

Vision Models for Wide Color Gamut Imaging in Cinema

Syed Waqas Zamir, Javier Vazquez-Corral, and Marcelo Bertalmío

Abstract—Gamut mapping is the problem of transforming the colors of image or video content so as to fully exploit the color palette of the display device where the content will be shown, while preserving the artistic intent of the original content’s creator. In particular in the cinema industry, the rapid advancement in display technologies has created a pressing need to develop automatic and fast gamut mapping algorithms. In this paper we propose a novel framework that is based on vision science models, performs both gamut reduction and gamut extension, is of low computational complexity, produces results that are free from artifacts and outperforms state-of-the-art methods according to psychophysical tests. Our experiments also highlight the limitations of existing objective metrics for the gamut mapping problem.

Index Terms—Gamut mapping algorithms, wide gamut imaging, color reproduction, vision models for color and contrast, gamut mapping for cinema.

1 INTRODUCTION

THE range of colors that a device is able to reproduce is called its color gamut. A very common and convenient way of describing colors is to ignore their luminance component and just represent the chromatic content on a 2D plane known as the CIE xy chromaticity diagram, shown in Fig. 1. In this figure the tongue-shaped region corresponds to the chromaticities of all the colors a standard observer can perceive. Most existing displays are based on the trichromacy property of human vision, creating colors by mixing three well-chosen red, green and blue primaries in different proportions. The chromaticities of these primaries determine a triangle in the CIE xy chromaticity diagram, and this triangle is the color gamut of the display in question. Therefore, for any given three-primary display there will be many colors that we could perceive but the display is not able to generate, i.e. all the colors with chromaticities outside the triangle associated to the display. Also, devices with different sets of primaries will have different gamuts. For this reason, in order to facilitate inter-operability a number of standard distribution gamuts have been defined, and for cinema the most relevant ones are shown in Fig. 1: DCI-P3 [1] is the standard gamut used in cinema postproduction and recommended for digital cinema projection, BT.709 [2] is used for cable and broadcast TV, DVD, Blu-Ray and streaming, and BT.2020 [3] is a very wide color gamut for the next generation UHDTV, currently only achievable by some state-of-the-art laser projectors. Fig. 1 also shows Pointer’s gamut [4], which covers all the frequently occurring real surface colors; we can see how only BT.2020 is able to completely include Pointer’s gamut.

The adaptation to a standard gamut implies altering the range of colors (and contrast) of the original content. This

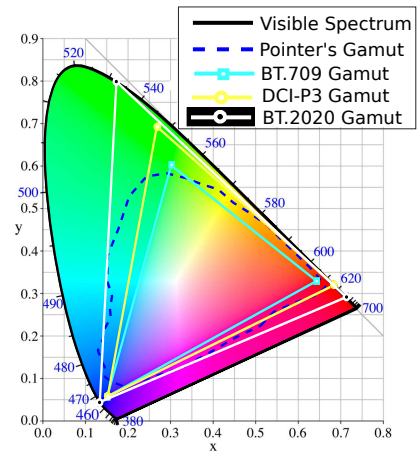


Fig. 1: Gamuts on CIE xy chromaticity diagram.

process is either carried out within the camera in live TV broadcasts (or low-budget movie productions), or performed off-line by expert technicians in the cinema industry. In practice, for the purpose of modifying the movie gamut, colorists at the post-production stage build 3D look-up-tables (LUTs) for each movie or specific scenes in it. These LUTs contain millions of entries and colorists only specify a few colors manually, while the rest are interpolated regardless of their spatial or temporal distribution [5]. Subsequently, the resulting movie may have false colors that were not originally present. To tackle this problem, colorists usually perform intensive manual correction in a shot-by-shot, object-by-object basis. This process is difficult, time consuming and expensive, and therefore it makes an automated procedure called gamut mapping (GM) very desirable: GM transforms an image so that its colors better fit the target gamut.

In general, there are two types of GM procedures. First is gamut reduction (GR), in which colors are mapped from a larger source gamut to a smaller destination gamut. A

Syed Waqas Zamir is with the Inception Institute of Artificial Intelligence, UAE. Email: waqas.zamir@inceptioniai.org,

Javier Vazquez-Corral and Marcelo Bertalmío are with the Universitat Pompeu Fabra, Spain. Email: javier.vazquez@upf.edu; marcelo.bertalmio@upf.edu.

common situation where GR is necessary is when a movie intended for cinema viewing is displayed on a TV [6], [7]. Second is gamut extension (GE), that involves mapping colors from a smaller source gamut to a larger destination gamut. For example, wide-gamut state-of-the-art displays often receive movies that are encoded with limited-gamuts as a precaution against regular (or poor) display devices; therefore, we cannot exploit the full color rendering potential of these new devices unless we use a GE procedure [5]. The process of GE is gaining importance with the introduction of new display technologies and laser projectors [8], [9]. These new displays use pure (very saturated) color primaries which enable them to cover much wider gamuts, so now a tablet screen may have a DCI-P3 gamut for instance, while all the content it shows comes in the smaller BT.709 standard.

