

Gamut Mapping in Cinematography through Perceptually-based Contrast Modification

Syed Waqas Zamir, Javier Vazquez-Corral, Marcelo Bertalmío

Abstract—Gamut mapping transforms the colors of an input image to the colors of a target device so as to exploit the full potential of the rendering device in terms of color rendition. In this paper we present spatial gamut mapping algorithms that rely on a perceptually-based variational framework. Our algorithms adapt a well-known image energy functional whose minimization leads to image enhancement and contrast modification. We show how by varying the importance of the contrast term in the image functional we are able to perform gamut reduction and gamut extension. We propose an iterative scheme that allows our algorithms to successfully map the colors from the gamut of the original image to a given destination gamut while keeping the perceived colors close to the original image. Both subjective and objective evaluations validate the promising results achieved via our proposed algorithms.

Index Terms—Gamut Mapping (GM), Gamut Mapping Algorithm (GMA), Gamut Reduction (GR), Gamut Extension (GE), color contrast, variational methods

I. INTRODUCTION

IN the film industry, an important problem at the post production stage is to reproduce colors for different display devices such that a person watching a movie on any of these devices perceives the same colors [9]. To reach this objective, colorists, by using 3D Look Up Tables (LUTs), map the colors of the movie material to the colors of each type of target device (digital cinema projector, HD TV, etc.). However, there are millions of entries in each LUT and colorists only specify a few colors manually, and the rest of the colors are interpolated without taking care of their spatial and temporal context [8]. Subsequently, the resulting video may have false colors that were not present in the original material and intensive manual correction by a skilled colorist is usually necessary. An automated process called Gamut Mapping (GM) is therefore needed: GM transforms colors from an input to an output device gamut, whereas the gamut of a device is the set of colors that this device is able to reproduce. Intensive research has been carried out in the GM area; however, it is still an open field due to the difficulty of the challenges involved. One of the major tasks is to retain the perceptual color appearance of the original image into the reproduced image and most of the algorithms in the literature lack in this regard.

There are two major types of GM: Gamut Reduction (GR) and Gamut Extension (GE). Gamut reduction involves the

mapping of colors from a larger source gamut to a smaller destination gamut. For example, cinema footage needs to be passed through a gamut reduction method in order to be displayed on a television screen [8], [23]. On the other hand, gamut extension refers to the transformation of colors from a smaller source gamut to a larger destination gamut. Gamut extension is needed by state of the art digital cinema projection systems, which often receive a cinema signal that is encoded with a limited gamut as precaution against commonly poor projectors; therefore, in order to realize the full potential of these projectors in terms of colors, a gamut extension procedure is required [9].

The contribution of this paper is twofold. First, we propose a perceptually inspired gamut reduction algorithm where gamut reduction is achieved through contrast reduction, by adapting a framework [11], [10] that is inspired by the properties of contrast perception in the human visual system and closely related to the Retinex theory of color vision [28]; an earlier, reduced version of this first contribution was presented as a conference paper in [46]. In that conference paper, experiments were performed only on 17 images, none of them coming from videos. The evaluation on videos has been performed exclusively for this publication. Second, we adapt the aforementioned framework [11], [10] to develop an algorithm for gamut extension using a contrast enhancement approach. Our algorithms outperform state of the art techniques [5], [29] both subjectively and according to the recently presented perceptually-based metric of [30].

This paper is organized as follows; first, the related work is summarised in section II. The image energy functional presented in [11] for perceptually inspired contrast enhancement is described in section III. Then, in section IV, the contrast enhancement model is adapted to obtain GR and GE algorithms. In section V, experiments and results are discussed. Finally, the paper is concluded in section VI.

II. RELATED WORK

A plethora of Gamut Mapping Algorithms (GMAs) exists in the literature and the interested reader is referred to the excellent book by Morovič [35]. The GMAs that are dedicated to perform gamut reduction, to which we refer as Gamut Reduction Algorithms (GRAs), are classified into two broad categories. The first category consists of global (also called non-local or non-adaptive) GRAs [15], [22], [38], [45] that involve point-to-point mapping of colors (usually through a

Authors are with the Department of Information and Communication Technologies, Universitat Pompeu Fabra, 08018, Barcelona, Spain. e-mail: {waqas.zamir, javier.vazquez, marcelo.bertalmio}@upf.edu

predefined lookup table) from source to destination gamut. The standard non-local GRA, Hue Preserving Minimum ΔE (HPMINDE), was proposed by Murch and Taylor [38] where, in order to reproduce the image, the out-of-gamut colors are clipped to closest points on the target gamut boundary along the lines of hue. Non-adaptive GRAs involve either clipping or compression and completely ignore the spatial color configuration in the source image. In contrast, the second category involves the GRAs that take into account the spatial color information of the original image while fitting the color gamut of an image into the gamut of a given device [4], [5], [7], [25], [34], [36], [47]. However, these algorithms are often computationally expensive, or based on many assumptions, and may report halo artifacts. McCann [32], [33] proposed a Retinex-inspired framework that performs spatial comparisons to reproduce the image while preserving the local gradients at all scales as in the original image. A similar multi-resolution GR approach that adapts to the original image content is suggested by Farup et al. [16]. A cluster based approach is defined in [29] for the optimization of GM. Alsam and Farup [5] proposed an iterative GRA that at iteration level zero gives a result identical to gamut clipping, whereas, by increasing the number of iterations, the solution approaches spatial gamut mapping. Unlike the global GRAs, the potential of spatial gamut reduction methods is to preserve the color gradient between two out-of-gamut colors instead of mapping them to the same in-gamut color. Another fundamental motivation behind spatial gamut mapping, in order to emulate the color perception properties of the human visual system, is the possibility to formulate a strategy where two out-of-gamut colors with identical lightness and chromaticity map to two different in-gamut colors depending on their spatial context in the image [25].

The case of gamut extension is different from that of gamut reduction: only a handful of Gamut Extension Algorithms (GEAs) exist in the literature. One could think of simply taking any GRA and use the one-to-one mapping in the reverse direction to perform gamut extension, as Morovič comments in [35]. Hoshino [19] proposed the first GEA to map printed images to the color gamut of high-definition TV using lightness and chroma mapping. Kang et al. [21] presented a GEA that was based on numerical fitting of subjective experimental data to linearly map lightness and chroma. Anderson et al. [6] presented a user assisted method where an expert was asked to expand the gamut of some key frames and the rest of the frames were corrected accordingly. Lui et al. [31] presented a gamut expansion method in the Luv space that extends from an anchor point while respecting the hue lines. Kim et al. [24] described a GEA with three type of extension strategies: chroma mapping, mapping along lines from the origin (which they called vector mapping) and adaptive mapping that is a compromise between the first two strategies. Casella et al. [13] introduced a sigmoid function into the linear mapping techniques so that it does not consider objects of low chroma equivalent to those of high chroma. Subjective evaluations were performed in Mujis et al. [37], and Laird et al. [26]. In the latter, they proposed and evaluated

the following five GEAs:

- True-color, where they maintain the color information between source and destination media.
- Same Drive Signal (SDS), that is used to reproduce the colors of a small gamut on a wider gamut display using linear extension.
- Hybrid Color Mapping (HCM), that performs GE by using a saturation-dependent linear combination of the true-color and SDS algorithms.
- Chroma, that maps the color of source gamut to the reproduction gamut along the lines of chroma axis while keeping the lightness and hue constant.
- Lightness-chroma adaptive, inspired by the work of [24], which performs GE by altering both lightness and chroma while keeping the hue constant.

Recently, Heckaman et al. [18] suggested an algorithm that exploits the use of memory colors such as the green of the grass or the blue of the sky while preserving the flesh tones. Jung et al. [20] proposed the search of color corrections between source and target gamuts into the 1931 CIE XYZ space.

One of the major problems in GM is the evaluation of the results. Usually, subjective comparisons are performed [12], [14]. The most common subjective method is the pair comparison, where observers are asked to choose which of two different gamut-mapped versions of an image is more faithful to the original. However, subjective measures are time consuming, involve complexities, and do not provide clear cues to improve the given GMA. Recently, a perceptually-based Color Image Difference (CID) metric [30] has been proposed that puts particular emphasis on the assessment of gamut-mapped images. It is based on evaluating the distortions in lightness, hue, chroma, contrast and structure of the gamut-mapped images with respect to the original images.

III. PERCEPTUALLY-BASED COLOR AND CONTRAST ENHANCEMENT

In this section we briefly recap the main concepts of the Retinex theory of Land (arguably the most popular approach for perceptually-based color correction), and its close relationship to the variational method for contrast enhancement which we will adapt for our gamut mapping purposes.

