https://doi.org/10.2352/ISSN.2470-1173.2017.18.COLOR-054

Extraordinary perceptual color stability in low cost, real time color image compression inspired by structure tensor analysis

Fritz Lebowsky, Mariano Bona; STMicroelectronics; Grenoble, France

Abstract

Cameras as well as displays of mobile phones, autonomously driven vehicles, PC monitors, and TVs continue to increase their native resolution to 4k by 2k and beyond. At the same time, their high dynamic range formats demand higher bit depth for the underlying color component signals. Subsequently, uncompressed pixel amplitude processing becomes costly not only when transmitting over cable or wireless communication channels, but also across on-chip image processing pipelines that access external memory units. In 2016 we introduced a low cost, real time, visually lossless color image compression concept inspired by structure tensor analysis which promises a highly adaptive and robust compression performance across a substantial range of compression ratios (between 1x and 3x) without significantly compromising perceptual image quality. We also noticed surprisingly strong perceptual color stability in spite of having processed each color component independently in RGB color space.

To manage a wider range of compression ratios as well as visually lossless image quality, we proposed a novel approach that converts image amplitudes into a pair of discrete structure and magnitude quantities on a pixel-by-pixel basis which had been inspired by structure tensor analysis. Graceful degradation of image information is controlled by a single parameter which aims at optimally defining sparsity – as a function of image content. Furthermore, we applied error diffusion via a threshold matrix to optimally diffuse the residual coding error.

Strongly encouraged by these findings, we continued implementing a version which combines structurally similar elements across RGB color components. As a result, we already achieve visually lossless compression with compression ratios above 4x with 8bit gamma pre-corrected color component signals while having to only analyze 4 nearest neighbors per pixel. We believe to have well identified a conceptual explanation for the algorithm's extraordinary perceptual color stability which we would like to present and discuss in detail. We also provide a detailed error distribution analysis across a variety of well-known, full-reference metrics which highlights the effectiveness of our new approach, identifies its current limitations with regard to high quality color rendering, and illustrates algorithm specific visual artifacts.

Introduction

We are greatly interested in developing image data compression architectures that enable context sensitive control of visually noticeable artifacts as a function of compression ratio. Compared to well-known image compression methods such as JPEG or MPEG standards, which can achieve relatively high compression ratios, we aim at low compression ratios in the range of 1.5 to 4. The fundamental architectural challenge arises from an implementation at a fraction of the cost of most well-known compression methods.

In [1] we had already introduced a novel image data compression method inspired by structure tensor analysis [2], which itself emerged from parametric functional compression across nonlinear quantities in image amplitude domain [3, 4, 5]. Surprisingly, at moderate compression ratios between 2.0 and 2.5 we did not encounter any critical visual artifacts compromising color fidelity, although each color component had been processed independently in RGB color space. Encouraged by the absence of visually annoying color artifacts as well as prospecting higher compression ratios we became interested in combining structural quantities across all color components while remaining in RGB color space. Such approach also mediates several significant advantages: (1) less cost in architecture implementation – especially when input and output color spaces are the same, (2) increased adaptation towards human visual fidelity of detecting any possible color artifacts at higher compression ratios and (3) performance comparison with color spaces that process luminance and chrominance components separately.

In the remaining sections we first summarize the key components of our latest image compression algorithm enabling enhanced stability of perceptual color image quality, followed by a statistical analysis demonstrating significant performance characteristics of the compression algorithm while comparing cumulative error distribution functions of two different metrics in luminance domain, namely YPSNR, SSIM. Then we also compare performance characteristics of our color image compression algorithm using a traditional color aware metric, the CIE $\Delta E76$ metric. Finally, we draw a conclusion and propose related work still ahead.

Key components of the image compression algorithm which provide increased stability of perceptual color image quality

The overall functional architecture of the image compression algorithm emerges from the idea to combine several generic methods which could offer improved system stability. After having efficiently converted image amplitudes into a parametric pair of discrete structure & gradient magnitude quantities at pixel resolution, we apply a classifier, an integral error control loop and an error diffusion method based on a threshold matrix.

Structure-magnitude pairs enabling classification into 3 major classes

Here, we briefly summarize the principal processing steps that estimate a structure quantity and a magnitude quantity for each pixel in the image. Processing is done in sequential raster scan

order, column after column, from top left to bottom right across the entire image. **Figure 1** shows the image processing kernel with its center pixel **p(0,0)** and its 4 local neighbors {p(-1,-1), p(0,-1), p(1,-1), p(-1,0)}which have already also been processed by the encoder.

Across all 4 local neighbors we determine the minimum gradient magnitude together with its maximum gradient magnitude in perpendicular direction and normalize the minimum gradient magnitude by the sum of minimum gradient magnitude and maximum gradient magnitude. Then, as illustrated in **Figure 2**, we (1) classify structure quantities in normalized gradient domain into three major categories, (a) contours (trajectories), (b) zones (homogeneous regions), and (c) extremes (local minimum or maximum); (2) operate an IntegralError control loop for each class separately; (3) prioritize the major categories in the order given above, (a) contours, (b) zones, and (c) extremes.

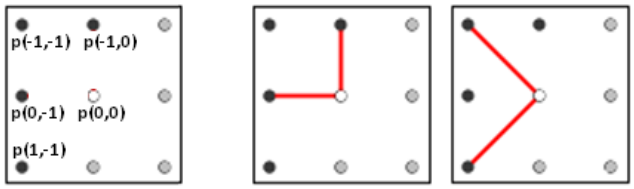

Figure 1. Image processing kernel with center pixel p(0,0) and its relevant four local neighbor pixels (left) in relative coordinates; illustration of possible min/max gradient pairs in 4-pixel neighborhood: horizontal-vertical (center) and diagonal (right).

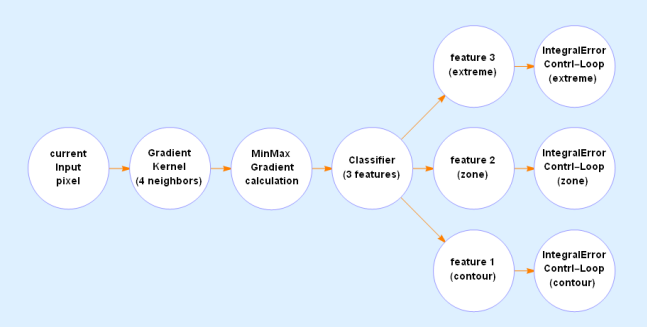

Figure 2. Preprocessing of pixel amplitudes to obtain structure/gradient magnitude quantities that will be fed to 3 individual IntegralError Control-Loops corresponding to 3 distinct classes (contour, zone and extreme values).

Integral error control loop

Figure 3 illustrates the building blocks of our integral error control loop. The resulting two normalized minimum gradient magnitude values **mG1** and **mG2** (of the MinMax Gradient calculation stage) will be passed on to the integral error control unit and quantized by the compressor's non-linear quantizer.

The quantizer itself is also configured by the dynamic Error density control as well as a local mean value that represents the intensity of the local image region (adaptation to perceptual tone curve phenomena). Together with the residual error **Err(n-1)** obtained from previous pixel encoding, we obtain a new current Error value **Err(n)** by selecting the normalized minimum gradient magnitude value which, for example, creates the minimum absolute value of the Sum of the associated center pixel amplitude after decode, Err(n-1), and the 'negated' original center pixel amplitude. Please also note that the decoder's functionality is entirely simulated within the compressor's encoder. Therefore, we actually estimate and minimize the current local Error value by also analyzing the mean absolute error MinAbsErr as well as the minimum local variance MinLocVar in differential context across the local neighborhood of 4 pixels. Not only do we consider the local neighborhood of input quantities, but we also take into account the encoded quantities across their local neighborhood (4 pixel positions).