At this point, we present a clarification on how gamut reduction and gamut extension differ practically. Gamut reduction is required, not optional, when the colors of the input image fall outside the display's gamut; without GR, the display will reproduce the image with artifacts and loss of spatial detail. On the contrary, gamut extension is not essential, rather it is considered as an enhancement operation [10]. For example, BT.709 footage presented as it is on a wide-gamut BT.2020 display device won't show any visual artifact, it's just that we will be missing the color rendering potential of the wide-gamut screen.

As a main contribution, in this paper we present a framework for gamut mapping that is based on models from vision science and that allows us to perform both gamut reduction and gamut extension. It is computationally efficient and yields results that outperform state-of-the-art methods, as validated using psychophysical tests. Another contribution of our research is to highlight the limitations of existing image quality metrics when applied to the GM problem, as none of them, including two state-of-the-art deep learning metrics for image perception, trained over large and very large scale databases (20,000+ images in one case, 160,000+ in the other) is able to predict the preferences of the observers.

We believe our results are of importance to the computer vision community for two main reasons. Firstly, because they provide another example that drawing insights from vision science and developing algorithms based on vision models can yield state-of-the-art performance for computer vision applications. And secondly, because our results demonstrate how deep learning approaches are not yet suitable to emulate perception with an adequate degree of accuracy, even when using very large databases with a huge number of human annotations. This begs the question of whether this failure is due to limitations in the network architecture, or rather a more intrinsic issue is at hand, as for instance it has been argued that the convolution-based spatial summation of artificial neural networks cannot constitute a proper model of how biological networks process information [11].

2 RELATED WORK

A large number of gamut mapping algorithms (GMAs) have been proposed in the literature, we refer the interested reader to the comprehensive book of Morović [10]. GMAs can be divided into two main categories: gamut reduction

algorithms (GRAs) and gamut extension algorithms (GEAs). Both GRAs and GEAs can further be classified into two subclasses: global and local. Global (also known as non-local or non-adaptive) methods map colors of an image to the target gamut independently, while completely ignoring the spatial distribution of colors in the image. Whereas local (also known as spatial) methods modify pixel values by taking into account their neighborhoods; as a result, two identical values surrounded by different neighborhoods will be mapped to two different values.

Global GRAs: One class of global GRAs consists of gamut clipping methods [12], [13], [14]. Gamut clipping is a very common approach to perform gamut reduction where colors that lie inside the destination gamut are left untouched while those colors that fall outside are projected onto the destination gamut boundary. In order to produce reduced-gamut images, gamut clipping techniques use particular strategies and mapping directions. For example, clipping chroma of the out-of-gamut (OOG) colors along lines of constant hue and lightness [15]; clipping each OOG color by seeking a minimum ΔE distance to the destination gamut boundary along lines of constant hue, this method is referred as hue preserving minimum ΔE (HPMINDE) [16]; clipping colors of low luminance and high luminance differently, thereby providing a special treatment to bright colors in order to avoid excessive chroma loss [17]. Clipping GRAs, due to their inherent behavior, project whole OOG color segments into single points on the target gamut, and therefore they may reproduce images with a visible loss of detail and color gradients. To avoid this sort of issue, another class of global GRAs are called gamut compression algorithms. These methods modify all colors of an input image, both inside and outside the target gamut. Such functionality enables compression GRAs to map OOG color segments to in-gamut color segments (instead of single points), though the results they produce may lack in saturation, specially when the difference between the source and target gamuts is large. Many compression GRAs have been proposed in the literature [18], [19], [20], [21], [22]. Some GRAs take an hybrid approach in which a combination of clipping and compression is used to perform gamut mapping; see for example [23]. A few GRAs [24], [25] make use of soft-clipping that squeezes colors near the target gamut boundary in order to accommodate the OOG colors.

Local GRAs: The frequency-based local GRAs [26], [27], [28] first reduce the gamut of the source image using a global method, and then in the second stage the high frequency image detail (obtained by using a spatial filter) is added to the reduced-gamut image. In these GRAs, another stage of gamut clipping is integrated to process the resulting image in case the spatial filtering operation places a few pixels outside the destination gamut. Local GRAs that are inspired from the Retinex framework perform spatial comparisons to retain source image gradients in the reproduced images [29], [30], [31], [32]. Some spatial GRAs [33], [34], [35], [36], [37] pose gamut mapping as an optimization scheme where, given a source image and its gamut mapped version, the aim is to keep perturbing the gamut mapped image until its difference w.r.t. the source image is minimized according to an error metric. Finally, an image energy functional [38] is introduced to decrease the contrast of the input image in

order to perform gamut reduction.