Land makes in [27] a very clear and detailed explanation of his Retinex theory and the experiments that led to its postulation. After scores of perceptual matching tests, his conclusion was that our perception of the color of an object had a physical correlate in what he called *scaled integrated reflectance*, which is defined, for each waveband (long, medium and short, corresponding to cone response sensitivities,) as a ratio: the integral of the radiance of the object over the waveband, divided by the integral over the same waveband of the radiance of a white object under the same scene illuminant. The scaling is a non-linear function that relates reflectance to lightness sensation. But this implies that in order to perceive the color

of an object somehow our visual system is comparing the light coming from the object with the light coming from a reference white, and Land wondered how we are able to find this reference white “*in the unevenly light world without reference sheets of white paper*” [27]. The sensation of white will be generated by the area of maximum radiance in all three bands (this is the von Kries’ model or “white-patch” assumption, although Land doesn’t cite it); this area could be used as reference, but Land didn’t know how our visual system could “*ascertain the reflectance of an area without in effect placing a comparison standard next to the area*” [27]. The solution he proposed consisted of comparing far-away points through paths: the ratio of the values at the two end-points of the path can be computed as the sequential product of the ratios of each point of the path with the following point. The Retinex algorithm consists of assigning, for each point and each waveband (long, middle, short), an estimate reflectance obtained as the average of the sequential products obtained on many paths, all ending in the given point. Land thought that this was a plausible explanation of how our visual system estimates reflectances but he didn’t want to venture where exactly this type of operations were being carried out, in the retina or at the cortex; therefore he chose the name “Retinex” for his approach. The Retinex algorithm is directly applied to digital images in a straightforward way, where the pixels will be the points and the three color channels R, G and B play the role of the wavebands. As proved in [42], Retinex always increases brightness so it can’t be directly applied to overexposed pictures; also, if the algorithm is iterated the results may improve but the convergence image is flat white, so there is some “sweet spot” of the number of iterations yielding the best output [10]. Another major source of problems is Retinex’s reliance on paths: their length, shape and number condition the results and many works have been proposed trying to optimize the selection of these variables.

The Automatic Color Enhancement (ACE) algorithm of Rizzi et al. [43] is also based on perception, and its relationship with Retinex will become clear shortly. ACE is designed to mimic some basic characteristics of the human visual system, like the white patch and the grey world mechanisms, lateral inhibition, the independence of chromatic channels, or the influence of spatial relationships in the scene. The authors perform experiments that show how ACE has several excellent properties: it allows to obtain very good color constancy, it increases the dynamic range of the input and, unlike Retinex, it can deal both with under- and over-exposed pictures, it can perform de-quantization (eliminating quantization artifacts produced by encoding an image with an insufficient number of bits per channel), and it can reproduce some visual perception illusions. Its main limitation is its computational complexity, $O(N^2)$ where N is the number of pixels.

In the work titled “*Perceptual color correction through variational techniques*,” Bertalmío et al. [11] start by recalling the variational histogram equalization method of Sapiro and Caselles [44], in which it is shown that the minimization of

the energy functional

$$E(I) = 2 \sum_x \left(I(x) - \frac{1}{2}\right)^2 - \frac{1}{AB} \sum_x \sum_y |I(x) - I(y)| \quad (1)$$

produces an image I with a flat histogram. The range of I is $[0, 1]$, x, y are pixels and A, B are the image dimensions.

In [11] the energy in Eq. (1) is interpreted as the difference between two positive and competing terms

$$E(I) = D(I) - C(I), \quad (2)$$

the first one measuring the dispersion around the average value of $\frac{1}{2}$ (as in the gray world hypothesis), the second term measuring the contrast as the sum of the absolute value of the pixel differences. But this measure of contrast is global, not local, i.e. the differences are computed regardless of the spatial locations of the pixels. This is not consistent with how we *perceive* contrast, which is in a localized manner, at each point having neighbors exert a higher influence than far-away points. Therefore, [11] proposes an adapted version of the functional of Eq. (1) that complies with some very basic visual perception principles, namely those of white patch, locality and not excessive departure from the original data:

$$E(I) = \frac{\alpha}{2} \sum_x \left(I(x) - \frac{1}{2}\right)^2 - \frac{\gamma}{2} \sum_x \sum_y w(x, y) |I(x) - I(y)| + \frac{\beta}{2} \sum_x \left(I(x) - I_0(x)\right)^2, \quad (3)$$

where α, β and γ are constant and positive weights, I is a color channel (R, G or B), $w(x, y)$ is a normalized Gaussian kernel of standard deviation σ , and $I(x)$ and $I(y)$ are two intensity levels at pixel locations x and y respectively.

By minimizing the image energy in Eq. (3) the aim is to maximize the contrast (second term of the functional), while not departing too much from the original image (third term) and also preserving the gray world hypothesis (first term). We can intuitively explain how the contrast term operates by considering the following. Since $\gamma > 0$, in order to minimize $E(I)$ we need to increase $\sum_x \sum_y w(x, y) |I(x) - I(y)|$, i.e. the local contrast, a weighted sum of local pixel differences. So if a pixel has a larger value than all the pixels in its neighborhood, then increasing the local contrast is achieved by enlarging these differences, increasing the pixel value even more. Conversely, when a pixel has a smaller value than all the pixels in its neighborhood, then increasing the local contrast is achieved by making the pixel value even smaller. If $\gamma < 0$, then the minimization of Eq. (3) *reduces*, not increases, the contrast, as pointed out in [10]. The radius σ of the kernel $w(x, y)$ controls how local the contrast enhancement is: a very large σ corresponds to the global case of Eq. (1), whereas with a small σ value the contrast term is computed over small pixel neighborhoods.

It is formulated in [11] that the solution to the minimization of Eq. (3) can be found as the steady state of the evolution equation

$$I^{k+1}(x) = \frac{I^k(x) + \Delta t \left(\frac{\alpha}{2} + \beta I_0(x) + \frac{\gamma}{2} R_{I^k}(x)\right)}{1 + \Delta t(\alpha + \beta)} \quad (4)$$

where the initial condition is $I^{k=0}(x) = I_0(x)$. The function $R_{I^k}(x)$ indicates the contrast function:

$$R_{I^k}(x) = \frac{\sum_{y \in \mathcal{J}} w(x, y) s(I^k(x) - I^k(y))}{\sum_{y \in \mathcal{J}} w(x, y)} \quad (5)$$

where x is a fixed image pixel and y varies across the image. The slope function $s(\cdot)$ is a regularized approximation to the sign function, which appears as it is the derivative of the absolute value function in the second term of the functional; in [11] they choose for $s(\cdot)$ a polynomial of degree 7.

In [11] it is shown that Eq. (3) has a single minimum and that it is also a fixed point of ACE. In other words, we can say that ACE is a numerical implementation of the gradient descent of Eq. (3). The minimization of Eq. (3) yields very good color constancy results and this method shares all the good properties and possible applications of ACE, plus the numerical implementation in [11] has a reduced complexity of $O(N \log N)$, where N is the number of pixels.

This method can be used for contrast enhancement since it produces good results without halos, contrast reversals, spurious colors or any other kind of visual artifact. The main reason for this, formally proven in [11], is that unless σ is too small, the minimization of Eq. (3) preserves the ordering of the level lines of the original image. Furthermore, the algorithm has a local but not pixel-wise nature: the value at each pixel is updated according to comparisons with all its neighbors. This comparison procedure has a sort of anchoring behavior, preventing pixel values from having sudden changes, which helps to explain why the same algorithm applied independently to the frames of a video produces results without temporal artifacts; we have corroborated this in our experiments, as we shall see. Finally, the algorithm does not produce hue shifts either, as it can be seen in the examples in Ferradans et al. [17]. A possible reason is that both large and small contrasts are modified in the same way and by a very similar amount. This amount is the maximum contrast modification allowed by the competition of the contrast term with the other two terms in the equation; therefore, it depends on the weights α , β , γ but not on the magnitude of the local contrast.

There is a very close connection between the formulation of [11] (Eq. (3)) and Retinex. In their kernel-based Retinex (KBR) formulation [10], Bertalmio et al. take all the essential elements of the Retinex theory (channel independence, the ratio reset mechanism, local averages, non-linear correction) and propose an implementation that is intrinsically 2D, and therefore free of the issues associated with paths. The results obtained with this algorithm comply with all the expected properties of Retinex (such as performing color constancy while being unable to deal with overexposed images) but don't suffer from the usual shortcomings such as sensitivity to noise, appearance of halos, etc. In [10] it is proven that there isn't any energy that is minimized by the iterative application of the KBR algorithm, and this fact is linked to its limitations regarding overexposed pictures. Using the analysis of contrast performed by Palma-Amestoy et al. [39], the authors of [10] are able to determine how to modify the

basic KBR equation so that it can also handle overexposed images, and the resulting, modified KBR equation turns out to be essentially the gradient descent of the energy of Eq. (3). In this way, the connection between Retinex, ACE, and the perceptual color correction of [11] (Eq. (3)) becomes explicit.

IV. GAMUT MAPPING VIA ENERGY FUNCTIONAL ADAPTATION

In this section, we adapt the image energy functional defined in Eq. (3) in order to perform gamut mapping from the gamut of the original image to the target gamut of a given device. Firstly, we replace the first term of the functional, $\frac{\alpha}{2} \sum_x (I(x) - \frac{1}{2})^2$, by $\frac{\alpha}{2} \sum_x (I(x) - \mu)^2$, where μ is the mean average of the original image, I_0 . In this way we are penalizing departures from the original mean, instead of imposing the gray world assumption, which could cause a change in the color palette of the result; this modification was already used by Ferradans et al. [17] for local contrast enhancement of tone mapped images. The resulting evolution equation is

$$I^{k+1}(x) = \frac{I^k(x) + \Delta t (\alpha \mu + \beta I_0(x) + \frac{\gamma}{2} R_{I^k}(x))}{1 + \Delta t (\alpha + \beta)} \quad (6)$$

Also, we recall from [10] that $\gamma \in \mathbb{R}$ in the functional is positive or negative depending on whether we want to maximize or minimize the contrast, respectively, and we will use positive values for gamut extension and negative values for gamut reduction.

