wrench into the works." In the UK they say, "throw a spanner." In the early nineteenth century, disgruntled French workers threw wooden shoes, called "*[sabots](#)*" into industrial machinery.
The items in this section are metaphorical monkey wrenches, spanners, or sabots.

We describe several issues that can adversely affect measurements. They are common enough so that ***everyone who tests cameras should be aware of them***. Some are beneficial for human vision when applied in moderation, but can distort measurements, especially when applied in excess — something we've seen all too often. Some of the issues are described in [Correcting Misleading Image Quality Measurements](#) [15].

Obsolete test charts

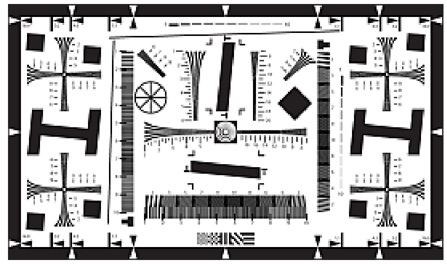

ISO 12233:2000 chart

Several obsolete charts should be avoided because more accurate and convenient choices are available. The most common is the ISO 12233:2000 chart, which has few usable edges, excessive contrast (that can cause problems with saturation), and is unsuitable for automatic region detection. Its drawbacks are described in the eSFR ISO Instructions. Another is the USAF 1951 chart, which was designed for visual analysis and is not supported by Imatest.

Oversharpening

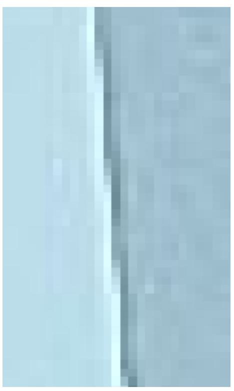

Some software sharpening is almost always beneficial for human vision, but excessive sharpening creates artifacts that can degrade images as well as making summary metrics like MTF50 or MTF50P unrealistically high, meaningless, or even impossible to measure, as shown on the right. It's a cheap way of boosting specifications without adding any information to the image.

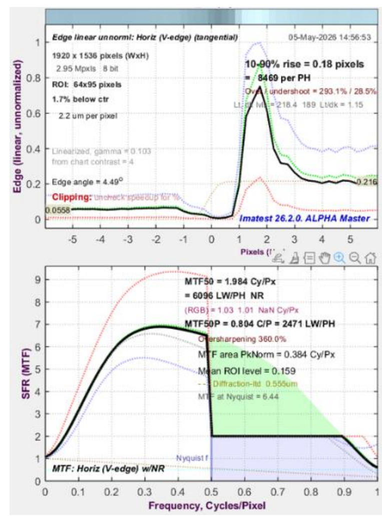

Severe oversharpening

Sharpening can be recognized by "halos" on the edges as well as a peaks in the edge and MTF response. They are acceptable if kept within reason (roughly 30%). Oversharpened images may look crisp on tiny displays like phones, but they should only be sent to the display; they should not be saved.

Bilateral filtering and texture

Bilateral filtering [6], which has been mentioned several times in this paper, is an edge-preserving noise reduction technique that makes pictorial images more pleasing by low-pass filtering relatively smooth regions of images to reduce visible noise while maintaining sharpness (or even sharpening images) near features such as edges, which are responsible for our perception of sharpness. Bilateral filtering

- Is nearly universal in JPEG images from consumer cameras (but is independent of JPEG compression).
- Can make a mess of measurements because it is nonlinear and nonuniform.

- Can be detected by the presence of a strong peak in the spatial noise plot, when information metrics are enabled. When it is present, information metrics such as *SNRi* cannot be reliably calculated.
- Removes fine texture from the image.

When there is concern about image texture (which only happens for bilateral filtering, additional test chart images may be analyzed.

- The best known is the Spilled Coins (Dead Leaves) chart, but results can be misleading. See Dead Leaves measurement issue.
- The Log F-Contrast chart provides a better indication of how bilateral filtering affects texture.

Tone mapping

(Local) tone mapping is a nonlinear, nonuniform ISP operation that lightens dark portions of High Dynamic Range (HDR) images while maintaining local contrast, so that important details in the dark portions are visible on normal displays in normal environments. It strongly impacts tonal response and dynamic range measurements.

