

Calculating & Comparing Image Quality

The quality of a received Slow Scan TV (SSTV) picture may be determined by comparing, pixel-by-pixel, for identical matches in RGB color space for the customary image displayed on an SSTV cam in JPG format.

If a ham received an image perfectly, there would be a 100% match. However, with analog transmissions, this will almost never happen. Considerable image degradation occurs as soon as the software, eg, MMSSTV, readies the image for transmission and losses are magnified as the signal encounters path loss when it moves through the ionosphere.

Aspects of this degradation problem appear in a research article, *SSTV Transmission Methodology*, written by this author.

The article here is presented so that hams may consider quantitative and qualitative aspects of signal reception and better understand factors of signal degradation. Then steps may be taken to significantly enhance signal reception for given transmission and ionospheric conditions.

ImageMagick is a powerful software package that can compare two images pixel-by-pixel, e.g., a 320x256 pixel JPG image from an SSTV transmission. This is a total of 81,920 pixels in many SSTV pictures. The ratio of matches to the total is expressed as a percentage. This may be viewed as a **measure of picture quality** received.

My son, Matthew, developed a routine so that any ham can upload pictures and quickly perform picture quality calculations. This program is integrated into my ham web site. It may be found in the L&M section at wb9kmw.com.

In practice there is significant deterioration. A 'fuzz' adjustment is available in ImageMagick. Fuzz is used to match colors which are **close** to the target colors in RGB space. Colors within this distance are considered equal. Fuzz may be expressed in absolute intensity units, or as we will do for SSTV image comparisons, expressed as a percentage of maximum possible intensity value of each pixel.

Without this fuzz factor, most SSTV received images would indicate virtually no exact match of pixels. To the naked eye, that would give a misleading and very poor measure of picture quality.

For example, numerous tests were made at various **fuzz levels**. Here is a typical example. The original SSTV transmission is on the left. The other picture was received 897 miles away by AE5TC in Texas.




Picture quality, or the identicalness of pixels, at various fuzz levels were calculated as:

Looking back at that image received by AE5TC, a 0.11% picture quality does not make sense. And further, this will not make sense when we later discuss the P-scale, a broad picture quality scale that is sometimes used by broadcast engineers to describe video quality.

Likewise, it starts to become a stretch to claim that this received image possesses 80.38% picture quality. So one may compromise with the arguable but reasonable 10% fuzz level as a normal default when expressing picture quality. All further discussion in this article will always assume use of a 10% fuzz factor.

Picture Fuzz Quality	
0%	0.11%
10%	63.84%
20%	80.38%
30%	87.94%
40%	93.02%

One begins to distort an image as soon as it is prepared for transmission. MMSSTV is a popular Slow Scan TV software program. It will take a stock image and adjust it as it is readied for transmission. In these examples, this JPG image is selected as the original photo.



Here are the ‘internal loop’ images as they have been readied for Martin 1 (from left), Martin 2 and Scottie 2 modes. Resulting picture qualities are, respectively, **90.30%**, **88.19%** and **92.11%** of the above original image. So there is 8 to 12 percent deterioration, and they have not even been sent to the soundcard yet for transmission over the air!



For this reason, it is better when making over-the-air comparisons to use the ‘internal image’ just as it has been prepared for transmission for the chosen SSTV mode.


Some hams do not **properly align their soundcard**. This introduces another element of image distortion before the RF signal ever leaves the antenna. A series of transmission studies were conducted over a brief 11-minute period, starting with perfect Tx image alignment and progressively altering the slant.

The received images from Steve NS3L's SSTV cam were analyzed. Here are received images for straight, 1 Hz and 2 Hz slants, yielding picture quality percentages of 72.71%, 60.95% and 45.86%, respectively.



Next the transmissions were further skewed 4 and 8 Hz. These pictures scored 27.52% and 16.74%. So even with rather uniform band conditions, just misaligning the transmission software creates significant distortion error. About the only thing going for Robin is that his ladder remains modestly straight!

Notice the call letters in the URL. When hams call CQ with such poor transmission alignment, others sometimes cannot make out the call sign of that station. Hence, their CQ goes unanswered.



Path loss from the ionosphere plays a crucial role in picture quality. Four signals from Pat 3B8FA were captured one evening from a distance of 9,872 miles over a 35 minute span. Pat later shared his original image which was used for these calculations. On the left is the original picture, one he uses frequently, and the first one received at WB9KMW with 62.31% picture quality.



Then in subsequent minutes he retransmitted, and three more pictures were received, which scored 52.60%, 27.93% and 11.62% as band conditions worsened. After that, transmissions failed to arrive.



On the receiving end, the manner in which ones **receiver and soundcard are configured** makes a big difference in decoding the SSTV signal. In this next experiment, the same transmission at the same time with the same Kenwood TS-590S transceiver was copied.

The difference: The image on the left received its audio from the headphone jack, fed it into the microphone jack of a Dell computer and used a stock RealTek soundcard to decode the signal.

With the better image on the right, audio came from the line output of the ACC2 jack in the Kenwood radio. This was fed as line input into an ultralow noise (-124 dB SNR) Sound Blaster ZxR soundcard mounted in an HP computer.