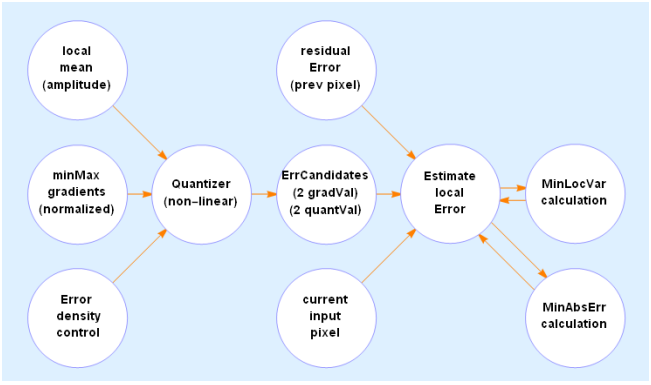

Figure 3. Estimation & minimization of local error within the integral error control loop.

Generic density control and anchor method

We would like to also introduce two principal functional components which play a significant role in stabilizing perceptual image quality, namely, error ***Density control*** and ***Anchor method*** along contours and zones. Both components use a simple threshold matrix control and can be easily generalized towards efficient re-use in application specific optimization tasks of almost any system parameter.

Density control

A Threshold Matrix offers a simple solution to controlling density in multidimensional context by using a single density parameter which is proportional to the desired overall density of any target quantity. **Figure 4** depicts a simplified example in spatial context. The Threshold Matrix is being repeatedly applied across the pixel raster of the image. A density parameter ranging from 0 to 4 covers 5 different levels of spatial density. For example, if the density parameter is set to 2, we select all the positions that are showing a value less than or equal to 2. In a second step we set the output value of all selected positions to one while values of the unselected positions remain zero. Now, we can modulate a desired target quantity by multiplying it with the output value in a pixel-by-pixel regime.

1	3	1	3	1	3	1	3	1
4	2	4	2	4	2	4	2	4
1	3	1	3	1	3	1	3	1
4	2	4	2	4	2	4	2	4

Figure 4. A highly simplified representation of a threshold matrix to demonstrate the concept of generic density control (see text for a detailed description).

In our image data compression task, for example, we also control the discrete representation of the normalized gradient magnitude and the structure orientation by such a density

parameter that translates into variable precision of signal quantization by just adding the value of the Threshold Matrix to the target quantity after normalization.

Anchor method

We baptized this method 'anchor method' because we literally anchor our compression algorithm at selectable pixel positions to prevent any undesirable drift in differential domain. The anchor method is considered to be a significant contributor to extraordinary perceptual color stability. Using the underlying density control as a function of compression effort lets us select pixel positions in the image where we decide to transmit the original image amplitude values. The anchor method is applied to contours and zones and, therefore, assists the contour and zone classes in achieving augmented compression performance.

A noteworthy side effect of the anchor method emerged with the fact that we can also control the lower bound of visual image quality by defining a lower limit of the associated density control parameter. Consequently, each class can control its contribution to visual image quality independently or relative to each other's performance goal.

Performance analysis of perceptual color stability

Surprisingly, we could not yet reveal critical visually noticeable color related artifacts, even at compression ratios beyond 3x. In other words, color fidelity was never compromised although each color component had been processed independently in RGB color space. We hypothesize that this is due to graceful pruning of prioritized local contrast quantities in context of visual masking. Such graceful performance seems to be enabled since we eliminated (1) any convolutional linear filtering which would increase visibility of blur, eliminated (2) visibility of blockiness and (3) minimized visibility of contouring.

Our underlying basic hypothesis can be summarized as follows: structure quantities that appear coherent across all color component channels can be transmitted as a single structure quantity without introducing inconsistent changes in hue quantities along contours that are visually noticeable. Or, formulated slightly differently, we transmit only one contour orientation coefficient per pixel – if such contour orientation quantity enables consistent color representation along a contour.

We definitely felt attracted by the surprisingly robust performance of the novel image data compression method inspired by tensor analysis which was already well noticeable during very early algorithm elaboration phases. How could we better understand why the basic concept of combining structural quantities at pixel resolution combined with feature classification and error diffusion suddenly led to a superior overall compression ratio remained a dominantly recurring question.

First of all, we would like to recall that we classified the structure quantities pixel-by-pixel into three major categories: contours (trajectories), zones (homogeneous regions with approximately constant gradients) and extremes (local minimum or maximum not part of a contour). The classification was motivated by the hypothesis that we find a higher level of stationary signal characteristics within each class than all classes combined. Especially strong, non-stationary signal behavior can be found for example when encountering a high contrast contour perpendicular to the direction of sequential pixel processing when transitioning from one zone to the next. Such scenarios create significant variability of variance quantities. In other words, signal power varies significantly in local context.

Traditionally, such critical scenarios have been addressed by using, for example, Maximum Likelihood or Sigma filters [6] which exhibit satisfactory performance in various challenging image processing tasks. However, their cost efficient use in real time image data compression tasks is highly questionable. Since such filters are mostly working in magnitude domain (vs. structure orientation domain) they also do not easily offer a cost efficient strategy to minimize the integral (coding) error.

Based on the observations mentioned above we elaborated the idea that combines structure quantities and integral error control on a pixel-by-pixel regime. At the same time, each of the three classes possesses its own error control loop. As a result, non-stationary signal behavior has been significantly reduced due to independently and separately processing contours, zones, and extremes. In summary, we manage error quantities at pixel resolution that demonstrate significantly improved stationary signal behavior as well as minimum latency – enabling elevated system stability and robustness especially when processing images consisting of natural scene content.

DISCUSSION: We would like to point out that a pixel-by-pixel control by structure quantities appears to be an efficient concept in comparison with block based transforms such as DCT or Wavelet Transform. In our approach, discrete structure quantities remain the same independent of compression ratio. We are just increasing (or decreasing) quantization of gradient magnitude quantities to achieve the overall target compression ratio. To the best of our knowledge, block based transforms do not yet process structure quantities concurrently with amplitude/ magnitude quantities. As a result, block based transforms face a specific dilemma. Control of phase coherency in time domain is highly neglected since the compression effort is only controlled by coefficient quantization in frequency domain.

In addition, several useful features are incorporated into the compression concept:

- Compression quality is optimized locally by using an integral error control loop at pixel level.
- Spatial error diffusion and spatial error density control with the help of a threshold dither matrix.
- The class attribute (3 types of classes) can be easily extracted from the compressed stream enabling independent processing of each class (without interfering with characteristics of another class that is considered incompatible). For example, contours (isophotes) can be identified on a pixel-by-pixel basis without having to re-estimate their location. This could also simplify edge/contour processing tasks to be carried out after de-compression.

We are quite confident about having constructed a novel, conceptual foundation enabling the algorithm's extraordinary perceptual (color) stability: (1) we ***classify*** structure quantities in normalized gradient domain into three major categories, (a) trajectories (contours), (b) homogeneous regions and (c) local extremes; (2) we ***prioritize*** the major categories in the order given above, (a) trajectories, (b) regions and (c) extremes); for example, if image quality needs to be degraded, we still predominantly protect trajectories; (3) for each category we assign a ***weight of global contribution*** to image quality; (4) we process structure quantities concurrently across all color components which

maximizes perceptual color stability; (5) we **track integral error quantities** separately for each category, and (6) with the help of **a threshold dither matrix** we apply ***spatial error diffusion*** and ***spatial error density control***.

Interestingly, with the above constraints the compression algorithm is able to determine a maximally feasible compression ratio with an associated ***estimated*** image quality that appears to be largely independent of image complexity. At the same time the achievable upper limit of compression appears to be astonishingly well correlated with a constant overall ***perceptual*** image quality just moderately influenced by the underlying mathematical image complexity!

Comparison of visual fidelity with luminance based image quality metrics

Now we can also estimate the image quality as a function of compression ratio since the local error control in combination with spatial error density control provides extraordinary stability of image quality across the full range of selectable compression ratios.

Figure 6(a-h) shows a series of image patches taken from 2 reference images of the Kodak image data base **(Figure 5**). Each image patch triplet is the result of compression at 3 different compression ratios.

Figure 5. Two representative high quality test images of the Kodak image data base from which we extracted a set of test patches to carry out compression performance analysis with regard to perceptual color quality analysis as well as detailed statistical error analysis based on several full-reference metrics.

The compression algorithm processes the image patches in sequential raster scan, column after column, from left to right and top to bottom within each column. We added a mirrored image patch on the right to visualize possible side effects emerging from image quality control or sequential raster scan processing.