A. Gamut Reduction Algorithm

In order to perform a gamut reduction operation we minimize the contrast of the source image using Eq. (6) until the colors of the larger gamut converge to the smaller destination gamut. In cases where we need to apply GR, the contrast coefficient γ will be negative and the value of α will be small. The evolution Eq. (6) has a steady state for each particular set of values for α , β , Δt and γ . For example, in Fig. 1a, a chromaticity diagram is shown with different gamuts (visible spectrum, sRGB gamut, source gamut, target gamut and reproduced gamut). It can be seen that when $\beta = 1$, $\alpha = 0$, and $\gamma = 0$ the steady state of the evolution equation is equivalent to the original image. In the same figure we show that as γ decreases and α increases following a $\frac{|\gamma|}{20}$ slope, the steady state of Eq. (6) has a gamut which is gradually smaller. Fig. 1a shows that, just by selecting a small enough value for γ ($\gamma = -3.21$ in this case) we are already performing gamut reduction. However, in this case, colors that were originally inside the target gamut move inwards too much, and the appearance of the image becomes washed-up, as Fig. 1b shows.

To improve the previous result, we present an iterative method in terms of the contrast coefficient γ . At each iteration, we run Eq. (6) for some particular α , β , and γ until we reach the steady state. The steady state of each iteration will provide us with some pixels of the final result. At iteration 1, we set $\beta = 1$, $\alpha = 0$, and $\gamma = 0$, and therefore the original image

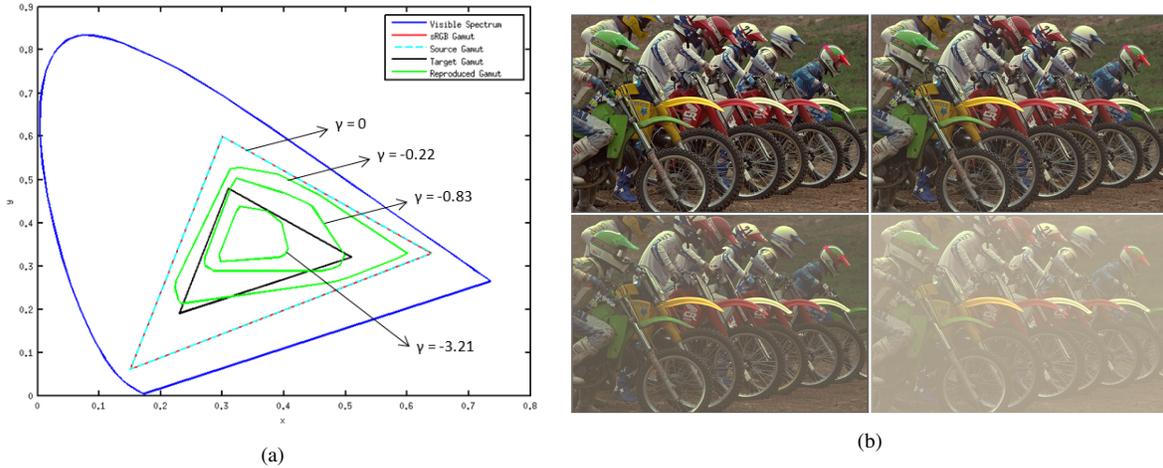


Fig. 1: Perceptual GR Approach. (a): Gamuts on chromaticity diagram. (b): Contrast reduction results. Top left: original image. Top right: $\gamma = -0.22$. Bottom left: $\gamma = -0.83$. Bottom right: $\gamma = -3.21$. As gamma becomes smaller the image becomes more washed out.

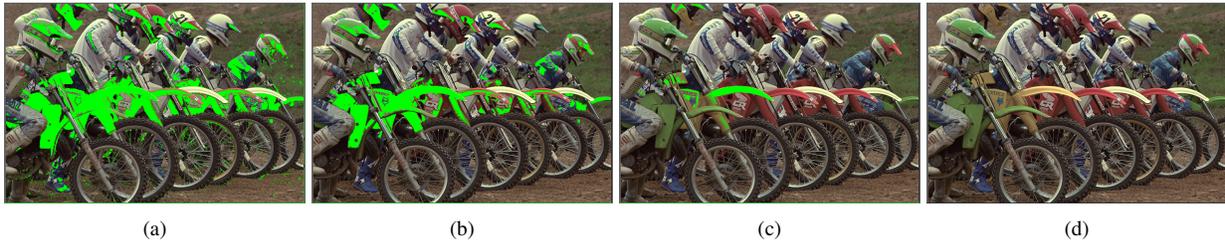


Fig. 2: Gradual mapping of colors. Out-of-gamut colors (in green) when (a): $\gamma = 0$, (b): $\gamma = -0.22$, (c): $\gamma = -0.83$, (d): $\gamma = -3.21$. As gamma decreases the number of out-of-gamut pixels is reduced.

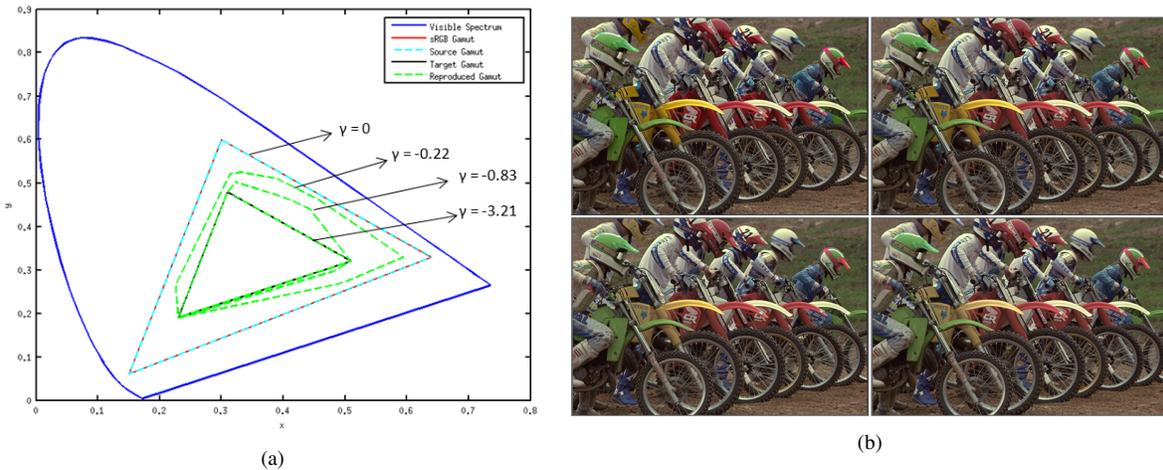


Fig. 3: Modified Perceptual GR Approach. (a): Gamuts on chromaticity diagram. (b): Top left: original image. Top right: $\gamma = -0.22$. Bottom left: $\gamma = -0.83$. Bottom right: $\gamma = -3.21$. In the modified perceptual GR approach, the reproduced image covers a wider range of gamut and appears pleasant in appearance.

is obtained as the steady state. We select the pixels that are inside the destination gamut for the final image and leave them untouched for the following iterations. We move to iteration 2, where we decrease γ (for example, setting $\gamma = -0.05$) and increase α in relation to γ by $\frac{|\gamma|}{20}$. We run again Eq. (6) until steady state, and we check whether any of the colors that were outside the gamut at the previous iteration have been moved

inside the destination gamut. If this is the case, we select them for the final image and leave them untouched for the following iterations. We keep iterating by decreasing γ (and increasing α accordingly) until all the out-of-gamut colors come inside the destination gamut. An example of this iterative procedure is shown in Fig. 2, where green pixels represent out-of-gamut pixels remaining in that iteration. It can be seen in Fig. 3a

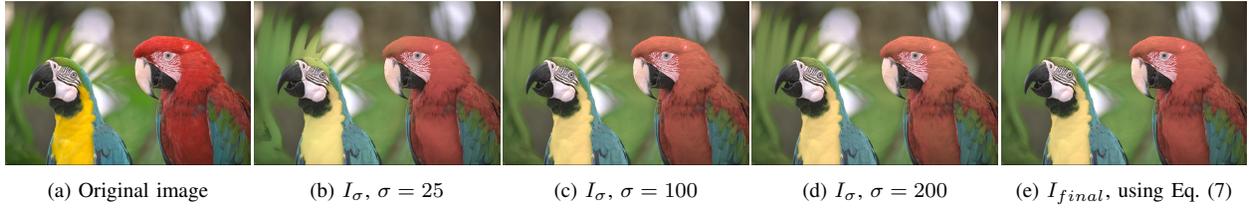


Fig. 4: Effect of standard deviation (σ). The small values of standard deviation (σ) of the Gaussian kernel (w) preserve colors but introduce artifacts. The larger values of σ produce the gamut reduced images that are free from artifacts, but where colors are less saturated.

