

Another Look at SSIM Image Quality Metric

Yuriy Reznik
Brightcove Inc.
Boston, USA
yreznik@brightcove.com

Abstract—We review the design of the SSIM quality metric and offer an alternative model of SSIM computation, utilizing subband decomposition and identical distance measures in each subband. We show that this model performs very close to the original and offers many advantages from a methodological standpoint. It immediately brings several possible explanations of why SSIM is effective. It also suggests a simple strategy for band noise allocation optimizing SSIM scores. This strategy may aid the design of encoders or pre-processing filters for video coding. Finally, this model leads to more straightforward mathematical connections between SSIM, MSE, and SNR metrics, improving previously known results.

Keywords—Image quality, PSNR, SSIM, SNR

I. INTRODUCTION

The Structural Similarity (SSIM) metric [1,2] has been around for nearly two decades and has become one of the most established and frequently used metrics for image and video quality analysis. Many research papers have followed, extending SSIM in multiscale, 3D, and temporal dimensions [3-6], analyzing its mathematical properties [7-9], and discussing its applications for improving the design of encoding algorithms [10-13] and streaming systems [14-16].

Yet, despite all the progress and broad acceptance in practice, what specifically SSIM means from a physical, mathematical, and signal processing standpoint is not completely understood. The original papers [1,2] explain SSIM as a generalized mean of 3-types of distortion criteria: changes in luminance, contrast, and structure. However, in the final formula in [1,2], some factors cancel out, producing a somewhat mysterious fraction

$$\frac{2\sigma_{xy}}{\sigma_x^2 + \sigma_y^2},$$

which no longer represents a proper measure of correlation (structure) or change in contrast.

Adding to the mystery, several suggestions have been made over the years that there must be some simple relation between SSIM and other distortion metrics, such as PSNR or MSE [7,8]. For instance, in 2010, A. Horé and D. Ziou [8] suggested that:

$$\frac{1 - SSIM}{SSIM} \sim \frac{MSE}{2\sigma_{xy}},$$

where MSE is a mean square error and σ_{xy} is a covariance between image patches x and y . However, since σ_{xy} depends on both images, this formula does not reduce SSIM (as a distance metric between x and y) to MSE!

In this paper, we will look at SSIM again and offer an alternative model for its computation, leading to a more straightforward interpretation of the SSIM metric and its

connection to MSE, SNR, and some other metrics. In Section II, we will bring definitions, introduce our main results, and discuss their consequences. In Section III, we will perform experimental validation of our proposed model for computing SSIM.

II. MAIN RESULTS AND THEIR CONSEQUENCES

A. SSIM definition

Recall that at a patch level, the SSIM between images x and y is defined as follows [1]:

$$SSIM(x, y) = \frac{2\mu_x\mu_y + C_1}{\mu_x^2 + \mu_y^2 + C_1} \cdot \frac{2\sigma_{xy} + C_2}{\sigma_x^2 + \sigma_y^2 + C_2} \quad (1)$$

where $\mu_x, \mu_y, \sigma_x, \sigma_y, \sigma_{xy}$ represent patch-level statistics:

$$\mu_x = \mathbf{E}[x] = \sum_{k=1}^N w_k x_k, \quad (2)$$

$$\sigma_x^2 = \mathbf{E}[(x - \mu_x)^2] = \sum_{k=1}^N w_k (x_k - \mu_x)^2, \quad (3)$$

$$\sigma_{xy} = \mathbf{E}[(x - \mu_x)(y - \mu_y)] = \sum_{k=1}^N w_k (x_k - \mu_x)(y_k - \mu_y), \quad (4)$$

and C_1, C_2 are some small constants.

In all these formulae, $\mathbf{E}[\cdot]$ denotes expectation operators with some density w superimposed over the patch. The original SSIM implementation [1] uses 11x11-pixels patches and circular-symmetric Gaussian density with a standard deviation $\sigma = 1.5$. The weights w_k are normalized: $\sum_i w_i = 1$.