It can be recognized by lower than expected values of gamma (where $Digital\ Number\ (DN) \approxeq illumination^{(encoding\ gamma)}$), which is typically around 0.45 to 0.5 for interchangeable images encoded in standard color spaces (sRGB, etc.). Dynamic Range measurements of tone-mapped images made with standard test charts are generally not valid. The Contrast Resolution chart can provide useful dynamic range measurements, focusing on object visibility over a range of tones, for tone-mapped images.

Saturation or clipping

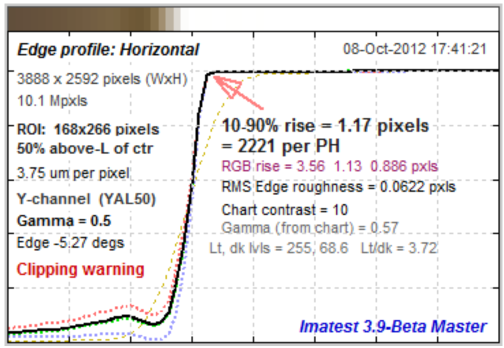

Abrupt clipping

Saturation or clipping takes place when the scene illumination range is too high to be captured by the image sensor, i.e., when the sensor reaches its full well capacity or maximum Digital Number (*DN*). (Clipping can also happen on the dark side.) If the onset of clipping is abrupt (**red arrow →** on the right), the resulting sharp corner contains significant high frequency energy that boosts MTF measurements, making them better than reality. If the onset is gradual and the corner isn't sharp, the effect on MTF will be much lower. The **Clipping warning** is merely an alert that doesn't indicate how severely MTF is affected.

Clipping can be minimized or eliminated with good exposure and with low-contrast targets. Slanted edges with 4:1 contrast ratio are recommended where available. Excessive sharpening makes it worse. Careful exposure is required for Chrome-on-Glass charts, which have a minimum contrast ratio of 10:1.

Dynamic Range inflation

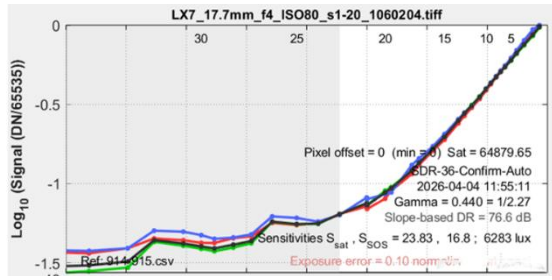

Response of 36-patch Dynamic Range chart, showing "plateaus" that can be misinterpreted. Full details above in [DR36 Charts](#).

Some High Dynamic Range (HDR) sensors have specified dynamic ranges as high as 150 dB, but once they're inside a camera behind a lens, stray light limits the actual dynamic range to around 100 dB. The most common issue with older DR36 test charts is stray light in the darkest regions near the bottom of the chart causing response "plateaus" that can be misinterpreted as part of the signal. Full details are given above in ***DR36 Charts***. Newer versions of the chart with light patches near the center row are better, but not perfect.

For best results — most representative of real camera performance, we recommend the ***Contrast Resolution*** or ***InfoDR (Information-based Dynamic Range)*** charts. But they won't give the 120+ dB measurements coveted by marketing departments.

Summary

We have presented a strategy for measuring camera performance that consists of three pillars.

I. **Spatial measurements:** For a test chart that contains multiple slanted edges distributed over the image field, measure sharpness, information capacity (especially, C_4), and if needed, lateral chromatic aberration, for each edge,

II. **Tonal measurements:** For a test chart that contains grayscale patches with a wide range of densities,
   a. Measure the signal, noise, and Signal-to-Noise Ratio (SNR) from each patch,
   b. If the recommended InfoDR chart is available, measure sharpness and information capacity, C_4, from the 4:1 contrast slanted edges in each patch group to obtain
      1. performance (C_4) as a function of illumination, for a wide range of light levels,
      2. information-based dynamic range, which is a better predictor of camera performance than traditional SNR-based dynamic range.

III. **Other measurements:** As needed — color, uniformity, and defects, to cover omissions from I and II. And it never hurts to ***look at the image***, just in case a rare defect shows up that the measurements didn't catch. (Let us know if this happens.)