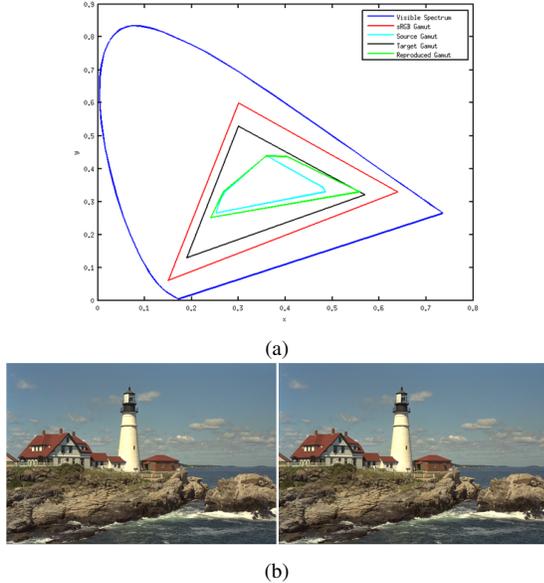


Fig. 5: Perceptual GE Approach. (a): Gamuts on chromaticity diagram. (b): Left: input image. Right: result of first trial of perceptual GEA. As any of the colors from the source image touches the boundary of the target gamut, the reproduction is achieved.

that the reproduced gamut is covering a much wider range of colors than previously. It is shown in Fig. 3b that the colors are better preserved as compared to the previous example (see Fig. 1b).

We also want to note that the standard deviation σ of the Gaussian kernel w is of great importance; we observe in Fig. 4 that a small value of σ leads to the preservation of colors but may introduce a few artifacts, whereas for the larger values of σ each color pixel is strongly influenced by the surrounding colors. Therefore, our method computes several gamut mapped images \mathcal{I}_σ by using four different values of standard deviations $\sigma \in \{\sigma_1, \dots, \sigma_4\}$. Subsequently, in order to obtain a final gamut mapped image \mathcal{I}_{final} , we select for each pixel x , a value out of four gamut mapped images $\mathcal{I}_\sigma(x)$ which has the minimum Lab ΔE distance with compared to the original image value $\mathcal{I}_{orig}(x)$

$$\mathcal{I}_{final}(x) = \arg \min_{\mathcal{I}_\sigma} (\text{Lab}(\mathcal{I}_\sigma(x)) - \text{Lab}(\mathcal{I}_{orig}(x)))^2, \forall x, \\ \sigma \in \{\sigma_1, \dots, \sigma_4\}. \quad (7)$$

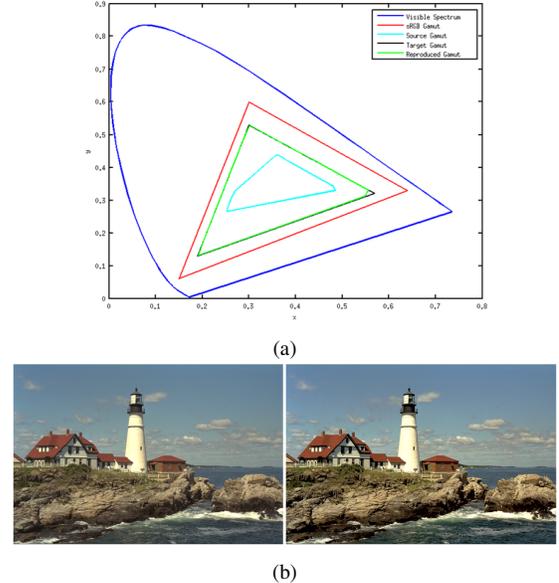


Fig. 6: Modified Perceptual GE Approach. (a): Gamuts on chromaticity diagram. (b): Left: input image. Right: result of proposed perceptual GEA. The result of GEA are pleasingly saturated (see the grass, water, roof of the house) and covers a wider range of gamut.

This procedure increases the computational cost, but the quality of the reproduced image improves to a great extent.

B. Gamut Extension Algorithm

In this section, the goal is to develop a gamut extension algorithm that is capable of accurately expanding the colors from a smaller source gamut to a larger destination gamut. To this end, we enhance the contrast of the source image by running the evolution Eq. (6) to the steady state while keeping the sign of γ positive and defining α as $\frac{|\gamma|}{20}$. To select the value of γ that is adequate for the extension, we pick the minimum γ for which any of the colors touches the boundary of the target gamut as it is presented in Fig. 5a. The drawback in this case is that the reproduced image (Fig. 5b) is less saturated, so we can't realize the full potential of the reproduction device in terms of colors. To overcome this problem, we apply Eq. (6) to larger values of γ and let the original gamut exceed the destination gamut upto a certain threshold level (T_e). This threshold T_e controls the level of saturation; a large value of T_e

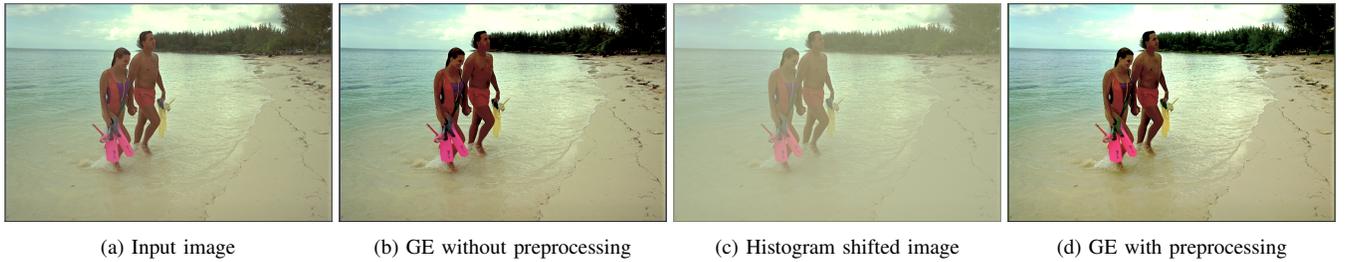


Fig. 7: Effect of preprocessing stage in GEA. Notice the trees in the background; the colors are preserved when GE is applied after preprocessing.

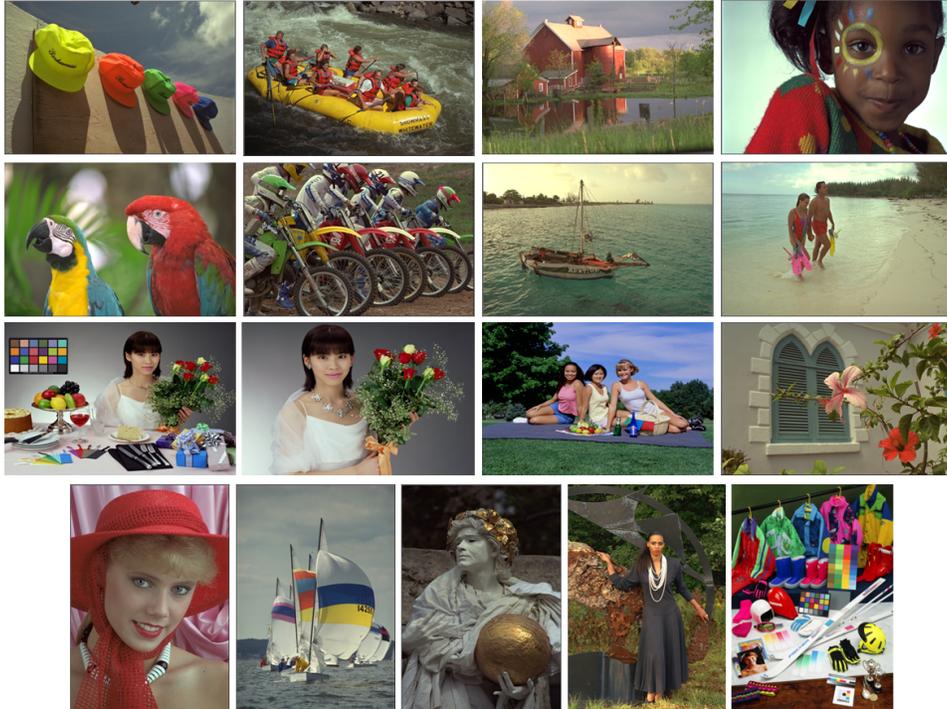


Fig. 8: Original sRGB images. The first 3 images of row 3 and the fifth image of row 4 are from CIE [14]. Rest of the images are courtesy of Kodak.

indicates a higher saturation level, whereas a small value of T_e yields a less saturated output (however, the higher saturation level comes at the cost of less perceptually valid results). After this, the colors that were placed outside the destination gamut in the previous stage are mapped back inside using our gamut reduction algorithm. It is shown in Fig. 6b that the reproduced image is much more pleasing and exhibits a wider gamut (Fig. 6a) as compared with the gamut-extended image in Fig. 5b.