To facilitate simplified visual comparison among a representative set of image patches together with their cumulative error distribution functions, we compiled the associated figures at the end of the paper (Figures 6a-6h). The representative set of image patches is motivated by the following criteria (at high compression ratios):

VillageK_PatchA: absence of visually noticeable artifacts
VillageK_PatchB: strong loss of structure (roof tiles)
VillageK_PatchC: differential non-linearities in CDFs
VillageK_PatchD: no loss of visually noticeable structure
LadyK_PatchA: Visually noticeable degradation of texture
LadyK_PatchB: slight loss of structure, varying contrast
LadyK_PatchC: mixed content – skintone & eye regions
LadyK_PatchD: skintone exposed to 'contouring'

In search for a suitable metric that could perhaps confirm our surprising result of extraordinary perceptual color stability, we first tried to analyze and visualize significant similarities, differences, and shortcomings between the SSIM and the PSNR metric. Most importantly, we should not neglect that both metrics have been developed by focusing on monochromatic context only. In other words, they do not explicitly describe color specific sensitivity of the human visual system.

So, why should we even consider a detailed performance analysis between PSNR and SSIM? We argue that an enhanced understanding of any metric is best achieved through parametric functional comparison with meaningful means of multi-dimensional visualization in context of differential quantities. In a second step, we can perhaps more easily identify and compare highly suitable metrics that account for color specific sensitivity of the human visual system. In summary, the analysis of suitable image patches in this subsection leads to a valuable comparison with simple CIE ΔE metrics which we discuss in the succeeding subsection dedicated to color aware metrics.

Our basic strategy was the following: we (1) created a set of image patches composed of objects with high relevance to the perceptual sensitivity of the human visual system, (2) analyzed the performance of our image compression algorithm upon subjective quantities derived from our own expert knowledge and (3) tried to correlate the result with the metric under evaluation. We were specifically interested in revealing systematic strengths as well as systematic errors and their root causes of our compression method. Visually analyzing the associated performance as a function of compression ratio seemed to be a promising approach – especially when focusing on variability in differential signal space.

Variability in differential signal space can be comfortably revealed and visualized by a cumulative distribution function (CDF), for example. Let us quickly recall that the CDF is bound to a **monotonic** signal dynamic that always starts at 0=f(0) and always ends at 1=f(1) when using normalized input quantities. In other words, the higher the (error) density and the smaller its mean (error), the later the probability converges toward 1. This property lets the CDF emerge as a great candidate for accumulative error signal analysis in 2-dimensional graphical representations. Subtle changes in nonlinear error variability manifest themselves as highly visible (local) slope changes. The higher the local (error) density, the steeper the local slope. Such local regions are of significant interest when searching for the root cause, especially in context of functional parametric interdependence, for example, with regard to our (perceptual) error space by varying the compression ratio.

We can also easily identify the source of the errors within the CDF if we keep track of each error element by adding an attribute that describes its pixel location. However, this feature has been left to the reader's imagination since it proves only really powerful, for example, in an environment aimed at truly interactive data exploration.

Figure 6 demonstrates the monotonic progression of image quality degradation with the associated cumulative distribution function (CDF) of PSNR quantities as well as SSIM quantities applied to a representative set of image patches.

CDF of PSNR

The CDF of PSNR exhibits an S-shaped curve with variable maximum slope when similar image content is evenly distributed

across the image patch. This appears to be a general characteristic and can be well observed across all image patches described in our current paper. We would like to point out that our compression algorithm already creates coding errors at a compression ratio of 1.0. However, since it performs indisputably visually lossless on practically any natural image content (@ compression ratio 1.0) so far, we consider the underlying shape of the CDF as an 'ideal' or 'optimal' reference shape. Interestingly, the compression performance spans across approximately 13dB between this ideal reference mean PSNR value and the visually lossless compression performance with compression ratios in the range of 2.6 to 2.7.

Since the CDF of PSNR shifts left as a function of decreasing mean PSNR, we also defined an arbitrary maximum upper mean PSNR bound equal to 60dB that corresponds to 1=F(1). Consequently, any PSNR value greater than 60dB will be clipped in the graphical representation.

The legend shows the mean PSNR value together with the achieved compression ratio in parenthesis. The spacing between each plotted curve is driven by the target compression ratio, evenly incremented by 0.4 in the range of 1.0 to 5.0, therefore, generating 11 independent curves. The inset illustrates the spatial occurrence of the error quantities at highest compression ratio. The darker the pixel in the inset image is, the stronger the associated error value. Now, we show evidence that the 2-dimensional curves provide valuable detailed insight into nonlinear differential variability which cannot be explicitly expressed by the single scalar mean PSNR value.

At the same time, we can compare the variability of these curves with the visual observations we can draw from the result images shown at three different compression ratios (Fig. 6). The 'ideal' compression performance at compression ratio 1.0 is illustrated in the top row. The visually lossless compression performance (at expert level) is shown in the middle row and the maximally achievable compression ratio is shown in the bottom row. The target compression ratio equals $1/10^{th}$ of the number following the keyword 'obj' and the achieved compression ratio equals $1/10^{th}$ of the number following keyword 'real.' The result image of the maximally achievable compression ratio is not only intended to illustrate the emerging visual artifacts but to also illustrate the *extraordinary perceptual color stability* of the current version of the compression method.

CDF of SSIM

We also propose to compare our subjectively motivated score of compression performance with the largely accepted full reference metric SSIM [7]. In comparison to PSNR, SSIM is already well normalized to 1 which transforms to location 1=f(1) in CDF space (see fig. 6). However, for the purpose of visualizing variability with increased detail, we preferred using a logarithmic scale along the CDF's ordinate axis. This also conveniently renders the CDF of SSIM very often as an approximately linear slope that varies its slope coefficient as a function of compression ratio. As with the graph of CDFs of PSNR, the inset image depicts the SSIM values in spatial context at maximum compression ratio.

Interestingly, the CDF of SSIM seems to change curvature across mid-range errors - from concave (exponent > 1) to convex (exponent < 1) - when we increase our visually lossless compression ratio beyond 2.7. So, can this curvature phenomenon therefore be highly associated to error values at visual threshold?

The image patch 'VillageK_PatchA' at a compression ratio of 4.5 demonstrates a special scenario in which SSIM clearly indicates that visually lossless performance has already been compromised while visually noticeable artifacts remained absent. Can we perhaps conclude that SSIM as a full-reference metric clearly indicates loss of structure compared to the reference image, but it cannot easily distinguish between visually noticeable artifacts and generic loss of structure quantities when visually analyzed in no-reference context?

Color aware image quality metrics

We can confidently state that perceptual color stability doesn't appear to be compromised in any of the carefully selected image patches in spite of a plurality of emerging visual artifacts at maximally achievable compression ratios. Highly stupefied by this phenomenon, we now would like to analyze the phenomenon in detail by using the rather simple, color centric metric CIE ΔE76 (CIEDE76). We remain mostly interested in revealing meaningful relative quantities in differential context (intra-CDF), instead of trying to argue about absolute quantities (inter-CDF) across comparable metrics, especially when considering the still questionable performance of SSIM with respect to absolute quantities (linear geometric mean) which we demonstrated by our set of color test patches. Since we are currently not interested in (re-)using absolute quantities of perceptual full-reference metrics, we also disregard the CIE ΔE2000 metric [8].

CDF of CIEDE76

The fundamental separation into luminance, chrominance, and hue domain already present in the ΔE76 metric offers valuable insight on relative perceptual sensitivity between the three domains. Considering the still questionable performance of SSIM with respect to absolute quantities (linear geometric mean) as, for example, revealed by our exemplary set of color test patches, we can also much more easily discuss the general performance characteristics of the ΔE76 components.

First of all, the CDF error profiles of our compression algorithm exhibit highly similar differential characteristics in luminance domain when we compare between SSIM and CIEDE76. Therefore, we adjusted the linear gain of error values so that they match the same dynamic range for simplified visualization in log-linear signal space. **Figure 7** shows the variability of CIEDE76 CDF in **luminance** domain across all 8 image patches; each image patch is represented by a pair of CDFs - with a unique color assigned - illustrating the minimum compression ratio (lower bound) and the maximum compression ratio (upper bound). Although the graph appears highly cluttered, it confirms the basic observations we have already drawn from the CDF of SSIM. Across all image patches we find concave to linear curvatures at low compression ratios and linear to convex curvature at high compression ratios.