The image quality on the left is only 44.01% of the better one that was decoded on the right. So with superior configuration in this example, there is a 2.27 times improvement in quality (100/44.01). This is a 3.56 dB improvement, or in ham terms, approximately one-half S-unit better reception!



Subsequently, the Dell computer was moved from the Kenwood headphone jack to a **shared connection for line audio output from the ACC2 jack, and fed as line-in audio to the computer**. The most significant remaining difference was the RealTek soundcard, versus the Sound Blaster.

The difference in picture quality of a sample image was measured. Now the one on the left with a RealTek soundcard scores 90.87%, compared to the right image.



This analysis is not exhaustive, but merely comparative. Quality percentage can be greater or less, depending on the signal-to-noise ratio for the particular received transmission.

Yet having superior configuration can make an important difference in the quality of a QSO, as the following transmission from Terry EA3EWO displays. His original transmission is on the left, the RealTek fed via microphone in the middle and Sound Blaster ZxR on the right.

Now compared to the original transmission, the picture quality percentages are, respectively, 31.33% and 47.81%.




Sometimes one gets to see the **near extremes of reception** when observing signals as received from afar and published by SSTV cams. In this case, very good reception (91.58%) was obtained by VE6PW in Calgary, Canada (1,252 miles away) and very poor quality (13.49%) was an image detected by NX0S in Missouri (393 miles).



That brings us now to a **video quality rating system** that has been used for years, based on a **P scale**. That scale is:

- P5 = broadcast quality
- P4 = typical quality with slight or minimal noise
- P3 = usable but somewhat noisy
- P2 = barely usable with lots of noise
- P1 = barely see the text
- P0 = unusable

One needs to find some way to align the picture quality percentages with these six P-scale factors. This can be accomplished by utilizing a normal statistical distribution curve for this problem, and adopting the following percentages from major standard deviation breakpoints.



The following approach is proposed for relating broader ranges of picture quality percentage measurements to P signal levels. Through repeated trials and observations, this quality percentage/P signal relationship holds up fairly well for a 10% fuzz factor setting on ImageMagick.

Quality	Standard Deviation	P Signal	Description
100-97.7%	Greater than 2σ	5	Broadcast quality
97.7-84.1%	1σ to 2σ	4	Good, some noise
84.1-50%	0 to 1σ	3	Usable, noisy
50-15.9%	-1σ to 0	2	Barely use, noisy
15.9-2.3%	-2σ to -1σ	1	Barely see text
2.3-0%	Less than -2σ	0	Unusable

So let's see how well this scale works with a variety of images received when compared to the transmitted original, which will consistently appear on the left hand side.

P5: broadcast quality Some transmissions, such as the VE6PW reception (91.58%) shown on a previous page come very close to broadcast quality. However, even if noise bars are not apparent, there is enough pixel degradation to mark most images down to the P4 level. True P5 picture quality, which lies more than 2 standard deviations above the statistical mean, is indeed rare with analog SSTV.

I was able to active a perfect 100% match under very controlled conditions. This involved transmitting at 5 watts into a 40 meter ICOM 706 MKIIG transceiver in Martin 1 mode, with one very important adjustment: The frequency from the Kenwood TS-590S transceiver was lowered by 60 Hz until it matched the actual frequency of the ICOM rig. This can be determined by examining the 1200 sync tone in the MMSSTV software.



For a 320x256 px image, the signal-to-noise ratio = $10 \log (320 \times 256) = 49.13$ dB. This is the point where every pixel in the received image exactly matches every pixel in the original transmitted image. In the days of B&W 405 line television, studio quality images from iconoscope television cameras had a signal-to-noise quality of 55 dB in a fully lit studio.

One might conclude that in most instances a true P5 picture, measured as more than two SD above the mean is perhaps not achievable for analog SSTV, and like Camelot, represents but a state of idealized beauty, peacefulness and enlightenment.

P4: good, some noise This reception from N9UWE has just a hint of noise, and serves as a nice example of a P4 image with 87.70% picture quality.



P3: usable, noisy A P3 picture quality is more commonly received. Some with a bit more and some with a bit less noise than displayed on this representative picture as received by VE1DBM, measured at 63.40% picture quality.



P4 and P3 pictures are normally reasonably good. Once one ventures below 50% quality level, the images become sketchy, noisy and sometimes one might consider that there is marginal distinction between P2 and P1 levels. However, remaining with the theoretical statistical distribution, we will describe these lower quality images.

P2: barely use, noisy This image was received by VE6PW while the ionosphere was in the process of fading out for the evening. The measured quality is 25.78%.



P1: barely see text A few minutes later, VE6PW could hardly copy the next transmission. Now the picture quality is only 7.64%.



P0: unusable Most of the time, signals with a quality score more than two standard deviations below the mean, ie, lower than 2.3%, are impossible to fully copy. The MMSSTV software may lock onto the transmission for several seconds, but cannot maintain image copy for the full minute plus from most SSTV transmissions. While not quite P0, this example has a picture quality of 5.96% compared to the original.