Global GEAs: While the majority of the published GMAs deal with the problem of gamut reduction, the case is very different for gamut extension: only a few works have been proposed in this direction. One simple solution to perform gamut extension is to take any compression GRA and use it in the reverse direction [39], [40], [41]. However, this way of approaching GE may yield images that are unnatural and unpleasant in appearance. The pioneering global GEAs [42], [43] map limited-gamut printed images to the wide gamut of HDTV in two stages: firstly the lightness is mapped using a non-linear tone reproduction curve, and secondly the chroma is extended along lines of constant hue and lightness. A few methods [44], [45] perform gamut extension using functions learned from user studies. Unlike the aforementioned GEAs, some global methods [46], [47], [48] first classify the colors of the input image according to a criterion, and then perform gamut extension differently for each class. For example, labelling each color of a given image as skin or non-skin [46]; dealing with objects of low chroma and high chroma differently [47]; identifying certain memory colors such as green grass and blue sky, and rendering them independently [48]. Other approaches [49], [50] propose three types of extensions: chroma extension, extension along lines from the origin, and adaptive mapping that is a compromise between the first two strategies. Some global GEAs [51], [52], [53] aim at preserving skin tones in the reproduced images.

Local GEAs: The local GEAs extend colors by taking into account their spatial distribution in the input image. This property certainly makes local GEAs adaptive and flexible but at the same time far more complex and computationally expensive than global GEAs. The multilevel GEA [54] in its first stage extends the source gamut using a non-linear hue-varying function, and in the second stage applies an image-dependent chroma smoothing operation to avoid an over-enhancement of contrast and to preserve detail in the final image. Recent works [38], [55], [56] perform spatial gamut extension using partial differential equations. In particular, the contrast of the input image is enhanced by minimizing an energy functional [38]; a monotonically increasing function [55] is applied on the saturation channel of the input image in HSV color space that allows to increase contrast without decreasing the image saturation values; and the GEA [56] operates only on the chromatic components of CIELAB color space, while taking into account the analysis of distortions in hue, chroma and saturation.

In this paper we propose a novel framework that overcomes several issues of current GM approaches, performing both GR and GE at a low computational cost and where the results are free from spatio-temporal artifacts. In the next section we will start by briefly mentioning some facts and models from the vision science literature that form the basis of our GM framework, that is going to be introduced in Section 4.

3 SOME VISION FACTS AND MODELS FOR GAMUT MAPPING

Light reaching the retina is transformed into electrical signals by photoreceptors, rods and cones. At photopic light levels, rods are saturated and the visual information comes from

cones, of which there are three types, according to the wavelengths they are most sensitive to: L (long), M (medium), and S (short). The response of all photoreceptors is non-linear and, for a single cell without feedback, can be well approximated by the Naka-Rushton equation [57], which is a particular instance of a divisive normalization operation [58], i.e. a process that computes the ratio between the response of an individual neuron and some weighted average of the activity of its neighbors, and this in turns allows the photoreceptor response to adapt to the average light level therefore optimizing its operative range.

Photoreceptors do not operate individually though, they receive negative (inhibitory) feedback from horizontal cells, which receive excitatory input from cones and generate inhibitory input to cones. Cone output goes to bipolar cells, that also receive lateral inhibition from horizontal cells and from another type of retinal neurons called amacrine cells. Bipolars feed into retinal ganglion cells (RGCs), which also receive input from amacrine cells, and the axons of the ganglion cells form the optic nerve, sending visual signals to the lateral geniculate nucleus (LGN) in the thalamus, where the signals are re-organized into different layers each projecting to a specific layer in the cortex. There are numerous axons providing feedback from the cortex to the LGN, but their influence on color vision is not known [59].

The lateral inhibition or center-surround processing, in which a cell's output corresponds to the difference between the activity of the cell's closest neighbors and the activity of the cells in the near (and possibly far) surround, allows to encode and enhance contrast therefore being key for efficient representation, and is present at every stage of visual processing from the retina to the cortex. The size of the receptive field (the visual region to which a neuron is sensitive to) tends to increase as we progress down the visual pathway. Lateral inhibition is often modeled as a linear operation, a convolution with a kernel shaped as a difference of Gaussians (DoG). In recent studies, the surround receptive field of RGCs is modeled as a sum of Gaussians [60]. RGCs produce an achromatic signal by combining information coming from the three cone types (L+M+S), and produce chromatic opponent signals by performing center-surround processing on signals coming from cones of different types: (L+M)-S roughly corresponds to "Yellow - Blue" opponency, and L-M to "Red - Green". Achromatic and color-opponent signals are kept separate in the LGN and onto the cortex.

There are two types of bipolars, one that is excited by light increments but does not respond to decrements, and the other that responds only to light decrements; they are organized in parallel channels that separately transmit lightness and darkness, and that are maintained separate from the retina to the cortex throughout the whole visual pathway.