The constants C_1, C_2 in the original SSIM design [1] are set to $C_1 = (0.01 \cdot I_{\max})^2$, and $C_2 = (0.03 \cdot I_{\max})^2$, where I_{\max} is the maximum pixel value (e.g., 255 for 8-bit pixels). As explained in [1], the purpose of these constants is to avoid numerical instabilities in cases when signals approach 0.

Considering the whole images, the mean SSIM is computed as the average of patch-level SSIMs at each pixel location (i, j) :

$$\overline{SSIM}(x, y) = \frac{1}{WH} \sum_{i=1}^W \sum_{j=1}^H SSIM(x(i, j), y(i, j)). \quad (5)$$

W and H denote image width and height, respectively.

B. Proposed Alternative Form

By looking at formula (1), we first notice that SSIM is essentially a product of two nearly identical functions:

$$\xi(x, y) = \frac{2\mathbf{E}[xy] + C}{\mathbf{E}[x^2] + \mathbf{E}[y^2] + C} \quad (6)$$

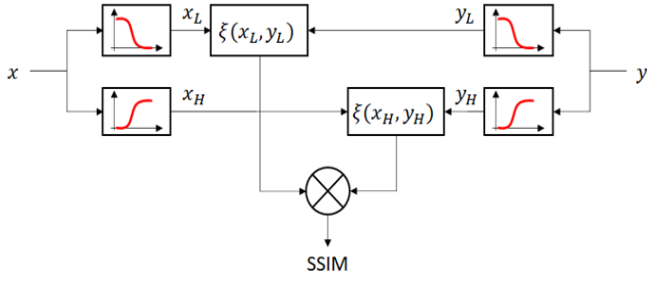


Fig. 1. Computation of SSIM by using formulae (7). SSIM is presented as a product of two identical distance measures $\xi(\cdot, \cdot)$ computed for low-pass and high-pass filtered versions of input signals x, y .

applied to different signals and at different scales. In the first term in the SSIM expression (1), this operation is applied to DC values μ_x, μ_y , treated as scalars ($N = 1$). In the second term in (1), it is applied to residual signals $x - \mu_x$, and $y - \mu_y$, observed in 11×11 patches ($N = 11^2$).

However, if we examine the derivation of DC values μ_x, μ_y in (1,2), we notice that they can also be understood as pixel values taken from some low-pass-filtered versions of input images x_L, y_L . Similarly, we also realize that the residual signals in each path $x - \mu_x$, and $y - \mu_y$ must be similar to signals taken from high-pass versions of the same images: $x_H = x - x_L$, $y_H = y - y_L$. And finally, we also notice that since DC values are coming from low-pass filtered images x_L, y_L , then the relative distance between them as scalars $\xi(\mu_x, \mu_y)$ or computed over surrounding patches $\xi(x_L, y_L)$ must be very small. The low-pass removes local variations making the patch-average results almost identical.

In other words, by combining all these observations, we can conjecture that patch-level SSIM can be computed in a *fully symmetric manner*, as follows

$$SSIM(x, y) \approx \xi(x_L, y_L) \cdot \xi(x_H, y_H) \quad (7)$$

where x_L, y_L are the patches in low-pass filtered images, $x_H = x - x_L$, $y_H = y - y_L$ are the patches in high-pass filtered images, and where $\xi(\cdot, \cdot)$ are identical distance functions (6) computed over patches in low-pass and high-pass images.

Figure 1 shows the flow-diagram explaining computations according to (7). In Section III, we present an experimental study, indicating that this process yields very similar results to the original SSIM formula. Our experiments further show that to split x, y into low-frequency and high-frequency components, it is sufficient to apply Gaussian filter with pixel-level standard deviation $\sigma = 3$. This value is 2x larger than the standard deviation used in windows for computing patch-level distances $\xi(x_L, y_L), \xi(x_H, y_H)$.

C. Illustration of operation

One of the immediate benefits of the proposed model (7) is a simple signal processing interpretation of how SSIM works, and why it is better than full-band metrics, such as MSE or PSNR. We illustrate this in Figure 2.