A preprocessing might be needed in cases where the source image has a medium-to-large dynamic range, since in this case when we apply the GEA the colors with low intensity levels tend to go towards black, something which is evident in Fig. 7b (see trees and forearms of the man). This indicates that the image content has a significant impact on the results of the GEA. To cope with this, we introduce a preprocessing step in our GEA where we give a shift to the histogram of each channel of the source image, from low towards high intensity values. The shift in histograms depends on the difference

between the source and the target gamut; the shift is larger if the difference between gamuts is large, whereas a small shift is required when the difference between gamuts is smaller. As we mentioned earlier, we let our GEA extend the color gamut of the source image outside a certain threshold (T_e) of the destination gamut. Once the threshold T_e is achieved, we then take the corresponding γ and estimate the shift value as

$$\text{shift}(c) = \frac{\gamma \bar{I}_c}{\gamma_{ref}} \quad c \in \{R, G, B\} \quad (8)$$

where \bar{I}_c denotes the mean of the color channel I_c . In our GE experiments we observe that even for a large difference between the source and the target gamuts, a maximum amount of contrast coefficient $\gamma = 2$ is required, therefore we fix the reference gamma value $\gamma_{ref} = 2$ to calculate the shift factor for the histograms. We perform preprocessing on the image by giving a shift to the histogram of each channel independently using the shift value of Eq. (8). Even though this improvement



Fig. 9: Results of GRAs on still images. Column 1: original images. Column 2: output of HPMINDE clipping [38]. Column 3: output of Lau et al. [29]. Column 4: output of Alsam et al. [5]. Column 5: output of our algorithm. Compare colors and details in the following regions: row 1, yellow bike’s mudguard and green helmet; row 2, colors at girl’s shoulder; row 3, color of red suit, beanie, socks, helmet and floor; row 4, fruit bowl, cake and gift boxes. Results of HPMINDE [38] and Lau et al. [29] are courtesy of Cheryl Lau.

in the model comes at the expense of some computational complexity, nonetheless the results are much more enhanced as shown in Fig. 7d.

V. EXPERIMENTS AND RESULTS

As stated in the introduction, good GMAs should not only reproduce the colors accurately but they should also generate results that match the inputs in terms of perception. In this section, we present both the subjective and objective evaluation of our GMAs on some very commonly used color images in the literature for gamut mapping applications. The original sRGB images are illustrated in Fig. 8 and these images exhibit a variety of spatial and chromatic characteristics. Furthermore, we perform experiments on videos, confirming that our GMAs can be applied independently on a frame-by-frame basis without producing temporal artifacts.

To compute the results for both GRA and GEA, we work in the RGB color space by fixing the parameters $\beta = 1$, $\Delta t = 0.10$ and $\alpha = \frac{\lfloor \gamma \rfloor}{20}$. We apply Eq. (6) to the source image until convergence, which we assume that has been reached when the difference between two consecutive steps falls below 0.5%. In case of GR, we decrease the gamma value with the change $\Delta\gamma = 0.03$ in order to obtain the final gamut mapped as explained in section IV-A. In the case of GE we need to select a γ value which is enough to extend the gamut of the source image to a certain threshold (T_e).

The primaries of gamuts that are used in this paper are summarised in Table I, and we will refer the gamuts using the naming convention mentioned in the same table.



Fig. 10: Detail preservation using GRAs on still images. Column 1: original cropped regions. Column 2: output of HPMINDE [38]. Column 3: output of Lau et al. [29]. Column 4: output of Alsam et al. [5]. Column 5: output of our algorithm. This is a closer (zoomed-in) view of the regions cropped from Fig. 9.

TABLE I: Primaries of gamuts.

Gamuts	Red Primaries		Green Primaries		Blue Primaries	
	x	y	x	y	x	y
BT.709/sRGB	0.640	0.330	0.300	0.600	0.150	0.060
DCI-P3	0.680	0.320	0.265	0.690	0.150	0.060
Simulated BT.709	0.610	0.330	0.330	0.530	0.150	0.060
Simulated DCI-P3	0.640	0.330	0.300	0.600	0.150	0.060
Toast	0.570	0.320	0.300	0.530	0.190	0.130
Toy	0.510	0.320	0.310	0.480	0.230	0.190

A. Results of GRA

To perform GR, we apply the procedure explained in section IV-A to map the colors of the original image inside the target gamut. We perform GR using four different values for the standard deviation $\sigma \in \{50, 100, 150, 200\}$ of the Gaussian kernel w , and combine all the four gamut reduced versions into a final gamut mapped image using Eq. (7).

1) *Subjective Quality Assessment*: In order to validate the efficiency of our GRA, we apply our method on a rather challenging target gamut named ‘Toy’ in Table I. Given an image in sRGB, our algorithm maps the gamut of the original image into the gamut ‘Toy’. The results presented in Fig. 9 show that our proposed framework works well in preserving the colors, texture and color gradients from the out-of-gamut regions while staying faithful to the perception of the original image. For example, in Fig. 10, rows 1 and 4, it can be seen that the colors reproduced by our GM algorithm (fifth column) are much more saturated than those of HPMINDE [38] (second column), and the state of the art algorithms of Lau et al. [29] (third column) and Alsam et al. [5] (fourth column). Similarly, in Fig. 10, row 2, our algorithm not only reproduces the color efficiently but also preserves a great amount of texture. In Fig. 10, row 3, we can see our method accurately represents the difference in the lightness of identical hue (see the pink socks



(a) Image sequence 1 (professional footage)



(b) Image sequence 2 (amateur video)

Fig. 11: Results of GRA. In both image sequences: top row, original frames; middle row, output of Alsam et al. [5]; bottom row, output of our algorithm. In the first 3 columns of image sequence 1, notice the color of the girl’s skin, her shirt and the poster on wall, whereas in the last 2 columns notice the colors of the door, cushions on sofas and the lamp. In image sequence 2, differences in perception and color reproduction can be seen in almost every region of the frames such as floor, inside the shop, and clothes of pedestrians and street artists. Original image sequence 1 is property of CBS Interactive and taken from [1]. Original image sequence 2 is from [2].

TABLE II: Quality assessment of GRAs on still images (Fig. 8): CID perceptual error.

	Mean	Median	RMS
HPMINDE Clipping [38]	0.0674	0.0514	0.0873
Lau et al. [29]	0.0665	0.0695	0.0807
Alsam et al. [5]	0.0472	0.0398	0.0627
Our GRA	0.0368	0.0285	0.0495

and pink beanie).

In Fig. 11 we show the results for our gamut reduction method

and for the state of the art algorithm of Alsam et al. [5] on two different image sequences, one professional and the other amateur, from which we have selected challenging shots that are colorful and have noticeable camera and object motion. Both GRAs are applied to each frame independently. We used the target gamuts ‘Toy’ and ‘Toast’ to reduce the gamut of image sequence 1 and image sequence 2 respectively. It can be seen that, for both sequences (selected frames) shown in the figure, our algorithm produces more saturated results as compared with the method of Alsam et al. [5]. For example,

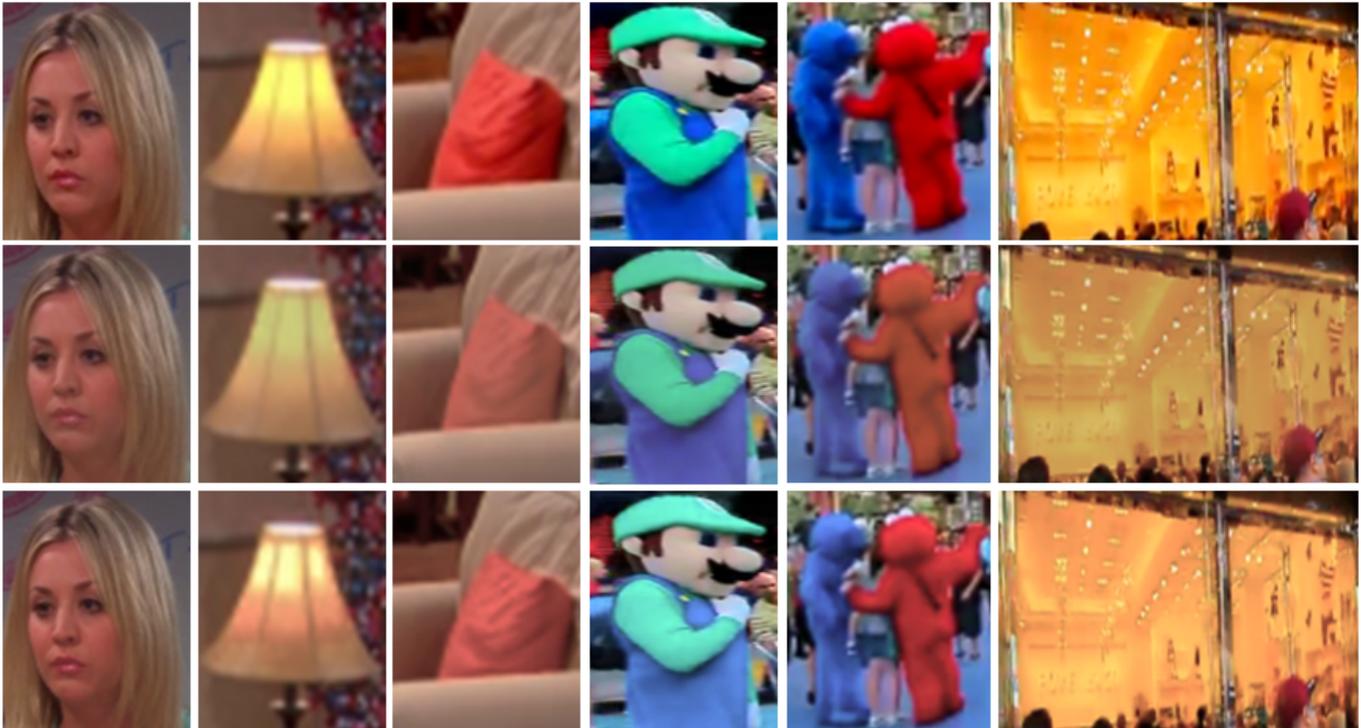


Fig. 12: Detail preservation using GRAs on video, all regions are cropped from Fig. 11. Top row: original cropped regions. Middle row: output of Alsam et al. [5]. Bottom row: output of our algorithm.