In the vision science literature the response of a cell is often (but not exclusively) modeled as a linear operation (weighted summation of the neighbors' activity, e.g. for lateral inhibition) followed by a non-linear operation (e.g. rectification, so as to consider only increments or decrements, but not both). For the linear part, DoG filters and oriented DoG (ODOG) filters are useful in predicting many perceptual phenomena [61], while common models for the non-linear part include rectification, divisive normalization and power-laws. For instance, non-linear photoreceptor response

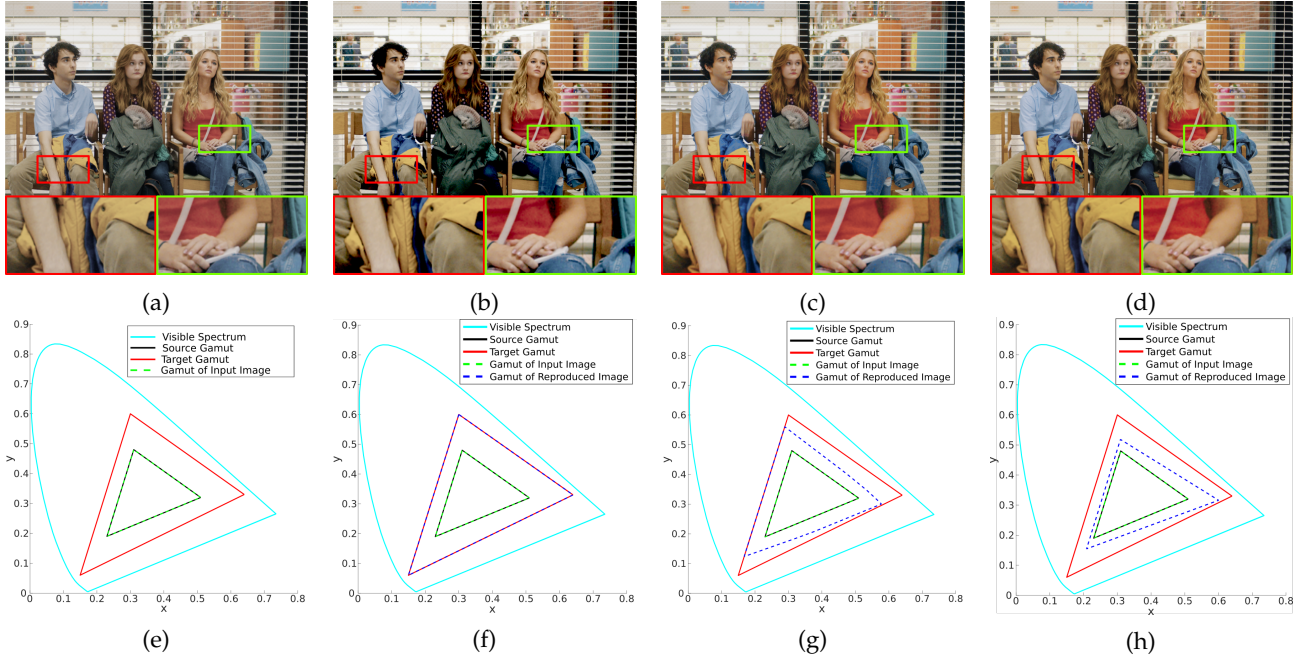


Fig. 2: Contrast enhancement produces gamut extension. Top row: (a) input image, (b) enhancing the contrast of all channels in RGB [38], (c) enhancing the contrast of chroma in CIELAB [56], (d) enhancing the contrast of saturation in HSV [55]. Bottom row: corresponding gamuts in CIE diagram. Note that the source gamut (black) and target gamut (red) are fixed as they correspond to the gamut of the display devices.

followed by linear filtering produces bandpass contrast enhancement that correlates with the contrast sensitivity function (CSF) of the human visual system [62].

The purity of a color is represented by its saturation S that expresses the amount of white that the color has: an achromatic color has $S = 0$, and blood red color has the same hue as pink, but higher saturation. The value of S can be computed from a combination of the achromatic and the chromatic signals. There is evidence in region V1 of the visual cortex, but not in the retina nor LGN, for cells tuned to S and for neural activity correlated with S [63], with a possible neural mechanism to this effect proposed in [64].

Finally, and very importantly for the GM problem, we shall mention the so-called Helmholtz-Kohlrausch effect, that implies that brightness perception depends on luminance and chrominance: color patches of the same luminance but different hue appear to have different brightness, as well as color patches of the same luminance and hue but different saturation (the higher the saturation, the higher the perceived brightness). As a consequence, if we were to modify the saturation of a color while preserving all other attributes, its brightness would appear to change. The interaction between achromatic and chromatic signals that produces brightness perception has been shown to happen at V1 [65], not before; there are some models for this, e.g. [66], [67], with a review in [68].

4 PROPOSED GAMUT MAPPING FRAMEWORK

In this section we first describe the basic functionality of our gamut extension and gamut reduction methods, and later provide implementation details in Section 4.3.