Fig. 3. Superimposed plots of Contrast Sensitivity Function CSF (black) and low- and high-pass filter responses (red and blue) in our model of computation of SSIM. Display Nyquist is assumed to be 40 cpd.

As input x in this example, we use image k04 from the Kodak data set [17]. The low-pass (x_L) and high-pass (x_H) versions of this image are shown in the top row in Figure 2. The input y in this example is a reconstruction of the same image after it was compressed with H.264 encoder [18] with QP=47. This image is highly distorted relative to the original. The low-pass (y_L) and high-pass (y_H) versions of this image are shown in the middle row. Finally, the last row shows the differences between these images.

Looking at the last row, we immediately notice that the direct difference between input images $x - y$ looks much "busier" compared to the difference images in low- and high-frequency domains. We see that magnitudes of errors in each subband domain become lower (particularly in the low-band), and their impacts become more obvious visually and conceptually. The low-pass filtering captures image changes in overall shapes, while high-pass shows differences in fine-grain details – contours, textures, etc.

In other words, we can see that the principal difference between SSIM and PSNR, MSE, and other simple metrics is that SSIM analyzes images in two subbands. Subband processing allows the separation of errors such that their impacts can be more accurately measured and incorporated in the final score.

D. Connection to CSF

Let us next look at the split between bands used in our model for computing SSIM. The pixel-level standard deviation of the Gaussian low-pass filter used to produce x_L and y_L is $\sigma = 3$. Consequently, the cut-off frequency of this filter is $f_c = \frac{F_s}{2\pi\sigma} \approx \frac{F_s}{18.85}$, where F_s is the sampling rate. By assuming that display Nyquist frequency approximately matches the human visual acuity limit (which is typically the case for ITU-R BT-500 [20] visual tests), this implies that F_s must be about 80 cpd (cycles per degree), and therefore $f_c \approx 4.244$ cpd.

As shown in Figure 3, this frequency approximately matches the peak of the Contrast Sensitivity Function (CSF) of human vision. This figure uses the Barten CSF model [19] and frequency responses of Gaussian low-pass and high-pass filters employed in our model for computing SSIM. Barten CSF model rendered for an observation angle $X_0 = 30^\circ$ and object luminance $L = 200$ cd/m².



Fig. 2. Visual illustration of steps in the proposed model for computation of SSIM. The top row shows an input image x and its low-pass x_L and high-pass $x_H = x - x_L$ versions. The middle row shows the second input image y , and its low-pass y_L and high-pass $y_H = y - y_L$ versions. This image is an encoded and decoded version of x , with significant distortions introduced by the codec. The bottom row shows the differences between the original images, as well as their filtered versions. It can be observed that full difference is much “busier” and less descriptive than low- and high-pass differences. Low-pass emphasizes changes in overall shapes, while high-pass emphasizes changes in textures, contours, and other fine features.

This observation brings another argument suggesting why SSIM may be effective. It does a proper signal separation considering the involved mechanisms of human vision. Recall that the decays in contrast sensitivity in low- and high-frequency bands are caused by significantly different phenomena. In the high-band, sensitivity decay is mainly caused by the eye optical MTF [22]. In the low-band, it is mainly the result of lateral inhibition [22,23]. The latter is a non-linear effect. The processing of low-band and high-band signals in the visual cortex is also different, as they turn into different sets of spatial frequency channels [23].

Even though in their motivations in [1], the authors of SSIM intended to design a metric that is not explicitly driven by the CSF and related phenomena of human vision, the use of patches and particular filters in the computation of SSIM, as we have shown above, leads back to this connection!

E. Impacts of differences in low- and high-frequency bands

By denoting each factor in SSIM expression (7) as $\xi_L = \xi(x_L, y_L)$, and $\xi_H = \xi(x_H, y_H)$, we can rewrite it as:

$$SSIM(\xi_L, \xi_H) = \xi_L \cdot \xi_H,$$

By using inequality between geometric and $-\infty$ means, and also noting that $\xi_L, \xi_H \leq 1$, we can show that:

$$\min(\xi_L, \xi_H)^2 \leq SSIM(\xi_L, \xi_H) \leq \min(\xi_L, \xi_H) \quad (8)$$

This inequality means that it is the smallest between subband differences ξ_L and ξ_H that has a limiting impact on the overall SSIM score!