Todd KD0TUU is only 250 miles from WB9KMW, so the station is normally in the skip zone and receives nothing at all.



Let's examine another received image trait: **the time stamp**. Many SSTV cams record a time stamp before publishing the picture on the Internet. How much does that distort the quality of the received image?

Let's consider the same received image with and without the time stamp inserted. Here they are, again with the originally transmitted image on the left.

The picture qualities are 93.22% with time stamp and 94.79% without. So in this example, the time stamp degraded the image by 1.57%. Almost all the time, SSTV cams include a time stamp, so it is difficult to avoid this added distortion of the originally transmitted image. The improvement is indeed small, ie, only 0.07 dB.



As discussed with the P5 image quality, an exact match between transmitter and receiver frequencies is important for soundcard frequency alignment. This allows black intensity to align precisely at 1500 Hz and white at 2300 Hz. Let's examine **how frequency mismatch affects the picture quality**, and indeed the visual look of the image to the naked eye.


This test image was transmitted 25, 50 and 75 Hz above and below the control frequency of otherwise perfectly aligned transmitter and receiver, and the effects were studied. A test was conducted for a 100 Hz frequency mismatch, and the MMSSTV software failed to decode the transmission. The AFC option was disabled for all tests. Martin 1 was the selected SSTV mode for enhanced fidelity.

Presented first are the received images when the transmitter was increased 25, 50 and 75 Hz above the perfectly matched frequency, beginning with the 25 Hz mismatch on the left.

With this skewing, the black tone, which is expected to be at 1500 Hz for a pure black pixel, comes in at 1525 Hz, for example, for the picture on the left. So the received image is decoded as being not quite as black. The more the mismatch, such as with 50 and 75 Hz offsets, the lighter the black becomes, and likewise the lighter all the pixel colors become during reception.


One will note at 75 Hz that MMSSTV fails to fully decode the image, and distortion occurs in the last few pixel rows. Perhaps you can see that the black in the upper left hand corner becomes progressively lighter, albeit slightly. You may notice the white in the first row becoming progressively broader as the transmitter frequency increases, and more of the orange bar turns towards yellow.

ImageMagick was used to calculate picture quality. It returned values of 94.27%, 92.27% and only 65.49% for the 75 Hz mismatch! The original is reproduced below the set of three for visual comparison.




One may visualize a 75 Hz error when viewing the spectral frequency display. Notice how the horizontal sync tone is considerably above its intended 1200 Hz target mark. Colors of darker intensity never get down to the 1500 black level, and the whiter intensities slop over their target 2300 Hz mark.

Because of this transmitter/receiver frequency mismatch, one gets the degraded picture that can be seen, and also can be measured for picture quality.



Next, turning to the frequencies 25, 50 and 75 Hz below the perfectly matched frequency, one sees the colors becoming progressively darker as the frequency mismatch increases. Picture qualities are recorded as 94.36%, 94.56% and 93.10%, a stunning contrast to the statistics generated for the progressively lighter images above! Now the black really starts to dominate in the upper color bar. The original is again reproduced below the set of three for visual comparison.



This last study shows the **importance of proper frequency tuning**. While it may not be such a factor for common SSTV QSOs on 14.230 MHz, when one participates in the World SSTV Club's 15 Meter Dash, where contestants transmit rapidly on unknown frequencies, one has but a +/- 75 Hz frequency tolerance in order to lock onto the transmitting station and successfully decode the signal.

If you suspect that your rig is not properly calibrated for frequency, you might read the author's article regarding *Rx Frequency Calibration*.

Up to this point, we have been studying images as decoded by the MMSSTV engine. But what about **other popular SSTV software programs**? Are they all equally effective?

To help answer this, MMSSTV, RX-SSTV, DM-780 and ChromaPix were operated simultaneously to decode SSTV transmissions. All operated on a fast HP computer using an ultralow noise Sound Blaster ZxR soundcard.


Here is an example with a reasonably strong signal. MMSSTV was found to consistently decode the best. So all picture quality statistics are relative to this strongest standard.

All four images look rather similar, but compared to MMSSTV on the top left, quality scores are as follows: RX-SSTV (which uses the MMSSTV engine): 98.52%, DM-780 (found inside Ham Radio Deluxe): 86.00% and ChromaPix (bottom right image): only 62.83%.

So for this particular study, MMSSTV is a mere 1.015 times better (100/98.52) than RX-SSTV, however, the differences are pronounced when comparing MMSSTV to DM-780 (1.163 times better) and ChromaPix (1.592 times better). These are 0.66 dB and 2.02 dB improvements using MMSSTV. Recall that 3 dB is one-half S-unit better reception. The value of MMSSTV will be even more evident when one looks at the weak image that is subsequently analyzed.



For a weak image, MMSSTV is found on the top left. Next is RX-SSTV: 73.21%. DM-780: 10.13%. ChromaPix failed to decode a single pixel. The clear winner for strong and weak signals is MMSSTV.



I hope this article inspires some hams to experiment and investigate other measurements of SSTV images, and perhaps use their findings to further improve their station set up for even better SSTV reception.

Larry Peterson WB9KMW @ WB9KMW.com