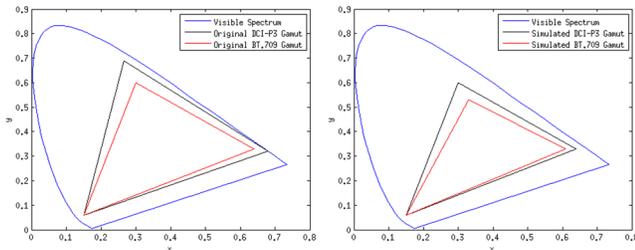


Fig. 13: Simulation of relationship between DCI-P3 and BT.709 gamuts. Left: original gamuts. Right: simulated gamuts.



Fig. 14: Spatial artifacts due to GRA. Top row: original images. Middle row: effect of extreme target gamut. Bottom row: effect of realistic target gamut. The relationship of target gamut and spatial artifacts can be seen at the forehead of the person, cup and on the pants.

in Fig. 12, the color reproduction of the girl’s skin (column 1) is perceptually more faithful in our results than in those of Alsam et al. [5]. Furthermore, in columns 2 and 6, we can see that the method of [5] has problems preserving the color of the very bright regions, whereas our algorithm efficiently reproduces those regions. It can be seen in column 4 that the algorithm of [5] introduces a shift in the hue (see blue color), while on the other hand our method represents colors more faithfully. Moreover, the colors reproduced by our algorithm are more saturated and pleasant as compared with [5] (see column 5).

Since we are using an extreme target gamut to map the colors of image sequence 1, we experience a few spatial artifacts in our results; similar artifacts are observed in the results obtained with the method in [5]. In cinematography, the content of digital video needs to be mapped from the camera gamut to

the BT.709 gamut, and these gamuts exhibit not so extreme differences with respect to each other. Therefore, to test our GRA under a real scenario, we obtain a new target gamut by simulating the difference between DCI-P3 and BT.709 gamuts inside the BT.709 gamut. The simulation of gamuts is shown in Fig. 13. It turns out that our GRA performs efficiently under realistic settings without introducing any spatial artifacts as it can be seen in Fig. 14. Moreover, we observed that it makes sense to apply our GRA on each frame independently, since there are no temporal artifacts and therefore no coherence

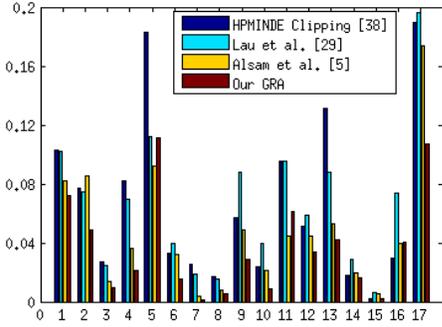


Fig. 15: Quality assessment of GRAs on the still images from Fig. 8 using the CID [30]. Horizontal axis: row-wise image number of Fig. 8. Vertical axis: perceptual error.

between frames is required. The full videos are available for the reviewers at www.ip4ec.upf.edu/node/91.

2) *Objective Quality Assessment*: Visually, the results presented so far underline the good performance of our GRA in terms of visual quality. This subjective outcome is backed by using the perceptual color quality measure presented in [30]: the Color Image Difference (CID) metric estimates the perceptual differences given by the changes, from one image to the other, in features such as hue, lightness, chroma, contrast and structure.

Comparisons using the CID metric in the still images of Fig. 8 are provided in Fig. 15. In this figure, we can see that our algorithm outperforms the other methods in 15 out of 17 test images. Moreover, the CID statistical data (mean, median and root mean square) is also presented in Table II, which justifies the reliable performance of our algorithm over the other approaches. In addition to objective quality tests on still images we also compare, using the CID metric, the perceptually-based efficiency of our algorithm and Alsam et al. [5] on the image sequences shown in Fig. 11 and on an animated video (where we perform reduction from the BT.709 to the ‘Toast’ gamut). The graphs shown in Fig. 16 and the CID statistical data presented in Table III depict that our method produces results with less perceptual error than the other method.

In conclusion, both subjective and objective assessments show that our GRA produces a gamut reduced video which is, perceptually, more faithful to the original video as compared with the other methods.

B. Results of GEA

1) *Subjective Quality Assessment*: To quantify the performance of our GE approach, we use ‘BT.709’ as a target gamut and set the saturation threshold $T_e = 10\%$ of the total number of pixels. We apply the proposed GEA on videos that have the starting color gamut equal to the ‘Toast’ gamut (for the sake of fair comparisons, these videos are reduced to the ‘Toast’ gamut using clipping in the Lab color space). We compare our

TABLE III: Quality assessment of GRAs on videos: CID perceptual error.

	Mean	Median	RMS
Alsam et al. [5]	0.0528	0.0413	0.0623
Our GRA	0.0372	0.0221	0.0474

algorithm with the SDS algorithm [26], which is similar to the industrial method of performing gamut extension in new TV screens [41]. We evaluate results of our GEA and SDS [26] by comparing them with the ground truth in order to analyze their efficiency. In Fig. 17, we present the results of our gamut extension procedure on videos where it can be seen that our expansion algorithm generates images that are pleasant and realistic in appearance. Fig. 18 shows zoomed-in details from Fig. 17. In columns 1 and 5, it can be seen that our GEA enhances the color of trees while keeping the sky color closer to the ground truth, whereas SDS [26] noticeably changes the color of sky regions. It is shown in columns 2 and 4 that the reproduction of skin color (face of the man and woman’s arms) by our GEA is realistic in appearance, but the SDS algorithm [26] reproduces flesh tones very much departed from the ground truth. The results also show that the SDS approach [26] over-saturates some regions, which in turn looks artificial. For example, in columns 3 and 6, it is clearly noticeable that our GEA achieves very good color reproduction (wall, hood of the shed and hat) without introducing any unnatural colors. On the other hand, the results produced by [26] are highly saturated and quite different from the original content. We have noticed that our GEA not only extends the color gamut perceptually but also performs contrast enhancement, which is evident in Fig. 18 (columns 3 and 6); if this were a problem, the original contrast could easily be restored with a very small amount of simple blurring.

Our GEA shows a robust performance in both image sequences. It can also be seen that the results obtained using our algorithm are less saturated than those of SDS [26] in a few regions. If we increase the saturation threshold (T_e) our algorithm will achieve more saturated colors, at the expense of increasing the perceptual error of the results.

2) *Objective Quality Assessment*: The primary purpose of our gamut extension algorithm is to process the colors of a digital video in order to display it with digital cinema projectors. In this case, a video should be mapped from an BT.709 source gamut to a DCI-P3 target gamut. However, paper evaluation of the results would require a wide gamut digital cinema projector to display the mapped video under special (cinematic) lighting conditions and/or a way to present the DCI-P3 results in the paper, as we are limited by sRGB for the article. Therefore, to quantify the performance of our gamut expansion method, we take an experimental approach where we simulate the difference between DCI-P3 and BT.709 gamuts inside the BT.709 gamut. Thereafter, we convert the BT.709 videos to the simulated sRGB gamut using a clipping on the xy components in the xyY color space while keeping the luminance component unchanged. Finally, we map the videos from the simulated BT.709 gamut to the simulated