Consequently, this also means, is that if one tries to balance codec-introduced errors to maximize the overall SSIM value, then the best such balance would be achieved when the magnitudes of such errors in both bands are the same:

$$\delta: \xi_L = \xi_0 + \delta, \xi_H = \xi_0 - \delta \Rightarrow SSIM(\xi_L, \xi_H) \Rightarrow \xi_L = \xi_H. \quad (9)$$

This observation offers a simple principle for encoder optimizations improving SSIM scores. This idea may also lead to the design of a pre-processing filter that removes some of the low-frequency features and thus allows the encoder to encode the remaining high-frequency content with fewer errors. With proper tuning, such a filter may allow existing encoders to achieve improved SSIM scores.

F. The relation between SSIM and other objective metrics

Finally, let us now take a closer look at quantities $\xi(x, y)$ in formula (7). For conceptual simplicity, we will discard constant C and will look instead at pure ratios:

$$\xi(x, y) = \frac{2\mathbf{E}[xy]}{\mathbf{E}[x^2] + \mathbf{E}[y^2]} \quad (10)$$

We first note that ratio $\xi(x, y)$ is not a proper measure of correlation. To measure correlation, one has to apply normalization by the geometric mean of $\mathbf{E}[x^2]$ and $\mathbf{E}[y^2]$:

$$\rho(x, y) = \frac{\mathbf{E}[xy]}{\sqrt{\mathbf{E}[x^2] \mathbf{E}[y^2]}} \quad (11)$$

But these quantities are related. By inequality between geometric and arithmetic means, we can see that:

$$\xi(x, y) \leq \rho(x, y). \quad (12)$$

Next, by looking at the reciprocal of (9), we observe that:

$$\frac{1}{\xi(x, y)} = \frac{\mathbf{E}[x^2] + \mathbf{E}[(x + (y - x))^2]}{2\mathbf{E}[xy]} = 1 + \frac{\mathbf{E}[(y - x)^2]}{2\mathbf{E}[xy]}$$

or, equivalently:

$$\frac{1 - \xi(x, y)}{\xi(x, y)} = \frac{MSE(x, y)}{2\mathbf{E}[xy]} \quad (13)$$

where $MSE(x, y) = \mathbf{E}[(x - y)^2]$ is the Mean Square Error.

The obtained formula (11) is similar to the result of A. Horé and D. Ziou [8], with an essential distinction that (13) is defined for subband signals. In their derivations, Horé and Ziou have assumed that DC differences are negligible (or equivalently, that $\xi_L = 1$), but this is not always the case in practice. Equation (13) establishes a more general and accurate relation.

However, as we already noted, formula (13) does not reduce $\xi(x, y)$ to $MSE(x, y)$, since $\mathbf{E}[xy]$ is a joint statistic of both signals. To produce a more direct relation, we rewrite (10) as:

$$\xi(x, y) = 1 - \frac{\mathbf{E}[(y - x)^2]}{\mathbf{E}[x^2] + \mathbf{E}[y^2]}$$

and then:

$$\frac{1}{1 - \xi(x, y)} = SNR(x, y) + SNR(y, x) \quad (14)$$

where $SNR(x, y) = \frac{\mathbf{E}[x^2]}{\mathbf{E}[(x-y)^2]}$ is the Signal to Noise Ratio.

Combining (14) and (7), we produce:

$$SSIM(x, y) \sim \left(1 - \frac{1}{SNR(x_L, y_L) + SNR(y_L, x_L)}\right) \cdot \left(1 - \frac{1}{SNR(x_H, y_H) + SNR(y_H, x_H)}\right) \quad (15)$$

where $SNR(x_L, y_L)$, $SNR(y_L, x_L)$, $SNR(x_H, y_H)$, $SNR(y_H, x_H)$ are signal-to-noise ratios computed for patches in low-pass and high-pass-filtered images respectively.