Rather expectedly, but still surprisingly, the variability of CIEDE76 CDF in **chrominance** domain and in **hue** domain shows mostly concave to linear curvatures and steeper linear slopes than in **luminance** domain across all 8 image patches.

The image patch *VillageK_PatchB* illustrating roof tiles with loss of structure at maximum compression is a highly suitable candidate to compare the CDFs of the metrics in luminance, chrominance, and hue domain. We expect a stronger differential non-linear progression towards convex curvature in luminance

domain versus chrominance and hue domain – especially across the roof tiles where luminance structure has been significantly lost. Astonishingly, the expected behavior is easily confirmed across the associated CDF in luminance domain at maximum compression ratio.

We also propose to discuss variability found in luminance, chrominance, and hue domain while comparing the image patch *VillageK_PatchA* with the image patch *VillageK_PatchB* at minimum and maximum compression ratios. While luminance changes curvature to convex, chrominance and hue do not. In addition, all chrominance and hue slopes exhibit a significantly steeper slope than the corresponding luminance slopes. We therefore correlate the underlying phenomenon with significantly higher perceptual color stability.

Conclusion

The novel, nonlinear method applied to image compression enables significantly better local adaptation to non-stationary image information in context of human visual perception while accomplishing low computational complexity and preserving finest levels of discrete structure (orientation) and magnitude information. Error quantities are being processed in several complementary differential domains enabling advanced visual masking strategies. We also maintain full spectral bandwidth fidelity and optimal phase coherency by having eliminated any linear filtering or median filtering which was traditionally used to improve prediction and coding performance. Needless to say that less filtering also favorably translates into a fewer number of operations per pixel. In addition, the new possibility of processing structure quantities and gradient magnitude quantities concurrently but independently enabled adjusting the underlying sparse representation of image information as a function of compression ratio guided by a novel error density control based on a threshold matrix concept. Surprisingly, the associated image quality is highly stable and relatively proportional to compression ratio. In addition, the compression algorithm can be configured to respect a lower bound of overall perceptual image quality. As a result, 8-bit gamma-pre-corrected high quality still images reach compression ratios between 3.9 and 4.8. First estimates, which however still need to be confirmed, predict that compression ratios of about 5x appear very reasonable for 10-bit video streams – without creating any visually noticeable artifacts. We presume that the substantially high computational efficiency has been achieved by focusing on discrete orientation (encoded with just 2 bits per pixel) as well as normalized minimum gradient magnitude that are easily derived from discrete pixel amplitudes already present on the underlying grid of the pixel matrix. Such a discrete pair of orientation quantities and minimum gradient magnitude quantities offers an excellent solution for representing a powerfully adaptive coefficient in local context. This can be efficiently quantized and seems to well imitate astonishingly well human visual local contrast perception.

A detailed analysis of error quantities across carefully selected image patches enabled us to underline the extraordinary perceptual color stability of the current image compression algorithm. Comparing three commonly known full-reference metrics, namely YPSNR, SSIM, and CIEDE76, in differential context of cumulative distribution functions proved extremely useful.

We also imagine that the presented concept can be advantageously applied to multi-dimensional and multi-scale data sets of many other challenging engineering tasks achieving efficient computational performance and dedicated precision – especially where structural quantities represent important information in local context.

Acknowledgment

The intuitive desire of steadily pursuing the underlying key ideas enabling the discovery of most select discrete multi-dimensional representation of image information in visually meaningful quantities, in quest of beneficially reaching beyond linear system theory, had been a highly challenging endeavor, not only from the engineering point of view, but also across project management and project financing matters. Replicating the importance of steadiness and perseverance, a huge thank you goes, once again, to the Electronic Imaging conference community members for their continuous mutual support spanning over more than 25 years. The extremely valuable interdisciplinary exchange of unparalleled scientific and technical knowledge has proven to be the very best platform to create – and follow up on – exaggerated ideas. May profound gratitude especially reach Bernice Rogowitz, John McCann, Reiner Eschbach, Gabriel Marcu, Alessandro Rizzi, Jan Allebach, Sheila Hemami, Damon Chandler, Mylene Farias, Michael Kriss, Peter Burns, Scott Daly, Al Ahumada, Beau Watson, and Sergio Goma for their fascinating personal interactive engagement as well as wonderfully stimulating and encouraging scientific support over so many years.

Randolph Fox, a ‘retired’ colleague at STMicroelectronics, provided magnificent handwritten tutorials on understanding tensor algebra in engineering applications already several years ago which enabled fascinating inspiration from structure tensor analysis. During frequent technical discussions - driven by his comprehensive personal interest in accompanying the project – he stimulated and shared many more powerful ideas based on his wonderful engineering experience in microelectronics. Mariano Bona’s excellent experience in discrete mathematics and control system theory together with his limitless curiosity in applying his wonderful knowledge to the domain of digital image processing let the latest discrete solution and its extraordinary perceptual color stability crystallize. A profound thank you shall acknowledge both of them.

Last but not least, the project would not have been feasible without substantial funding from within the NANO2017 framework. Therefore, we would like to also express ultimate gratefulness to Mario Diaz-Nava for having offered his extraordinary experience in project funding as well as having gained great confidence in the potential project outcome.

References

[1] Fritz Lebowsky, Mariano Bona, "How suitable is structure tensor analysis for real-time color image compression in context of high quality display devices," in Proc. IS&T - Electronic Imaging, Color Imaging XXI: Displaying, Processing, Hardcopy, and Applications, pp. 1-8(8), 2016.

[2] Thomas Brox, J. Weickert, B. Burgeth, P. Mrázek, "Nonlinear structure tensors," Image and Vision Computing, vol. 24, no. 1, pp. 41-55, 2006.

[3] Fritz Lebowsky, "Optimizing color fidelity for display devices using contour phase predictive coding for text, graphics, and video content," in Proc. SPIE 8652, Color Imaging XVIII: Displaying, Processing, Hardcopy, and Applications, pp. 86520X, 2013.

[4] Marina Nicolas, Fritz Lebowsky, "Optimizing color fidelity for display devices using vectorized interpolation steered locally by perceptual error quantities," in Proc. SPIE 9395, Color Imaging XX: Displaying, Processing, Hardcopy, and Applications, pp. 939502, 2015.

[5] Marina Nicolas, Fritz Lebowsky, "Preserving color fidelity in real-time color image compression using a ranking naturalness criterion," in Proc. IS&T - Electronic Imaging, Color Imaging XXI: Displaying, Processing, Hardcopy, and Applications, pp. 1-7(7), 2016.

[6] Hao Pan, Scott Daly, M. Ibrahim Sezan,"Analysis of the sigma filter using robust estimation," in Proc. SPIE 5672, Image Processing: Algorithms and Systems IV, 2005.

[7] Z. Wang, A. C. Bovik, H. R. Sheikh and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," IEEE Transactions on Image Processing, vol. 13, no. 4, pp. 600-612, 2004.

[8] G. Sharma, W. Wu, E. N. Dalal, "The CIEDE2000 Color-Difference Formula: Implementation Notes, Supplementary Test Data, and Mathematical Observations," Color Research and Application, vol. 30. No. 1, 2005.

Author Biographies

Fritz Lebowsky received his MS (1985) and PhD (1993) in electrical engineering from the Technical University of Braunschweig, Germany. He began his professional career as a research and teaching assistant at the Institute of Telecommunications of the Technical University of Braunschweig in 1985. From 1991 he worked as a research and development engineer in the field of digital video processing at Micronas in Freiburg, Germany. In 1995 he joined Thomson Consumer Electronics Components in Meylan, France, as a development engineer modeling video processor networks as well as digital acquisition sub-systems for DVD ROM drives. In 2000 he joined the Imaging and Display Division of STMicroelectronics Inc. in San Jose, CA, developing advanced display engines for the PC flat panel monitor market. Since 2004 he has been with STMicroelectronics in Grenoble, France, working on image quality improvement for consumer products as well as transferring the attained knowledge to other domains of microelectronic products, for example, automotive where quality and reliability at reasonable cost are essential.