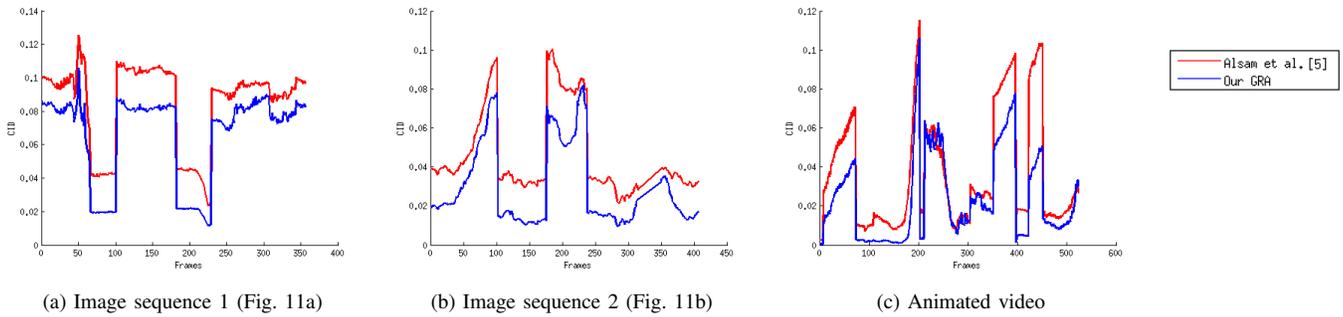
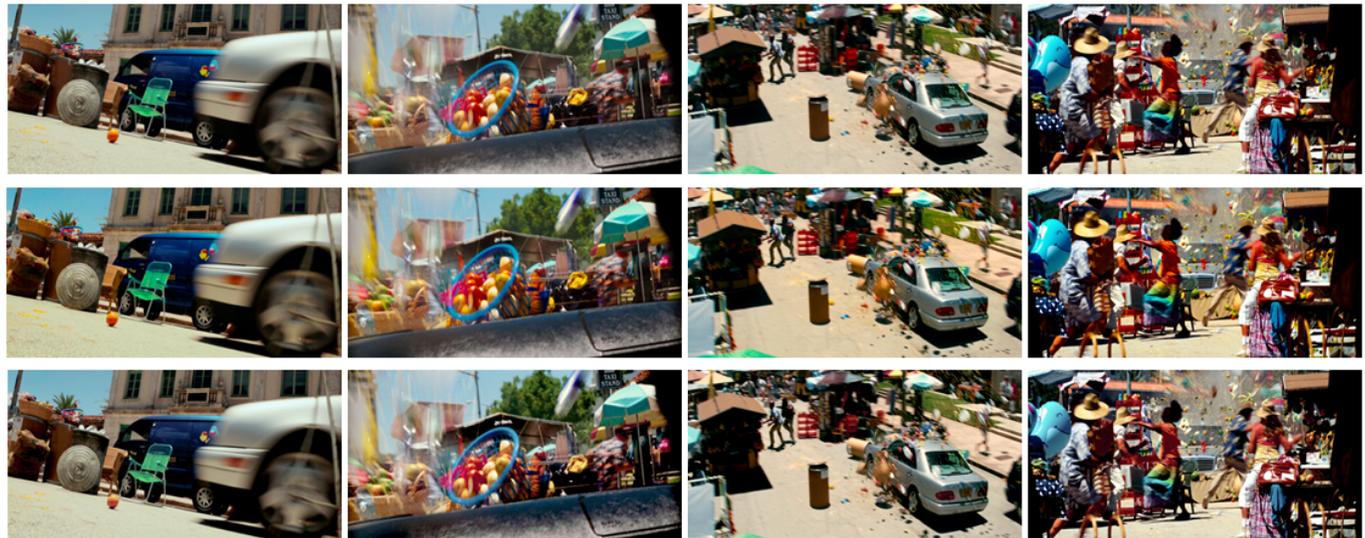


Fig. 16: Quality assessment of GRAs on videos using the perceptual difference measure CID [30]. Horizontal axis: time (in frame numbers). Vertical axis: perceptual error.



(a) Image sequence 1



(b) Image sequence 2

Fig. 17: Results of GEA. In both image sequences; top row, ground truth (BT.709 gamut); middle row, output of Same Drive Signal algorithm [26]; bottom row, output of our GEA. In both image sequences, the results of our GEA are more pleasing and perceptually faithful to the ground truth as compared with the output of the SDS algorithm [26]; see zoomed-in details in Fig. 18. Both image sequences are the property of Paramount Pictures and taken from [3].

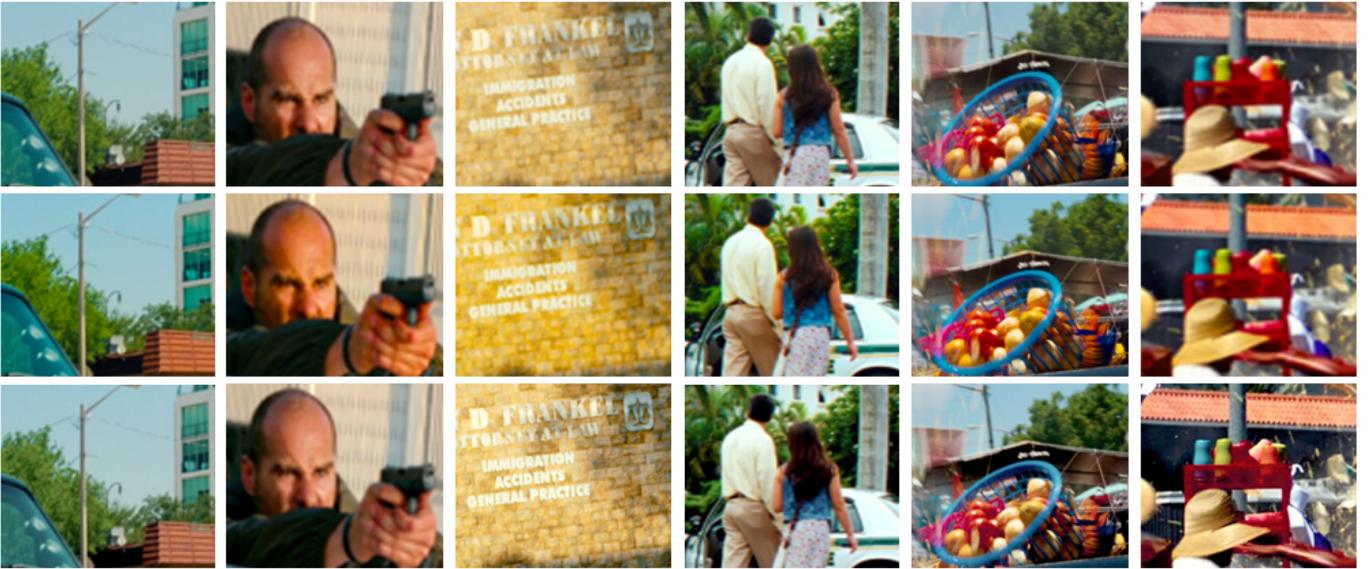


Fig. 18: Comparison of GEAs, details from Fig. 17; top row, ground truth; middle row, output of Same Drive Signal algorithm [26]; bottom row, output of our GEA. The results of the SDS algorithm [26] are highly saturated and off from the ground truth. See the reproduction of sky, trees, skin tone, wall, clothes, fruit bucket and hat.

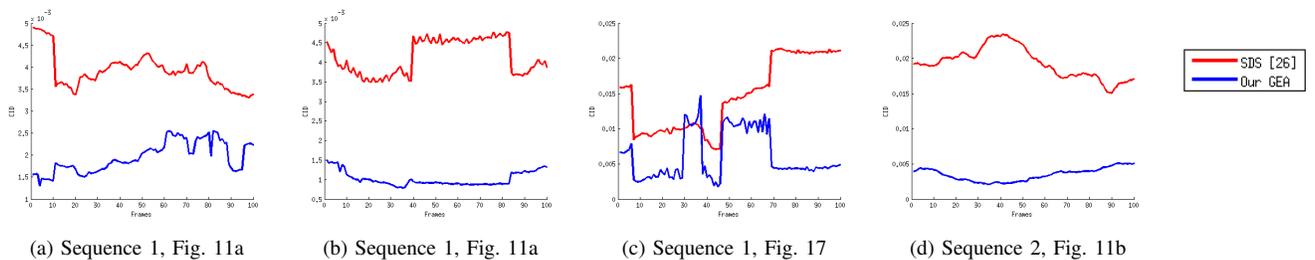


Fig. 19: Quality assessment of GEAs on videos using the perceptual metric CID [30]. Horizontal axis: time (in frame numbers). Vertical axis: perceptual error.

DCI-P3 gamut. Since in this case we have a ground truth to compare with (simulated DCI-P3 gamut is equal to the original BT.709 gamut), we can assess our gamut extension algorithm quantitatively using the CID metric. For this experiment we set $T_e = 3\%$ since source and target gamuts have small color differences. The results are presented in Fig. 19, where it can be seen that our reproduction is perceptually more faithful to the original content than the results of SDS. The CID statistical data is summarised in Table IV, which justifies the good the performance of our GEA on videos.

We have observed in our experiments that the quality of the input signal also matters; if there exist any artifacts due to compression or noise in the input video they may become prominent once passed through the GE module. A possible way to deal with such kind of input material is to perform preprocessing using the approach reported in [40].

Finally, we want to stress that, as it was the case with our GRA, no temporal artifacts appear in the results of our GEA either.

TABLE IV: Quality assessment of GEAs on videos: CID perceptual error.

	Mean	Median	RMS
SDS [26]	0.0105	0.0060	0.0128
Our GEA	0.0031	0.0024	0.0041

VI. CONCLUSIONS

We have presented gamut mapping algorithms based on a perceptually inspired variational framework, which is modified in order to perform gamut reduction and gamut extension. The main advantage of our method is its perceptual inspiration, that allows us to mimic some basic properties of the human visual system while performing the mapping, and this is corroborated by the good scores we obtain with a perceptual metric for color and contrast distortion. The proposed gamut mapping algorithms have shown a robust performance in a variety of scenes which indicates their strong potential to be used in color management pipelines both in the film and broadcast industries.