4.1 Gamut extension

In previous works we have shown how a contrast enhancement method, implemented as a partial differential equation that minimizes a certain energy, produces gamut extension when applied independently to the R, G and B channels of an input image [38], but also when applied just on the color opponent channels [56] or just on the saturation channel [55]. See Fig. 2 for an illustration.

Based on this gamut extension ability that contrast enhancement has, and considering some of the vision models enumerated in the previous section, we now propose the following basic gamut extension method: perform contrast enhancement on the saturation channel by center-surround filtering (using a model of lateral inhibition), followed by rectification so as to ensure that the saturation does not decrease (based on a model of nonlinear processing by single neurons and the existence of ON and OFF pathways), and finally modify the brightness to account for the Helmholtz-Kohlrausch (H-K) effect (using a modified model of brightness based on neural activity data).

Basic GE method:

- 1) The inputs are an image sequence, whose gamut we want to extend, and the specifications of the source and target gamuts.
- 2) Convert each input frame into HSV color space (see *appendix*). We will keep H constant.
- 3) Using the specifications of source and target gamuts, define a linear filter K_e similar to a DoG. This filter is then convolved with S , obtaining S_1 which is contrast-enhanced. Figure 3a (left) shows an example K_e filter.

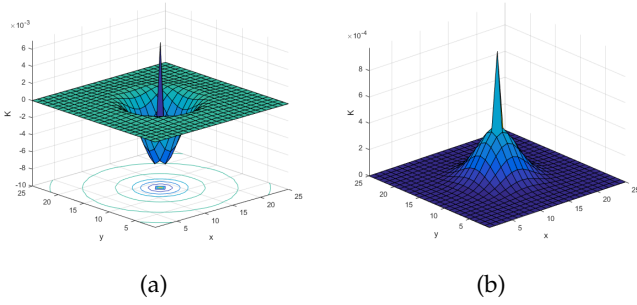


Fig. 3: Examples of kernel used in our framework: (a) for gamut extension. (b) for gamut reduction.

- 4) Add a constant value image C to S_1 , obtaining S_2 . This step attempts to preserve the mean of the original image.
- 5) Rectify $(S_2 - S)$ and produce $S_3 = S + \text{rectified}(S_2 - S)$, where $\text{rectified}(S_2 - S) = \max(S_2 - S, 0)$. This ensures that $S_3(x) \geq S(x)$ for each and every pixel x , i.e. that the process increases the saturation for all pixels with respect to their value in the original image.
- 6) Modify V to compensate for the Helmholtz-Kohlrausch effect, correcting V so that perceived brightness does not change for those colors whose saturation has been modified. This is done using a simplified version of the model by Pridmore [67] that yields $V_1 = V \left(\frac{S}{S_3}\right)^\rho$.
- 7) The final result is the image with channels (H, S_3, V_1) .

See Fig. 5 comparing the original image (left) with the intermediate result replacing S with S_3 (middle) and the final result with both S_3 and V_1 (right).

As an enhancement to the method we can add a logistic function τ after step (4) that linearly combines S_2 with the original S , giving more importance to S in the case of low-saturated values so as to preserve skin tones and other memory colors: $S_2' = (1 - \tau(S))S_2 + \tau(S)S$. The shape of function $\tau(S)$ is shown in Fig. 4 and the formula is:

$$\tau(S(x)) = 1 - \frac{1}{(1 + 0.55e^{-1.74S(x)})^2} \quad (1)$$

The values used in Eq. (1) have been chosen in the following manner. We first collect a dataset of several images containing skin tones, less saturated natural objects and memory colors. Next we perform gamut extension on these images. Finally, we empirically search for the parameter values in Eq. (1) such that the final reproduced images look natural and pleasant. Note that none of the images from this dataset were included in the final evaluation of GEAs.

4.2 Gamut reduction

Essentially, gamut extension can be seen as the inverse of the gamut reduction problem [10]. Since GE can be achieved by contrast enhancement, GR can be obtained by decreasing contrast, as we proved in [38]. In [69] Kim *et al.* showed that convolution with K_e minimizes a functional that has




Fig. 4: Logistic function used to give weights to each pixel of the input image.

a term for contrast enhancement; if we change the sign of this term, then the minimization of the functional performs contrast decrease and the solution is achieved by convolution with a new kernel K_r . Figure 3b (right) shows an example K_r filter. Following the idea presented in Section 4.1, where convolution of S with some kernel K_e yields GE, then GR could be performed by convolution with a kernel K_r that is the inverse (in Fourier space) of a kernel that would perform GE.

Basic GR method:

- 1) The inputs are an image sequence, whose gamut we want to reduce, and the specifications of the source and target gamuts.
- 2) Convert each input frame into *HSV* color space. We will keep H constant.
- 3) Use a linear filter K_r , similar to a sum of Gaussians, to convolve with S , obtaining S_1 which is contrast-decreased.
- 4) Add a constant value image C to S_1 , obtaining S_2 . This step attempts to preserve the mean of the original image.
- 5) Rectify $(S - S_2)$ and produce $S_3 = S - \text{rectified}(S - S_2)$, where $\text{rectified}(S - S_2) = \max(S - S_2, 0)$. This ensures that $S_3(x) \leq S(x)$ for each and every pixel x , i.e. that the process decreases the saturation for all pixels with respect to their value in the original image.
- 6) Modify V to compensate for the Helmholtz-Kohlrausch effect: $V_1 = V \left(\frac{S}{S_3}\right)^\rho$.
- 7) The final result is the image with channels (H, S_3, V_1) .