Mariano Bona received his PhD (1983) in applied mathematics from the Scientific University of Grenoble, France. Since 1985 he has been with STMicroelectronics, having primarily worked on front-end signal processing and speech compression. In 2015 he joined the research and development team focusing on innovative nonlinear image compression algorithms.

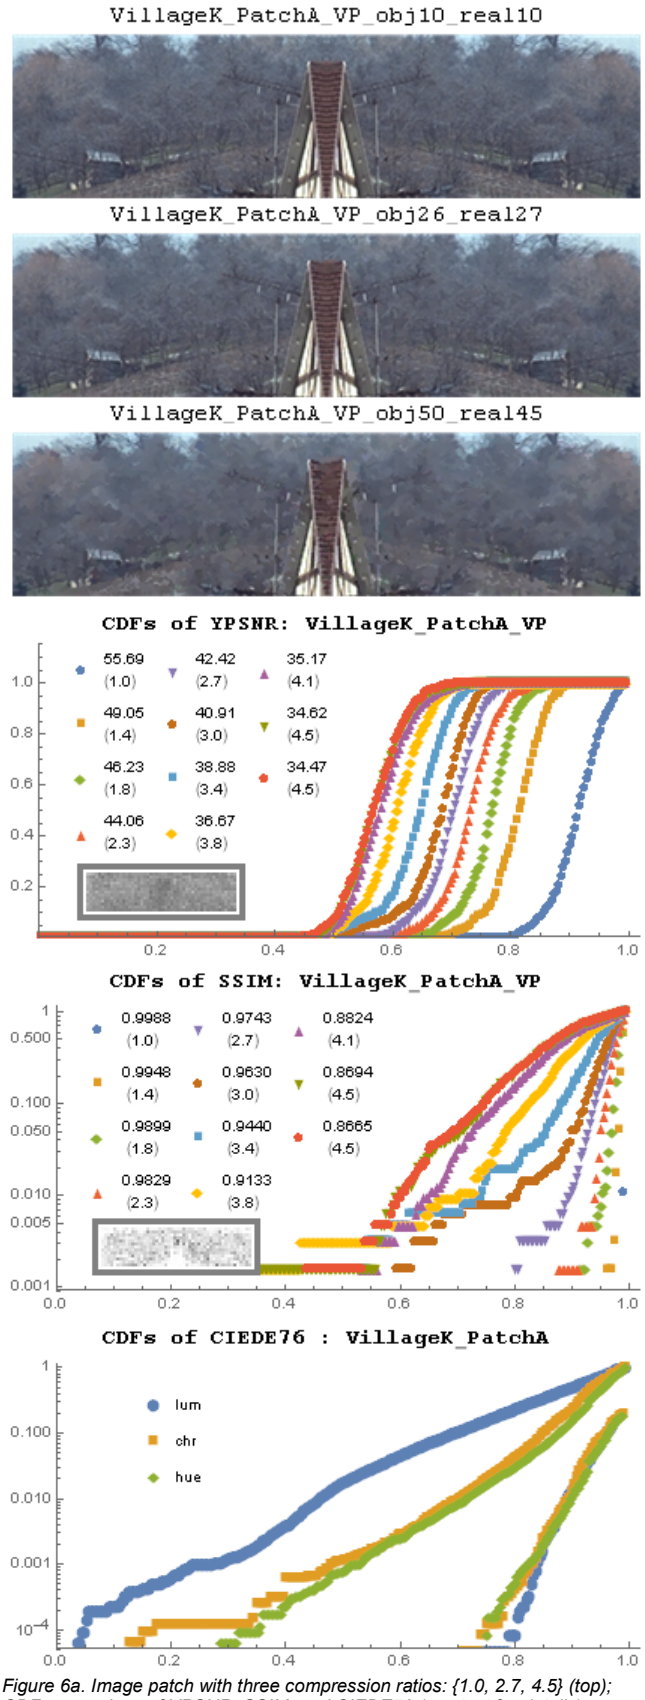

Figure 6a. Image patch with three compression ratios: {1.0, 2.7, 4.5} (top); CDF comparison of YPSNR, SSIM, and CIEDE76 (see text for details)

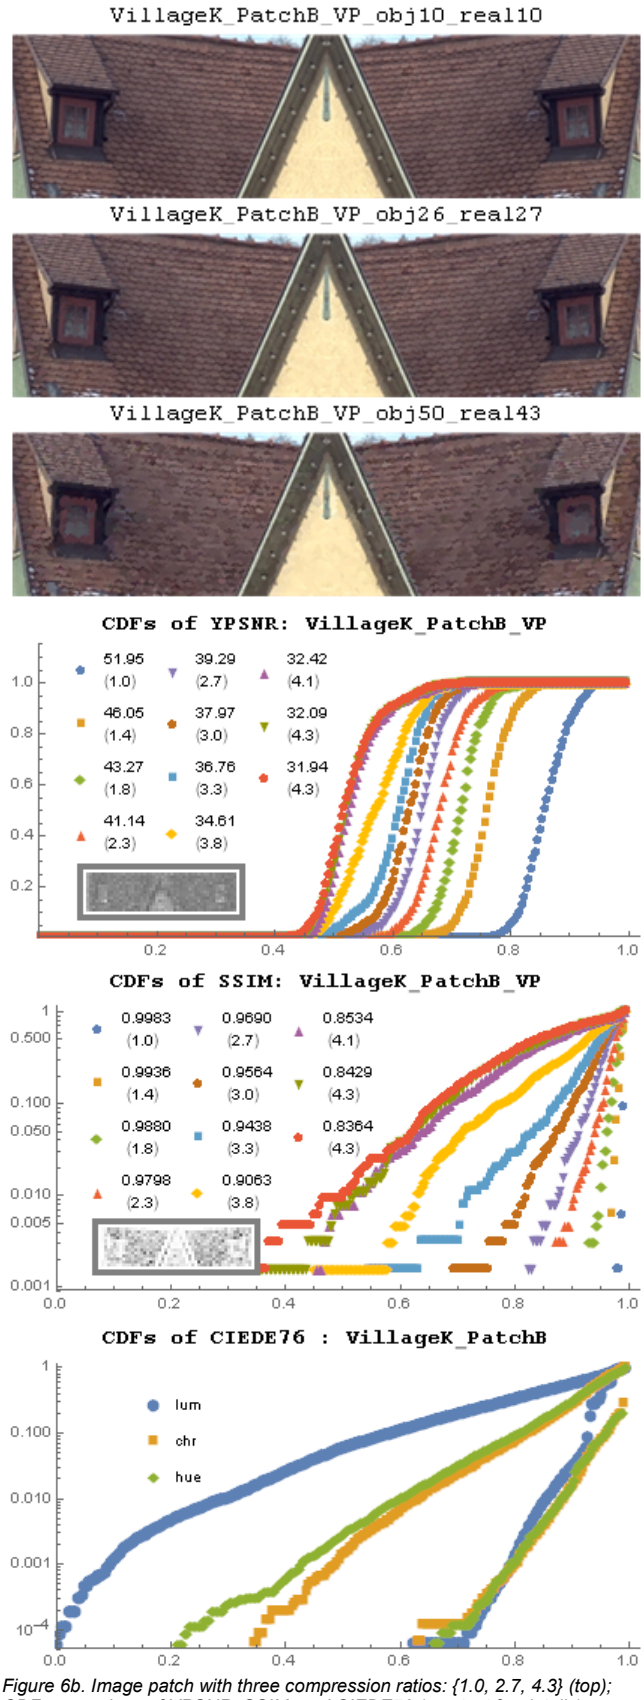

Figure 6b. Image patch with three compression ratios: {1.0, 2.7, 4.3} (top); CDF comparison of YPSNR, SSIM, and CIEDE76 (see text for details)