For future work we intend to make use of other perceptual metrics to optimize our gamut reduction algorithm. In the case

of gamut extension, we will perform an intensive subjective evaluation using special equipment in controlled viewing conditions.

ACKNOWLEDGEMENTS

This work was supported by the European Research Council, Starting Grant ref. 306337, and by Spanish grants ref. TIN2011-15954-E and ref. TIN2012-38112. The authors would like to thank Jan Morovič, Ivar Farup and Ali Alsam for their comments and suggestions and for sharing their codes. Many thanks go to Cheryl Lau for running experiments on a set of images.

REFERENCES

- [1] http://www.cbs.com/shows/big_bang_theory/.
- [2] <http://www.youtube.com/watch?v=ezyrSKgcyJw>.
- [3] <http://www.painandgainmovie.com>.
- [4] A. Alsam and I. Farup. Colour gamut mapping as a constrained variational problem. In *Proc. of 16th Scandinavian Conference on Image Analysis*, pages 109–118, 2009.
- [5] A. Alsam and I. Farup. Spatial colour gamut mapping by orthogonal projection of gradients onto constant hue lines. In *Proc. of 8th International Symposium on Visual Computing*, pages 556–565, 2012.
- [6] H. Anderson, E. Garcia, and M. Gupta. Gamut expansion for video and image sets. In *Proc. of the 14th International Conference of Image Analysis and Processing - Workshops*, pages 188–191, 2007.
- [7] R. Bala, R. Dequeiroz, R. Eschbach, and W. Wu. Gamut mapping to preserve spatial luminance variations. *Journal of Imaging Science and Technology*, 45:122–128, 2001.
- [8] D. Bankston. The color-space conundrum, part one. *American Cinematographer*, page 6, 2005.
- [9] M. Bertalmío. *Image Processing for Cinema*. CRC Press, Taylor & Francis, 2014.
- [10] M. Bertalmío, V. Caselles, and E. Provenzi. Issues about retinex theory and contrast enhancement. *International Journal of Computer Vision*, 83(1):101–119, 2009.
- [11] M. Bertalmío, V. Caselles, E. Provenzi, and A. Rizzi. Perceptual color correction through variational techniques. *IEEE Transactions on Image Processing*, 16(4):1058–1072, 2007.
- [12] N. Bonnier, F. Schmitt, H. Brettel, and S. Berche. Evaluation of spatial gamut mapping algorithms. In *Proc. of 14th Color Imaging Conference*, 2006.
- [13] S. E. Casella, R. L. Heckaman, and M. D. Fairchild. Mapping standard image content to wide-gamut displays. In *Sixteenth color imaging conference: color science and engineering systems, technologies, and applications*, pages 106–111, 2008.
- [14] CIE. Guidelines for the evaluation of gamut mapping algorithms. Technical report, 2004.
- [15] F. Ebner and M. D. Fairchild. Gamut mapping from below: Finding minimum perceptual distances for colors outside the gamut volume. *Color Research and Application*, 22:402–413, 1997.
- [16] I. Farup, C. Gatta, and A. Rizzi. A multiscale framework for spatial gamut mapping. *IEEE Transactions on Image Processing*, 16(10):2423–2435, 2007.
- [17] S. Ferradans, M. Bertalmío, E. Provenzi, and V. Caselles. An analysis of visual adaptation and contrast perception for tone mapping. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, October 2011.
- [18] R. L. Heckaman and J. Sullivan. 18.3: Rendering digital cinema and broadcast tv content to wide gamut display media. *SID Symposium Digest of Technical Papers*, 42(1):225–228, 2011.
- [19] T. Hoshino. A preferred color reproduction method for the hdtv digital still image system. In *Proc. of IS&T Symposium on Electronic Photography*, pages 27–32, 1991.
- [20] K. Jang, H. Cho, H. Kang, and B. Kang. A color gamut mapping system using the rgb primaries and white-point correction for a wide gamut display device. In *Proc. of International Conference on Future Information & Communication Engineering*, pages 559–568, 2013.
- [21] B. H. Kang, J. Morovic, M. R. Luo, and M. S. Cho. Gamut compression and extension algorithms based on observer experimental data. *ETRI journal*, 25(3):156–170, 2003.
- [22] N. Katoh, M. Ito, and S. Ohno. Three-dimensional gamut mapping using various color difference formulae and color spaces. *Journal of Electronic Imaging*, 8(4):365–379, 1999.
- [23] G. Kennel. *Color and mastering for digital cinema: digital cinema industry handbook series*. Taylor & Francis US, 2007.
- [24] M. C. Kim, Y. C. Shin, Y. R. Song, S. J. Lee, and I. D. Kim. Wide gamut multi-primary display for hdtv. In *Proc. of 2nd European Conference on color Graphics, Imaging and Vision*, pages 248–253, 2004.
- [25] R. Kimmel, D. Shaked, M. Elad, and I. Sobel. Space-dependent color gamut mapping: A variational approach. *IEEE Transactions on Image Processing*, 14:796–803, 2005.
- [26] J. Laird, R. Muijs, and J. Kuang. Development and evaluation of gamut extension algorithms. *Color Research & Application*, 34(6):443–451, 2009.
- [27] E. H. Land. The Retinex theory of color vision. *Scientific American*, 237:108–128, 1977.
- [28] E. H. Land and J. J. McCann. Lightness and retinex theory. *Journal of the Optical Society of America*, pages 1–11, 1971.
- [29] C. Lau, W. Heidrich, and R. Mantiuk. Cluster-based color space optimizations. In *Proc. of IEEE International Conference on Computer Vision, ICCV '11*, pages 1172–1179, 2011.
- [30] I. Lissner, J. Preiss, P. Urban, M. S. Lichtenauer, and P. Zolliker. Image-difference prediction: From grayscale to color. *IEEE Transactions on Image Processing*, 22(2):435–446, 2013.
- [31] Y. Liu, G. Song, and H. Li. A hue-preserving gamut expansion algorithm in cielv color space for wide gamut displays. In *Image and Signal Processing (CISP), 2010 3rd International Congress on*, volume 5, pages 2401–2404, 2010.
- [32] J. J. McCann. Lessons learned from mondrians applied to real images and color gamuts. In *Proc. of Color Imaging Conference*, pages 1–8, 1999.
- [33] J. J. McCann. A spatial colour gamut calculation to optimize colour appearance. In *Colour Image Science: Exploiting Digital Media*, pages 213–233, 2002.
- [34] J. Meyer and B. Barth. Color gamut matching for hard copy. In *Proc. of SID Digest*, pages 86–89, 1989.
- [35] J. Morovič. *Color gamut mapping*, volume 10. Wiley, 2008.
- [36] J. Morovič and Y. Wang. A multi-resolution, full-colour spatial gamut mapping algorithm. In *Proc. of Color Imaging Conference*, pages 282–287, 2003.
- [37] R. Muijs, J. Laird, J. Kuang, and S. Swinkels. Subjective evaluation of gamut extension methods for widegamut displays. In *Proc. of the 13th International Display Workshop*, pages 1429–1432, 2006.
- [38] G. M. Murch and J. M. Taylor. Color in computer graphics: Manipulating and matching color. *Eurographics Seminar: Advances in Computer Graphics V*, pages 41–47, 1989.
- [39] R. Palma-Amestoy, E. Provenzi, M. Bertalmío, and V. Caselles. A perceptually inspired variational framework for color enhancement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(3):458–474, 2009.

- [40] R. Palma-Amestoy, E. Provenzi, M. Bertalmío, and V. Caselles. A perceptually inspired variational framework for color enhancement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(3):458–474, 2009.
- [41] C. Poynton. Wide-gamut displays. *Information Display*, 23(7):10, 2007.
- [42] E. Provenzi, L. De Carli, A. Rizzi, and D. Marini. Mathematical definition and analysis of the Retinex algorithm. *Journal of the Optical Society of America A*, 22(12):2613–2621, December 2005.
- [43] A. Rizzi, C. Gatta, and D. Marini. A new algorithm for unsupervised global and local color correction. *Pattern Recognition Letters*, 24:1663–1677, 2003.
- [44] G. Sapiro and V. Caselles. Histogram modification via differential equations. *Journal of Differential Equations*, 135:238–266, 1997.
- [45] M. C. Stone, W. B. Cowan, and J. C. Beatty. Color gamut mapping and the printing of digital color images. *ACM Transactions on Graphics*, 7(4):249–292, 1988.
- [46] S. W. Zamir, J. Vazquez-Corral, and M. Bertalmío. Gamut mapping through perceptually-based contrast reduction. In *Proc. of the 6th Pacific-Rim Symposium on Image and Video Technology*, pages 1–11, 2013.
- [47] P. Zolliker and K. Simon. Retaining local image information in gamut mapping algorithms. *IEEE Transactions on Image Processing*, 16(3):664–672, 2007.