See Fig. 6 comparing the original image (left) with the intermediate result replacing S with S_3 (middle) and the final result with both S_3 and V_1 (right).

While one pass of this basic method already performs GR, we have found that it gives better results to iterate steps (2) to (4) with a sequence of filters K_r of progressively larger spatial extent, keeping fixed after each iteration all pixels whose colors have become in-gamut. Figure 7 shows the evolution of the filters, the image gamut and the image.

4.3 Implementation details

4.3.1 Computation of the convolution kernel

The kernel K_e for GE is computed as

$$K_e = \mathcal{F}^{-1} \left(\frac{1}{1 - \gamma \left(\frac{19}{20} - \mathcal{F}(\omega) \right)} \right), \quad (2)$$




Fig. 5: Comparison of gamut extension results: (a) input image, (b) extended-gamut image ignoring the H-K effect, (c) extended-gamut image considering the H-K effect.




Fig. 6: Comparison of gamut reduction results: (a) input image, (b) reduced-gamut image ignoring the H-K effect, (c) reduced-gamut image considering the H-K effect.

where \mathcal{F} denotes the Fourier transform, ω is a normalized 2D Gaussian kernel and γ is a positive constant. As mentioned above, the shape of K_e is similar to a difference of Gaussians, see an example in Fig. 3 (left).

The kernel K_r for GR is computed as

$$K_r = \mathcal{F}^{-1} \left(\frac{1}{1 - \gamma \left(\frac{21}{20} - \mathcal{F}(\omega) \right)} \right). \quad (3)$$

where γ is in this case a negative constant. The shape of K_r is similar to a sum of Gaussians, see an example in Fig. 3 (right). The motivation for the form of K_e and K_r kernels is given in the *appendix*.

4.3.2 Computation of optimal γ value

The basic GR method that we have proposed in Section 4.2 is already capable of mapping the colors of a wide-gamut image to a small destination gamut. However, we observed that the same method yields better results if used iteratively in the following manner. At iteration level one we apply steps (2)-(4) of the basic GR method with $\gamma = 0$. This will provide us back with the original image for which we check if there are some pixels that lie inside the destination gamut. If there are, we mark these pixels as a part of the final reduced-gamut image and these values will not be modified in subsequent iterations. We move to the next iteration and apply again steps (2)-(4) of the basic GR method but now with a slightly decreased γ value (for example, setting $\gamma = -0.05$) for the kernel in Eq. (3), and then check whether any of the pixels

that were outside the gamut at the previous iteration are now moved inside the destination gamut: we select those pixels for the final image and leave them untouched for the following iterations. We keep repeating this process until all the out-of-gamut colors are mapped inside the destination gamut. This iterative procedure implies that in the case of gamut reduction there will be a unique γ value that will provide us with an optimal reduced-gamut image. This γ value is the one that is just sufficient enough to bring all the out-of-gamut colors to inside the target gamut.

In the case of gamut extension, the gamut of the optimal result (in terms of appearance) usually does not extend to the target gamut boundaries, it lies somewhere in between the source and the target gamuts. This implies that for an input image there can be many γ values that we can use to produce corresponding extended-gamut images (some over-enhanced, some under-enhanced) and then we need to choose the one which is optimal. However, selecting the optimal image would require visual inspection. To address this issue, we present an automatic procedure to find a good-performing value for γ that allows our GEA to adapt itself for any combination of source and destination gamut according to the content of the input image.

Given a pair of source and destination gamuts (3-primaries triangles), we compute the area of the source gamut (SG_{area}) and the destination gamut (DG_{area}), and find the difference between them as:

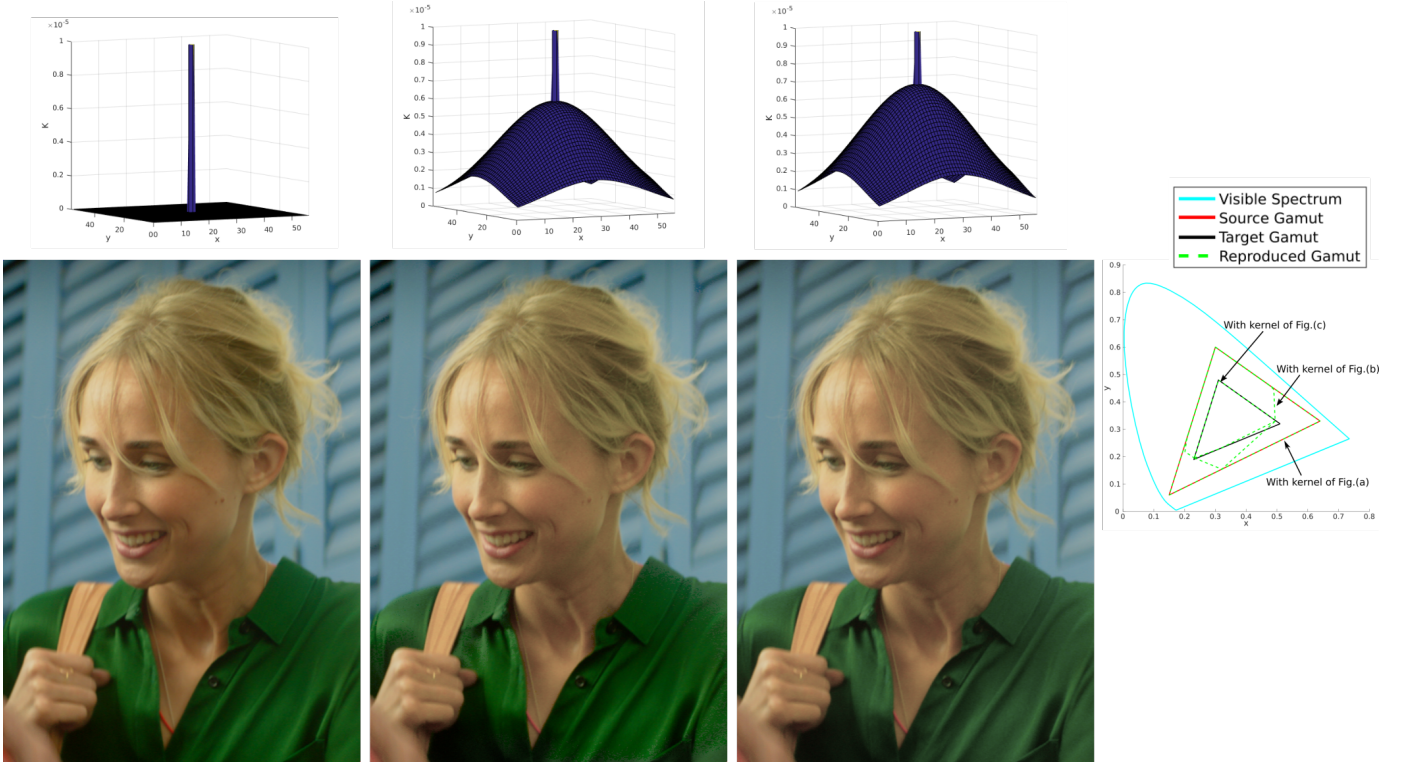


Fig. 7: Effect of increasing kernel size on image gamut. Row 1: example of kernels with progressively larger spatial extent. Row 2: reduced-gamut images corresponding to each kernel. Last column: evolution of image gamut that progressively decreases with an increase in the spatial extent of the kernel.

$$d_g = |SG_{area} - DG_{area}| \quad (4)$$

Considering the color differences between the source and target gamuts (d_g), we define γ_{base}

$$\gamma_{base} = \sqrt[3]{d_g}, \quad (5)$$

where the cube root function used in Eq. (5) has been chosen based on tests we performed for several combinations of source and destination gamuts, and it provides us with an initial γ_{base} value.

For each image, we created gamut extended images using several different values of γ . Then by subjectively comparing these reproduced images with the wide-gamut reference images, we manually found the most faithful (optimal) reproduction and therefore its corresponding γ value. While analyzing histograms of the saturation component of various optimal images, we observed a trend that our GEA requires a small γ value if the input image has a large percentage of low-saturated pixels. Whereas a large γ value is needed for the input image that has a large percentage of high-saturated pixels. Thus, we modify the γ_{base} value to obtain the optimal γ value as:

$$\gamma = \gamma_{base} + (T_S - P_{LS})\gamma_{base} \quad (6)$$

where T_S is a threshold to define the saturation level below which all the pixels will be considered as low-saturated colors, and P_{LS} denotes the percentage of low-saturated pixels in the input image. In this paper, considering that the

saturation channel of the input images are in the range $[0,1]$, we use $T_S = 0.3$. The images that we use for the computation of T_S are not included in the final evaluation of GEAs.

4.3.3 Temporal aspect

Both for GR and GE our method is applied independently to each frame of the video input. We have not needed to impose any sort of temporal consistency to our algorithm and this is due to the effectively large size of the kernels we use, which remove or strongly reduce the influence of sudden changes and make the results stable and the framework very robust.