The three pillars of image quality measurement

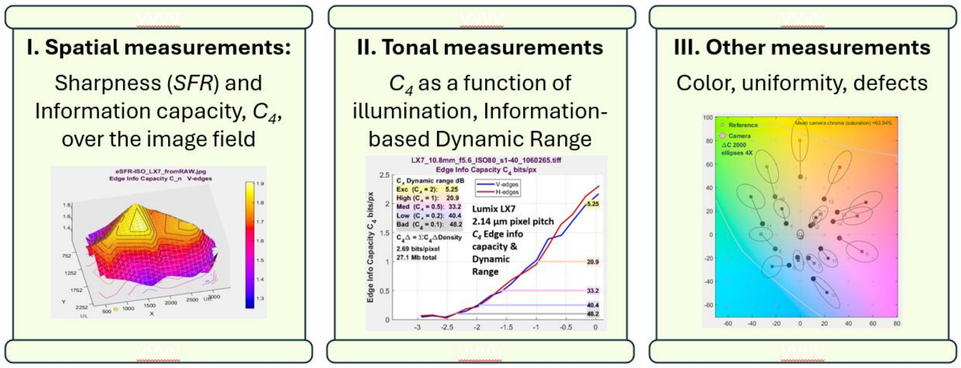

We have described measurements derived from information theory, most importantly the information capacity for 4:1 contrast objects, C_4, which is especially valuable for characterizing camera performance over a wide range of illumination, and also Ideal Observer SNR (*SNRi*) [7-10] and error probability, that quantify how well an object of a given size can be detected.

We emphasized that C_4, which is calculated from three factors — sharpness ($SFR(f)$), noise ($NPS(f)$), and signal amplitude ($V_{light}–V_{dark}$), is a *complete* pixel-level performance metric that can answer the question, "How good is the pixel or camera"

An internet search for Image detection, identification, and recognition turned up lots of clever algorithms, but almost nothing about camera quality. Pixel count was the most common camera specification. The industry hasn't even gotten to *SFR*, much less information content. We have our work cut out.

We encourage the use of information metrics such as information capacity C_4 and *SNRi* (*object detectability*) in addition to (or in place of) sharpness and noise for evaluating camera performance.

In other words, think in terms of *information bits* rather than *cycles*.

References

[1] C. E. Shannon, "A mathematical theory of communication," Bell Syst. Tech. J., vol. 27, pp. 379–423, July 1948; vol. 27, pp. 623–656, Oct. 1948.

[2] C. Shannon, "Communication in the Presence of Noise," Proceedings of the I.R.E., January 1949, pp. 10-21.

[3] N. L. Koren, "Measuring camera information capacity with slanted-edges (Invited)", Electronic Imaging, 2023, pp 454--1 - 454-9, https://doi.org/10.2352/EI.2023.35.8.IQSP-454.

[4] N. L. Koren, "Image Information Metrics From Slanted Edges: A Toolkit of Metrics to Aid Object Recognition, Machine Vision, and Artificial Intelligence Systems", Electronic Imaging, 2024, pp 256-1 - 256-17, https://doi.org/10.2352/EI.2024.36.9.IQSP-256. We recommend the revised version of this paper: https://www.imatest.com/wp-content/uploads/2024/03/Koren_Image_Information_Metrics_paper_final.pdf.

[5] ISO 12233:2024 Digital Cameras — Resolution and spatial frequency responses, https://www.iso.org/stanard/88626.html.

[6] Tomasi, C., and R. Manduchi. "Bilateral Filtering for Gray and Color Images". Proceedings of the 1998 IEEE International Conference on Computer Vision. Bombay, India. Jan 1998, pp. 836–846.

[7] Robin Jenkin, Paul Kane, "Fundamental Imaging System Analysis for Autonomous Vehicles," Proc. IS&T Int'l. Symp. on Electronic Imaging: Autonomous Vehicles and Machines, 2018, pp. 105-1 – 105-10, https://doi.org/10.2352/ISSN.2470-1173.2018.17.AVM-105.

[8] ICRU Report 54, Medical Imaging – The Assessment of Image Quality, Bethesda, MD: International Commission on Radiation Units and Measurements, 1966.

[9] Orit Skorka, Paul J. Kane, "Object Detection Using an Ideal Observer Model", Proc. IS&T Int'l. Symp. on Electronic Imaging: Autonomous Vehicles and Machines, 2020, pp 41-1 - 41-7, https://doi.org/10.2352/ISSN.2470-1173.2020.16.AVM-041.

[10] Paul J. Kane, "Signal detection theory and automotive imaging", Proc. IS&T Int'l. Symp. on Electronic Imaging: Autonomous Vehicles and Machines Conference, 2019, pp 27-1 - 27-8, https://doi.org/10.2352/ISSN.2470-1173.2019.15.AVM-027 .