The formula (15) shows that SSIM can be computed using a combination of SNR metrics between low- and high-pass-filtered images. These derivations confirm that connections between SSIM and other objective metrics exist, but they are not as simple as suggested earlier [7,8].

III. EXPERIMENTAL VALIDATION

In this section, we describe tests performed to validate the accuracy of the proposed method for the computation of SSIM.

For our experiments, we used 24 standard images from Kodak dataset [17]. We have converted them to BT.709 YUV files and then produced encoded and decoded versions using

TABLE I. REFERENCE SSIM VALUES AND DIFFERENCES BETWEEN THE REFERENCE AND SSIM COMPUTED BY THE PROPOSED METHOD (7)

| File | QP=17 | | QP=22 | | QP=27 | | QP=32 | | QP=37 | | QP=42 | | QP=47 | |
|------|--------|---------|--------|---------|--------|---------|--------|---------|--------|---------|--------|---------|--------|---------|
| | SSIM | Delta | SSIM | Delta | SSIM | Delta | SSIM | Delta | SSIM | Delta | SSIM | Delta | SSIM | Delta |
| k01 | 0.9949 | 0.0003 | 0.9867 | 0.0008 | 0.9697 | 0.0016 | 0.9337 | 0.0028 | 0.8687 | 0.0038 | 0.7766 | 0.0026 | 0.6648 | -0.0026 |
| k02 | 0.9928 | 0.0000 | 0.9784 | 0.0001 | 0.9447 | 0.0001 | 0.8962 | -0.0009 | 0.8318 | -0.0033 | 0.7569 | -0.0069 | 0.6843 | -0.0119 |
| k03 | 0.9879 | 0.0000 | 0.9768 | -0.0001 | 0.9669 | -0.0003 | 0.9506 | -0.0008 | 0.9251 | -0.0016 | 0.8947 | -0.0033 | 0.8617 | -0.0059 |
| k04 | 0.9906 | -0.0001 | 0.9782 | -0.0004 | 0.9603 | -0.0009 | 0.9327 | -0.0018 | 0.8937 | -0.0034 | 0.8453 | -0.0063 | 0.7928 | -0.0108 |
| k05 | 0.9941 | 0.0003 | 0.9869 | 0.0005 | 0.9757 | 0.0007 | 0.9551 | 0.0008 | 0.9168 | 0.0004 | 0.8552 | -0.0015 | 0.7619 | -0.0067 |
| k06 | 0.9930 | 0.0001 | 0.9851 | 0.0003 | 0.9709 | 0.0006 | 0.9427 | 0.0006 | 0.8920 | 0.0000 | 0.8126 | -0.0027 | 0.7108 | -0.0076 |
| k07 | 0.9889 | -0.0002 | 0.9781 | -0.0004 | 0.9707 | -0.0008 | 0.9599 | -0.0016 | 0.9427 | -0.0030 | 0.9173 | -0.0056 | 0.8811 | -0.0087 |
| k08 | 0.9956 | 0.0003 | 0.9853 | 0.0010 | 0.9626 | 0.0022 | 0.9326 | 0.0030 | 0.8876 | 0.0030 | 0.8180 | 0.0019 | 0.7076 | -0.0030 |
| k09 | 0.9893 | 0.0000 | 0.9623 | 0.0001 | 0.9320 | -0.0001 | 0.9180 | -0.0006 | 0.8989 | -0.0016 | 0.8708 | -0.0038 | 0.8376 | -0.0065 |
| k10 | 0.9898 | 0.0000 | 0.9645 | 0.0001 | 0.9328 | -0.0004 | 0.9149 | -0.0014 | 0.8931 | -0.0031 | 0.8649 | -0.0055 | 0.8321 | -0.0084 |
| k11 | 0.9918 | 0.0002 | 0.9799 | 0.0004 | 0.9611 | 0.0006 | 0.9264 | 0.0004 | 0.8701 | -0.0012 | 0.8019 | -0.0044 | 0.7361 | -0.0083 |