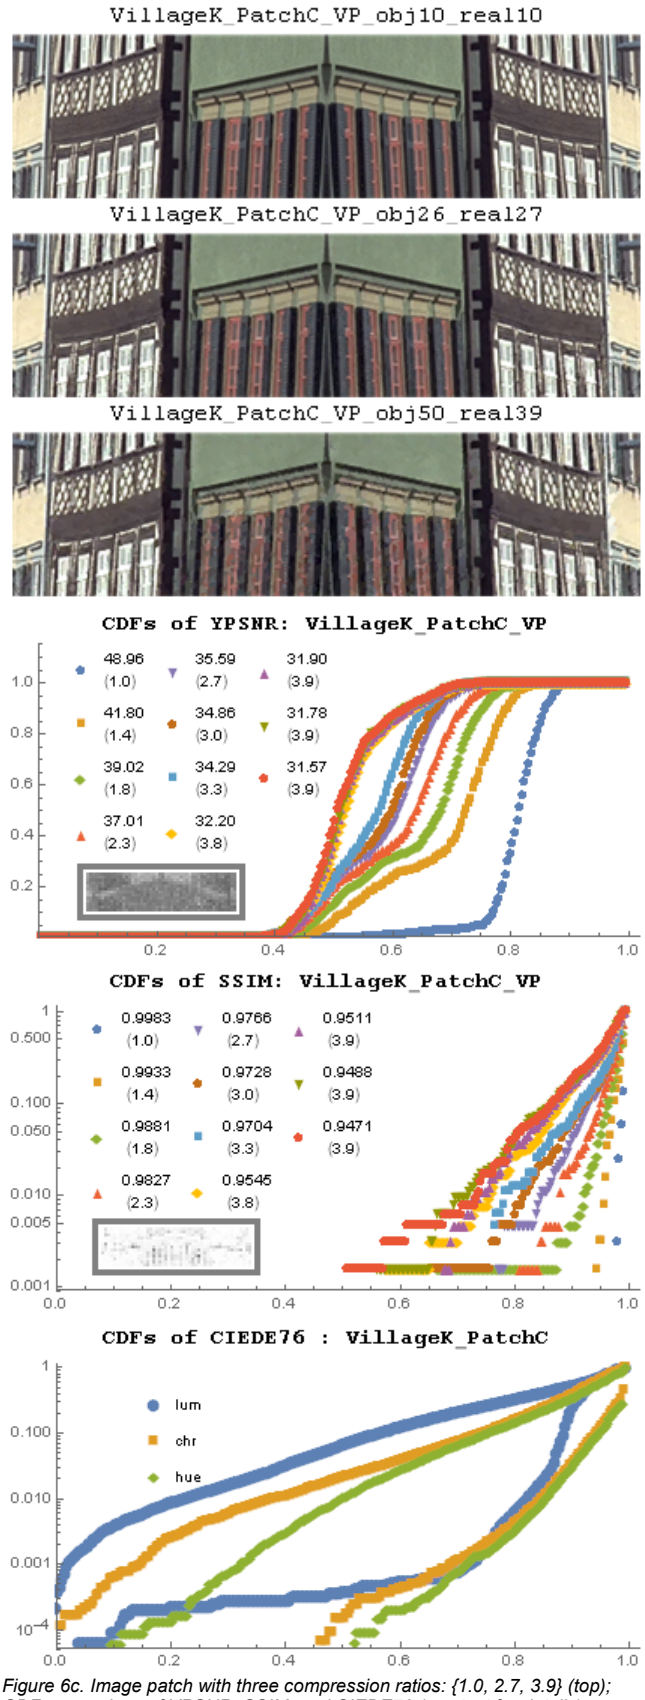

Figure 6c. Image patch with three compression ratios: {1.0, 2.7, 3.9} (top); CDF comparison of YPSNR, SSIM, and CIEDE76 (see text for details)

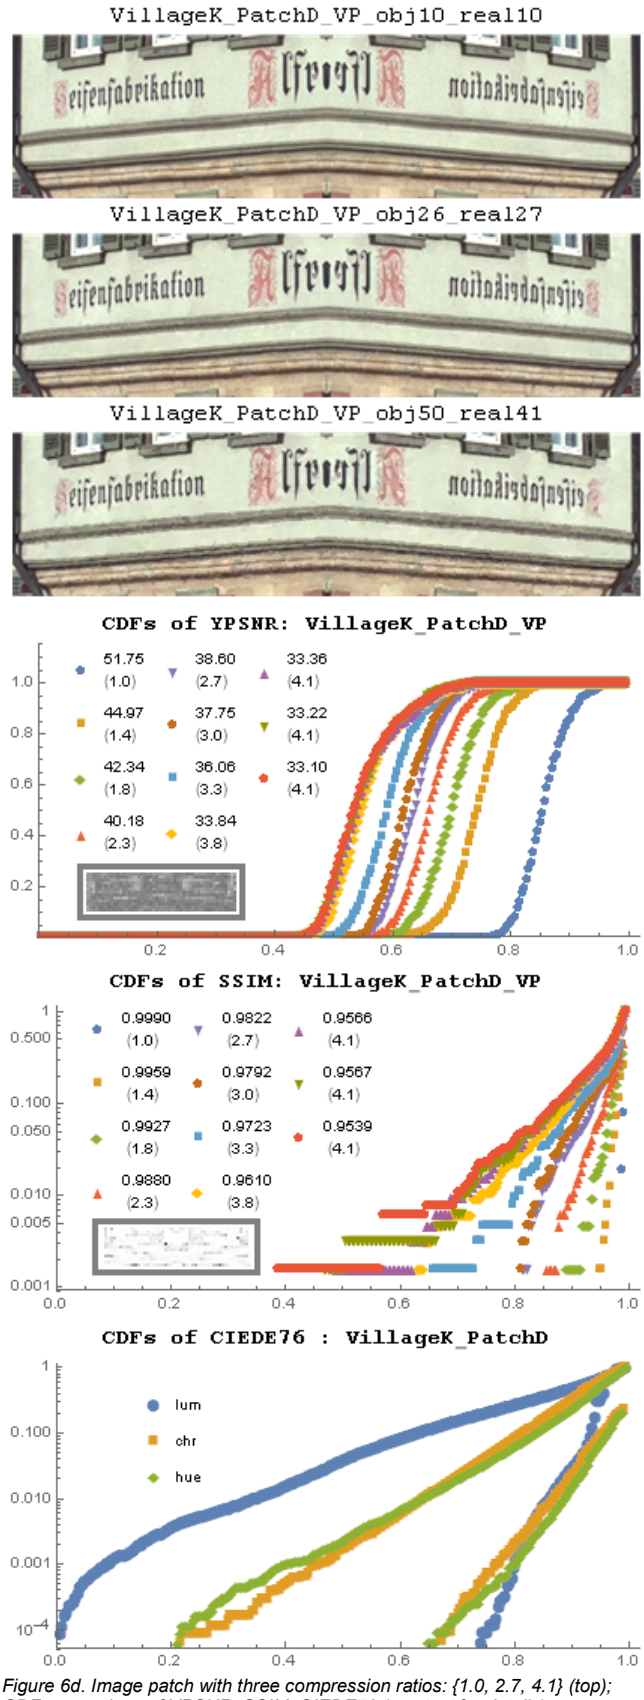

Figure 6d. Image patch with three compression ratios: {1.0, 2.7, 4.1} (top); CDF comparison of YPSNR, SSIM, CIEDE76 (see text for details)