[11] Norman Koren, Robert C Sumner, Henry Koren, "Measuring MTF with wedges: pitfalls and best practices", *Proc. IS&T Int'l. Symp. on Electronic Imaging: Autonomous Vehicles and Machines*, 2017, pp 6-11, https://doi.org/10.2352/ISSN.2470-1173.2017.19.AVM-451 (Quick summary: Using wedges to measure *MTF* is *rubbish*.)

[12] ISO 12232:2019 Photography — Digital still cameras — Determination of exposure index, ISO speed ratings, standard output sensitivity, and recommended exposure index, https://www.iso.org/standard/73758.html.

[13] Geever, D., Brophy, T., Molloy, D., Ward, E., George, R., Koren, N., Glavin, M., Jones, E., & Deegan, B. (2026). Information Capacity as a Predictor of Perception Performance. *IEEE Open Journal of Vehicular Technology*, *7*, 708-722. https://doi.org/10.1109/OJVT.2026.3655075

[14] J. R. Janesick, Photon Transfer DN → λ, SPIE Digital Library, 2007. https://doi.org/10.1117/3.725073

[15] Norman L. Koren, "Correcting Misleading Image Quality Measurements", *Proc. IS&T Int'l. Symp. on Electronic Imaging: Image Quality and System Performance XVII*, 2020, pp 242-1 – 242-10, https://doi.org/10.2352/ISSN.2470-1173.2020.9.IQSP-242

[16] MIT OpenCourseWare, 6.02 Draft Lecture Notes, 2012, Chapter 9: Noise, https://web.mit.edu/6.02/www/s2012/handouts/9.pdf

Author biography

Norman Koren became interested in photography while growing up near the George Eastman Museum in Rochester, NY. He received his BA in physics from Brown University (1965) and his Masters in physics from Wayne State University (1969), then worked in the computer storage industry simulating digital magnetic recording systems and channels. He founded Imatest LLC in 2003 to develop software, test charts, and lab hardware for measuring the quality of digital imaging systems. Since 2021 he has been obsessed with applying information theory to image quality measurement.

Appendix. *SNRi* and Error Probability

This calculation is included here because it's new, important, and was derived while writing this White Paper.

We use results from detection theory to calculate the total error probability, which is the sum of the probabilities of false positives and false negatives, as illustrated in the Receiver Operating Curve (ROC), below. For the gaussians with means μ_1 and μ_2 and identical standard deviations $\sigma = \sigma_1 = \sigma_2$, the minimum error probability occurs when the decision threshold, $\boldsymbol{x_0 = (\mu_1 + \mu_2)/2}$. Note that the system is symmetrical around $\boldsymbol{x = 0.5}$. More details on detection theory can be found in Wikipedia and in reference papers [7-10].

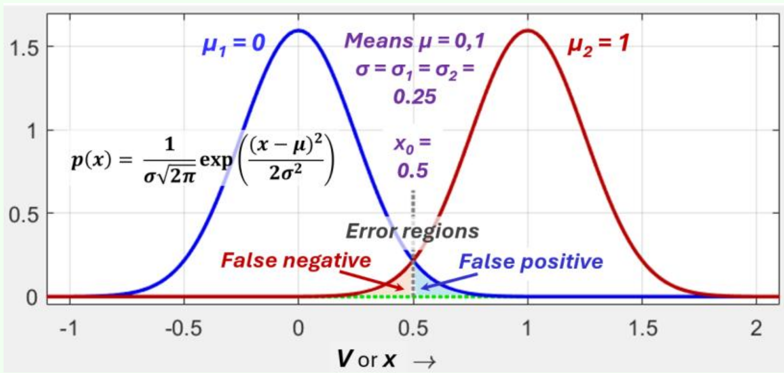

For the gaussian curves,

$$\boldsymbol{p(x) = \frac{1}{\sigma\sqrt{2\pi}}\exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)}$$

The key assumption is that *SNRi* can be treated as a standard Signal-to-Noise Ratio (SNR) with signal = $\Delta\mu = \mu_2 - \mu_1$ and noise = σ, illustrated above.

$$\boldsymbol{SNRi = (\mu_2 - \mu_1)/\sigma = \Delta\mu/\sigma = 1/\sigma}$$

Thanks to the symmetry of the ROC, we can use the left gaussian to calculate the error rate because its mean, $\boldsymbol{\mu_1 = 0}$, makes the calculation simpler. The error probability, which we also call Bit Error Rate, BER, is

$$\boldsymbol{Pr(x > x_0) = BER = Pr(\text{error}) = \int_{x_0}^{\infty} p(x)dx = \int_{x_0}^{\infty} \frac{1}{\sigma\sqrt{2\pi}}\exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)dx}$$

From https://www.gaussianwaves.com/2012/07/q-function-and-error-functions/

For the gaussian curve on the left, the decision threshold is $\boldsymbol{x_0 = (\mu_1 + \mu_2)/2 = 1/2}$.