| k12 | 0.9887 | 0.0000 | 0.9776 | -0.0002 | 0.9646 | -0.0005 | 0.9443 | -0.0015 | 0.9169 | -0.0029 | 0.8888 | -0.0049 | 0.8673 | -0.0063 |
| k13 | 0.9967 | 0.0001 | 0.9911 | 0.0003 | 0.9753 | 0.0006 | 0.9363 | 0.0011 | 0.8557 | 0.0005 | 0.7283 | -0.0034 | 0.5750 | -0.0123 |
| k14 | 0.9932 | 0.0001 | 0.9826 | 0.0001 | 0.9637 | -0.0001 | 0.9281 | -0.0014 | 0.8704 | -0.0046 | 0.7942 | -0.0102 | 0.7149 | -0.0162 |
| k15 | 0.9909 | -0.0001 | 0.9805 | -0.0003 | 0.9645 | -0.0007 | 0.9393 | -0.0017 | 0.9069 | -0.0033 | 0.8707 | -0.0050 | 0.8339 | -0.0073 |
| k16 | 0.9907 | 0.0000 | 0.9804 | 0.0000 | 0.9677 | -0.0001 | 0.9433 | -0.0003 | 0.8989 | -0.0009 | 0.8316 | -0.0029 | 0.7461 | -0.0078 |
| k17 | 0.9908 | -0.0002 | 0.9687 | -0.0007 | 0.9367 | -0.0018 | 0.9121 | -0.0035 | 0.8840 | -0.0055 | 0.8510 | -0.0079 | 0.8090 | -0.0101 |
| k18 | 0.9945 | 0.0000 | 0.9820 | -0.0001 | 0.9431 | -0.0004 | 0.8858 | -0.0015 | 0.8310 | -0.0038 | 0.7617 | -0.0077 | 0.6729 | -0.0136 |
| k19 | 0.9913 | 0.0001 | 0.9741 | 0.0000 | 0.9555 | 0.0000 | 0.9282 | 0.0002 | 0.8788 | -0.0003 | 0.8172 | -0.0023 | 0.7594 | -0.0041 |
| k20 | 0.9889 | 0.0000 | 0.9739 | -0.0001 | 0.9632 | -0.0001 | 0.9453 | -0.0003 | 0.9128 | -0.0011 | 0.8726 | -0.0033 | 0.8375 | -0.0052 |
| k21 | 0.9912 | 0.0000 | 0.9701 | 0.0001 | 0.9484 | 0.0001 | 0.9269 | 0.0001 | 0.8918 | -0.0003 | 0.8395 | -0.0024 | 0.7723 | -0.0082 |
| k22 | 0.9913 | -0.0001 | 0.9765 | -0.0002 | 0.9562 | -0.0006 | 0.9225 | -0.0017 | 0.8716 | -0.0040 | 0.8097 | -0.0076 | 0.7471 | -0.0108 |
| k23 | 0.9877 | -0.0004 | 0.9705 | -0.0010 | 0.9575 | -0.0017 | 0.9441 | -0.0026 | 0.9286 | -0.0037 | 0.9107 | -0.0046 | 0.8900 | -0.0048 |
| k24 | 0.9928 | 0.0002 | 0.9846 | 0.0003 | 0.9718 | 0.0002 | 0.9471 | -0.0004 | 0.9041 | -0.0024 | 0.8395 | -0.0069 | 0.7560 | -0.0139 |
| RMS | 0.9915 | 0.0002 | 0.9781 | 0.0004 | 0.9591 | 0.0009 | 0.9321 | 0.0016 | 0.8909 | 0.0028 | 0.8360 | 0.0052 | 0.7727 | 0.0091 |

H.264 video encoder [18], operating in Main Profile, "slow" encoding preset, and using fixed-QP rate control mode.

Table 1 shows the results. The columns "SSIM" list reference SSIM values computed as described in [1] for each file and QP setting, and the column "Delta" lists differences between the reference SSIM and the one calculated by our method. For band separation, we used a circular-symmetric Gaussian filter with $\sigma = 3$. The last row in Table 1 lists RMS average values across all files.