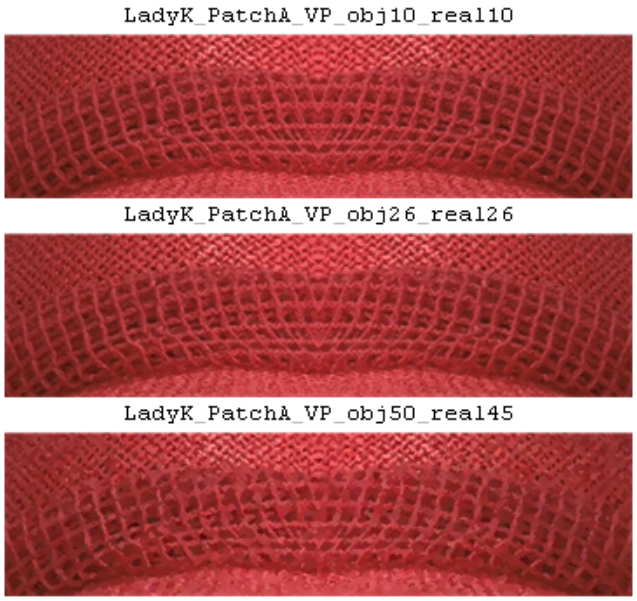

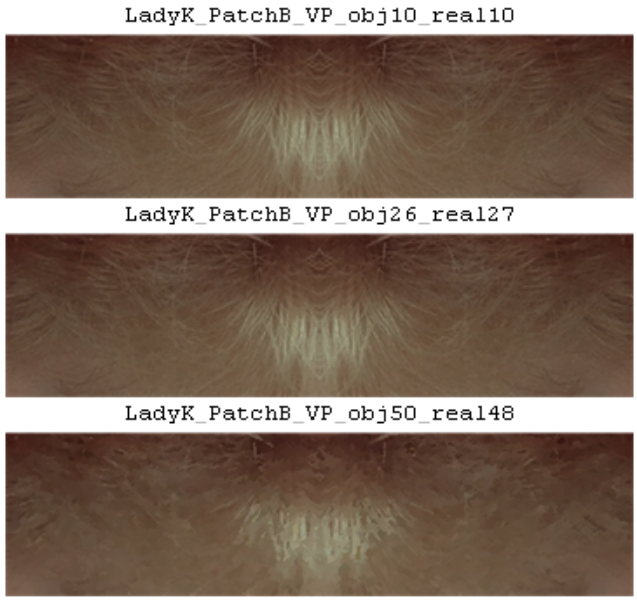

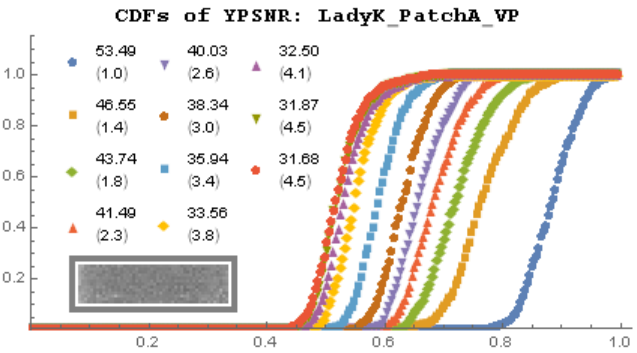

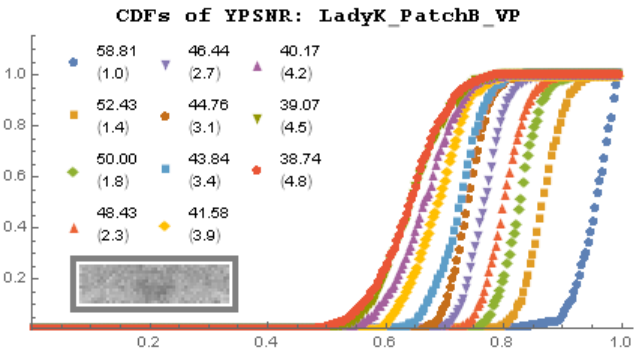

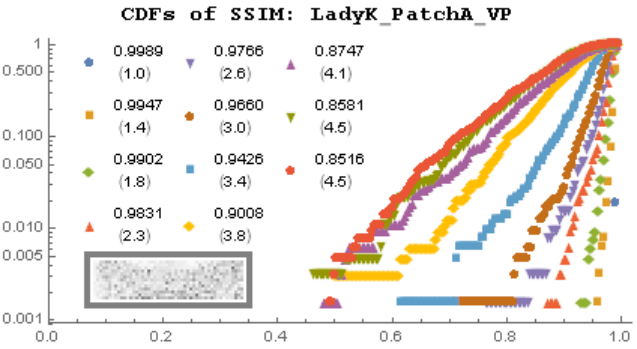

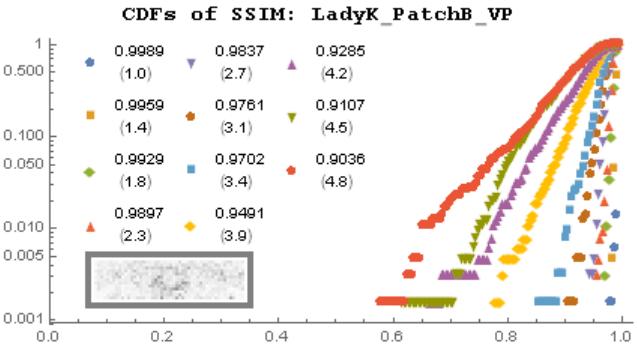

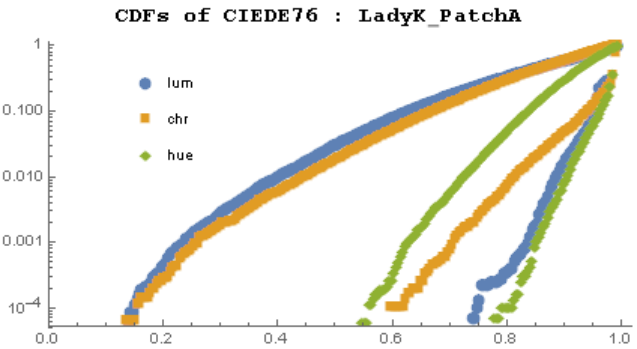

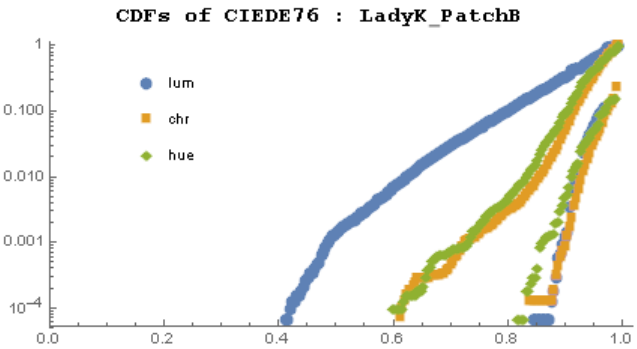

Figure 6e. Image patch with three compression ratios: {1.0, 2.6, 4.5} (top); CDF comparison of YPSNR, SSIM, and CIEDE76 (see text for details)

Figure 6f. Image patch with three compression ratios: {1.0, 2.7, 4.8} (top); CDF comparison of YPSNR, SSIM, and CIEDE76 (see text for details)