Let $\boldsymbol{u^2 = (x-\mu)^2/(2\sigma^2)}$. Then, $\boldsymbol{u = (x-\mu)/(\sigma\sqrt{2})}$, $\boldsymbol{du = dx/(\sigma\sqrt{2})}$, $\boldsymbol{dx = \sigma\sqrt{2}\,du}$, and $\boldsymbol{u_0 = 0.5/(\sigma\sqrt{2}) = 1/(2\sqrt{2}\,\sigma) = SNRi/(2\sqrt{2}\,\sigma)}$.

$$\boldsymbol{Pr(x > x_0) = Pr(\text{error}) = \int_{u_0}^{\infty} \frac{\exp(-u^2)}{\sigma\sqrt{2\pi}} \sigma\sqrt{2}\,du = \frac{1}{\sqrt{\pi}} \int_{SNRi/(2\sqrt{2})}^{\infty} \exp(-u^2)\,du}$$

From the MATLAB documentation for erfc, the MIT class notes [16], and other sources,

$$\boldsymbol{\text{erfc}(z) = \frac{2}{\sqrt{\pi}} \int_{z}^{\infty} \exp(-u^2)du}$$

$$\boldsymbol{Pr(x > x_0) = BER = Pr(\text{error}) = \frac{1}{2}\text{erfc}\left(\frac{SNRi}{2\sqrt{2}}\right)}$$

BPSK (Binary Phase Shift Keying) equation (for checking results) — BPSK is symmetrical around $V = 0$ with $\mu = \pm 0.5$, i.e., $V_0 = -\sqrt{E_b} = -1/2$ and $V_1 = \sqrt{E_b} = 1/2$ in the above diagram. The standard equation for BPSK is,

$$\boldsymbol{P_b = P(\text{error}) = \frac{1}{2}\text{erfc}\left(\sqrt{\frac{E_b}{N_0}}\right)}$$

This equation, which is inexplicably missing from the classic papers on SNRi [7-10], can be derived from texts on electronic communications under the topic, "Binary Phase Shift Keying (BPSK)."

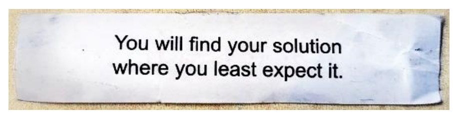

There was a discrepancy when we assumed that $\boldsymbol{N_0 = \sigma^2}$; $\boldsymbol{\sqrt{N_0} = \sigma}$, which is explained in the MIT class notes, https://web.mit.edu/6.02/www/s2012/handouts/9.pdf,

> "We will denote $2\sigma^2$ by N_0. It has already been mentioned that σ^2 is a measure of the expected power in the underlying AWGN process. However, the quantity N_0 is also often referred to as the **noise power**, and we shall use this term for N_0 too."

Based on this statement, it appears that for BPSK, $\boldsymbol{N_0}$, rather than $\boldsymbol{\sigma}$, is what is actually measured, so that $\boldsymbol{N_0 = 2\sigma^2}$, and $\boldsymbol{\sqrt{N_0} = \sigma\sqrt{2}}$. Using $\boldsymbol{SNRi = \Delta\mu/\sigma}$ and assuming $\boldsymbol{\sqrt{E_b} = \Delta\mu/2}$,

$$\boldsymbol{P_b = P(\text{error}) = \frac{1}{2}\text{erfc}\left(\sqrt{\frac{E_b}{N_0}}\right) = \frac{1}{2}\text{erfc}\left(\frac{\Delta\mu}{2\sqrt{2}\,\sigma}\right) = \frac{1}{2}\text{erfc}\left(\frac{SNRi}{2\sqrt{2}}\right)}$$

This resolved the discrepancy we struggled with when we assumed $N_0 = \sigma^2$.