It can be observed that the differences between SSIM produced by using our model and reference are very small, despite the broad range of variation of SSIM values, distortion, and content types. Table 1 provides results for images at 1536x1024 resolutions. The results for other resolutions are very similar and omitted due to paper size restrictions.

REFERENCES

- [1] Z. Wang, A. Bovik, H. Sheikh, E. Simoncelli, "Image quality assessment: from error visibility to structural similarity". IEEE Transactions on Image Processing, vol. 13, no. 4, pp. 600–612, 2004.
- [2] Z. Wang and A. C. Bovik, "A universal image quality index," IEEE Signal Processing Letters, vol. 9, no 3, pp. 81–84, 2002.
- [3] Z. Wang, E. P. Simoncelli and A. C. Bovik, "Multiscale structural similarity for image quality assessment," 37th Asilomar Conference on Signals, Systems & Computers, 2003, vol.2, pp. 1398-1402.
- [4] K. Zeng and Z. Wang, "3D-SSIM for video quality assessment," 19th IEEE International Conference on Image Processing, 2012, pp. 621-624.
- [5] AK Moorthy and A. C. Bovik. "Efficient motion weighted spatio-temporal video SSIM index," Human Vision and Electronic Imaging XV, vol. 7527, pp. 440-448. SPIE, 2010.
- [6] Y. Wang, T. Jiang, S. Ma, and W. Gao. "Spatio-temporal SSIM index for video quality assessment," 2012 Visual Communications and Image Processing, pp. 1-6. IEEE, 2012.
- [7] R. Dosselmann and X. D. Yang, "Existing and emerging image quality metrics," Proceedings of the Canadian Conference on Electrical and Computer Engineering, pp.1906-1913, 2006.
- [8] A.Horé, D. Ziou, "Image quality metrics: PSNR vs. SSIM," ICPR 2010.
- [9] D. Brunet, E.R. Vrscay, and Z. Wang, "On the mathematical properties of the structural similarity index," IEEE Transactions on Image Processing, 21(4), 2011, pp.1488-1499.
- [10] T. Richter and K. J. Kim, "A MS-SSIM Optimal JPEG 2000 Encoder," 2009 Data Compression Conference, 2009, pp. 401-410.
- [11] F. N. Rahayu, U. Reiter, T. Ebrahimi, A. Perkis, and P. Svensson, "SS-SSIM and MS-SSIM for digital cinema applications," Proc. SPIE 7240, Human Vision and Electronic Imaging XIV, 72400P (10 February 2009)
- [12] S. Wang, A. Rehman, Z. Wang, S. Ma and W. Gao, "SSIM-MotivatedRate-Distortion Optimization for Video Coding," in IEEE Transactions on Circuits and Systems for Video Technology, vol. 22, no. 4, pp. 516-529, April 2012
- [13] W. Wu and X. Zhang, "Code performance improvement scheme for X264 based on SSIM," 2012 3rd IEEE International Conference on Network Infrastructure and Digital Content, 2012, pp. 396-400.
- [14] Z. Wang, K. Zeng, A. Rehman, H. Yeganeh, and S. Wang, "Objective video presentation QoE predictor for smart adaptive video streaming," Proc. SPIE 9599, Applications of Digital Image Processing XXXVIII, 95990Y, 22 September 2015.
- [15] Y. Reznik, K. Lillevold, A. Jagannath, J. Greer, and J. Corley, "Optimal design of encoding profiles for ABR streaming," Proc. Packet Video Workshop, Amsterdam, The Netherlands, June 12, 2018.
- [16] Y. Reznik, "Average Performance of Adaptive Streaming," Proc. Data Compression Conference (DCC'21), Snowbird, UT, March 2021.
- [17] Eastman Kodak Company, Kodak lossless true color image suite (PhotoCD PCD0992).
- [18] x264 encoder project, <https://www.videolan.org/developers/x264.html>.
- [19] P. G. J. Barten, "Formula for the contrast sensitivity of the human eye," Proc. SPIE 5294, Image Quality and System Performance, 18 December 2003.
- [20] Rec. ITU-R BT.500-13, "Methodology for the subjective assessment of the quality of television pictures," ITU-R, 2014.