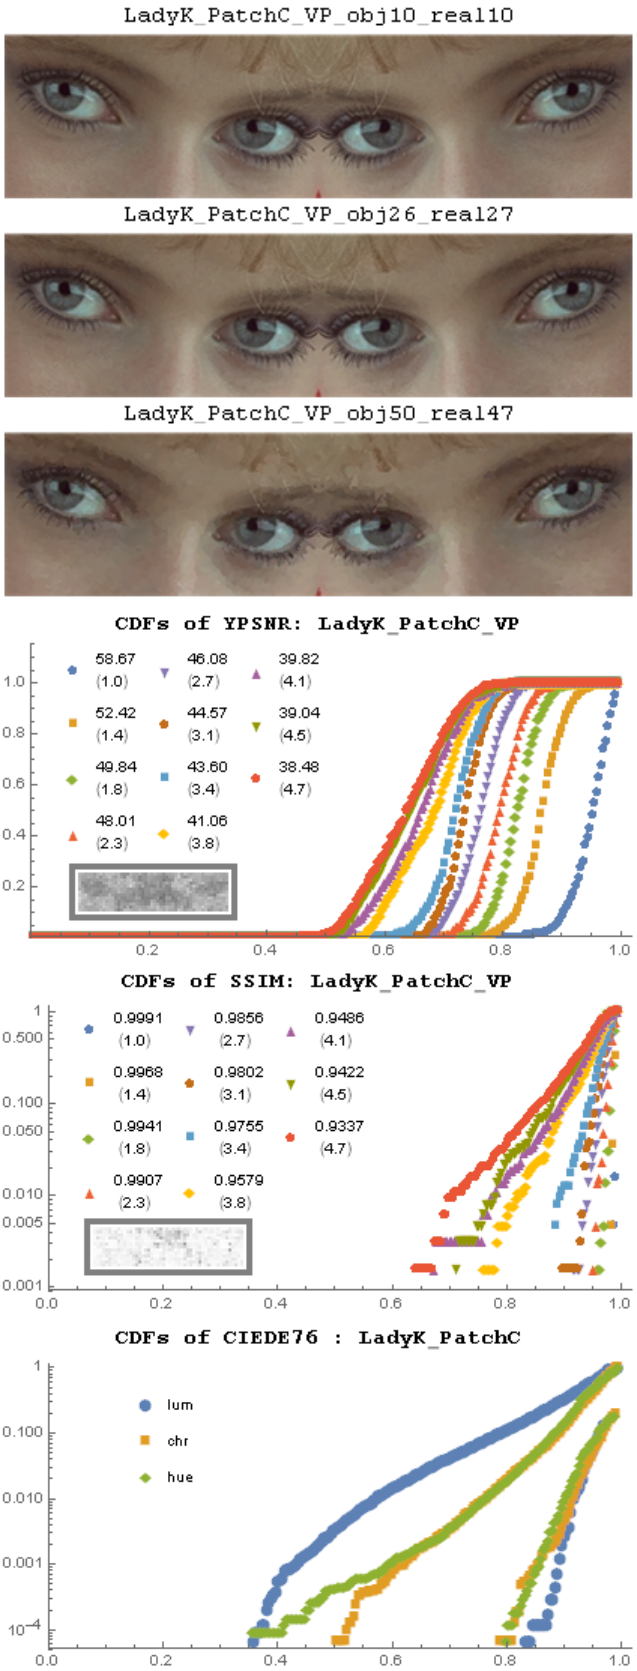

Figure 6g. Image patch with three compression ratios: {1.0, 2.7, 4.7} (top); CDF comparison between YPSNR, SSIM, and CIEDE76 (see text for details)

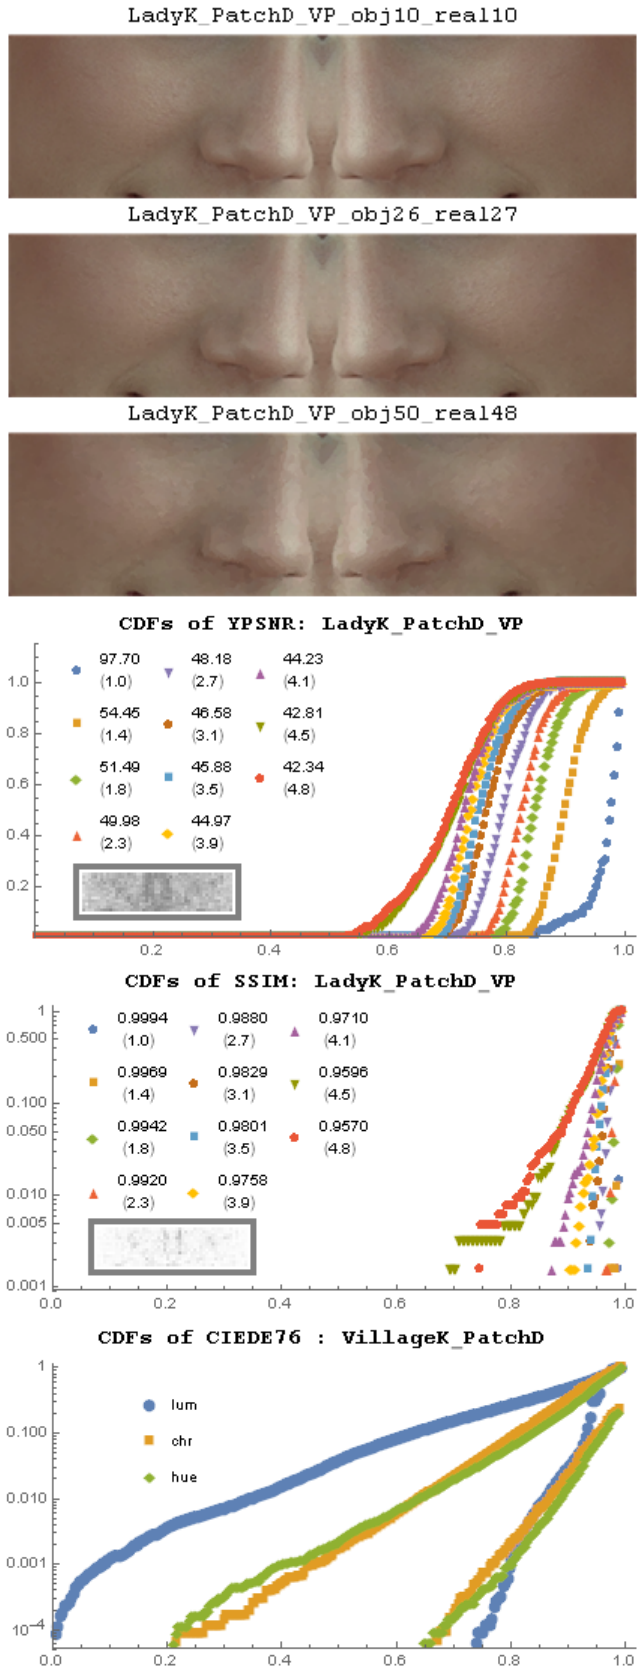

Figure 6h. Image patch with three compression ratios: {1.0, 2.7, 4.8} (top); CDF comparison of YPSNR, SSIM, and CIEDE76 (see text for details)

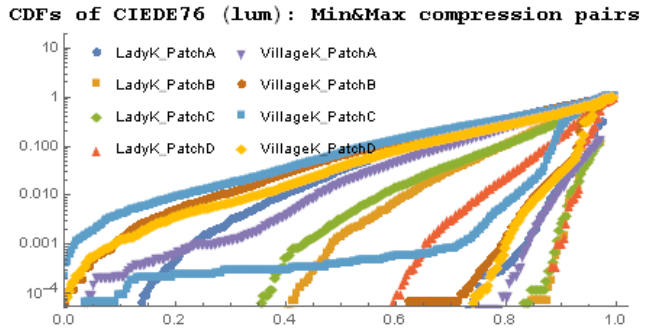

*Figure 7a. Variability of CIEDE76 CDF in **luminance** domain across image patches; each image patch is represented by a pair of CDFs - with a unique color assigned - illustrating the minimum compression ratio (lower bound) and the maximum compression ratio (upper bound)*

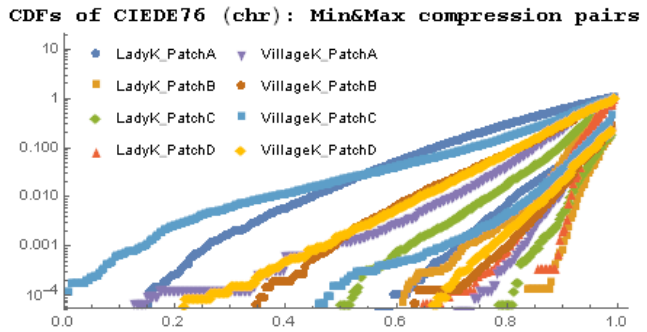

*Figure 7b. Variability of CIEDE76 CDF in **chrominance** domain across image patches; each image patch is represented by a pair of CDFs - with a unique color assigned - illustrating the minimum compression ratio (lower bound) and the maximum compression ratio (upper bound)*

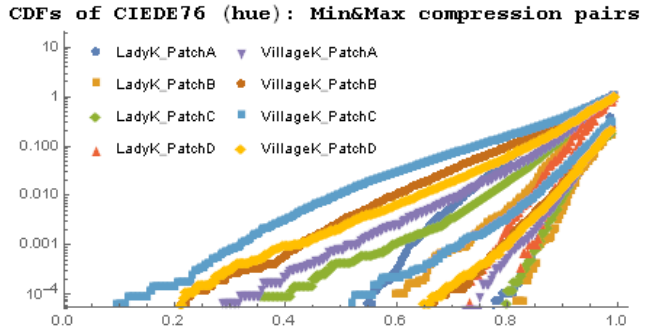

*Figure 7c. Variability of CIEDE76 CDF in **hue** domain across image patches; each image patch is represented by a pair of CDFs - with unique color assigned - illustrating the minimum compression ratio (lower bound) and the maximum compression ratio (upper bound)*