4.3.4 The HDR case

The input is assumed to be of standard dynamic range (SDR), as it is done in the GM literature. This assumption would correspond (going back to the vision fundamentals mentioned in Section 3) to having as input image for our GM framework the signal generated by the photoreceptors. The reason for this is the well known fact that the Naka-Rushton equation that models photoreceptor responses optimizes the performance efficiency of cones and rods by adapting the possibly high dynamic range (HDR) input intensities to the SDR representation capabilities of photoreceptors [70]. In fact several successful tone mapping approaches (that convert HDR images into SDR) in the computer graphics and image processing communities use non-linear curves based on the Naka-Rushton model (e.g. [71], [72]).

Therefore, if the input video to be gamut-mapped is in HDR, our framework requires that it is tone-mapped first




Fig. 8: Wide-gamut test images used in the evaluation of GEAs. First three original images are from [76], images 4-7 are from [77], images 8-18 were captured by the authors and the rest of the images are from mainstream movies.

and then processed with our GM algorithm. This is consistent with the workflow for GM of HDR content proposed in [73].

The output of our GM method will also be in SDR form. If it were required for it to be in HDR, then an inverse tone mapping method should be applied to the output, preferably respecting the artistic intent present in the material as in the case of [74].

5 PSYCHOPHYSICAL EVALUATION

The goal of gamut mapping for cinema is to develop GMAs that reproduce content respecting as much as possible the artist’s vision, because it is an important feature that a GMA should have in order to be adopted by the movie industry. This could be achieved by including the reference images in the psychophysical tests that act as a stand-in for the content creator’s intent. Therefore we conduct psychophysical experiments in order to compare the performance of the proposed GMAs with other methods in cinematic conditions using a digital cinema projector (Barco-DP-1200 [75]) and a large projection screen.

5.1 Viewing Conditions and Evaluation Protocol

To emulate real cinema-like conditions, we used a large hall with the ambient illuminance of 1 lux and the illumination measured at the screen was around 750 lux. During the experiments there was not any strong colored object present in the observers’ field of view. We used a glare-free screen that was 3 meters wide and 2 meters high. Each observer was instructed to sit approximately 5 meters away from the screen.

In this study, we used a forced-choice pairwise comparison technique to gather raw experiment data (independently) for both gamut reduction and gamut extension problems. Each observer was shown three images simultaneously on the projection screen: the reference image (in the middle) and a pair of reproductions (one image on the left side and the other on the right side of the reference image). We asked each observer to make selections according to these instructions: a) if there are any sort of artifacts in one of the reproductions, choose the other, and b) if both of the reproductions have artifacts or are free from artifacts, choose the one which is perceptually closer to the reference image. In pair comparison




Fig. 9: Wide-gamut test images used in the evaluation of GRAs. First five original images are from [77], the last four were captured by the authors and the rest of the images are from mainstream movies.

evaluation [21], in order to calculate differences among n chosen GMAs, observers need to compare $n(n-1)/2$ number of pairs for each test image. For a given pair of reproductions, a score of 1 is given to the reproduction which is selected by an observer, and a score to 0 to the other reproduction. For each test image, the responses of an observer were stored in a $n \times n$ raw matrix where the value in column i and row j denote the score given to GMA i as compared with GMA j .

To compute accuracy scores from the raw psychophysical data, we use the same approach as in the work of Morovič [21] (see Chapter 5 of his thesis), that is based on Thurstone’s law of comparative judgment [78].

Finally, a corpus of 15 observers participated in each of the experiments we performed in our lab, as this is the number of observers for pair comparison tests that is suggested by several technical recommendation documents (e.g. [79], [80]). All the observers (12 male and 3 female with ages in the range of 23 to 36 years) had normal color vision as tested with the Ishihara’s test of color deficiency.

5.2 Image media

The DCI-P3 wide-gamut test images that we used in evaluating GEAs and GRAs are, respectively, shown in Fig. 8 and Fig. 9. (Note that these are sRGB images because we are limited to the sRGB standard to show results on paper.) Some of these test images were taken from the publicly available datasets [77], [76], while others were from [56] and from mainstream feature films.

5.3 Experiment 1: Evaluation of GEAs

In the case of the psychophysical evaluation of GEAs, the first step is to create limited-gamut input images. This is achieved by applying a clipping operation on the DCI-P3 reference images in order to map the out-of-gamut colors to the boundary of the BT.709 gamut (or any other desired gamut). In order to perform clipping we used the xyY color space, and clip chromaticities of the out-of-gamut colors of a given image to the boundary of the destination gamut towards a focal point that is the white point ‘D65’. The experimental gamuts that we use in this paper are depicted in Fig. 10, and their primaries are listed in Table 1. The procedure for computing the γ value that we use for the kernel in Eq. (2) is presented in Section 4.3.2, and the standard deviation for ω is set to one-third of the number of rows or columns of the input image (whichever is larger). In the




Fig. 10: Gamuts on chromaticity diagram.

formula of brightness modification in step 5 of the basic GE (see Section 4.1), we set $\rho = 0.35$.

5.3.1 Experimental Setups for GE

We have defined the following two different experimental setups for the evaluation of GEAs.

- 1) *Setup 1: Mapping from small gamut to DCI-P3 gamut.* As quantum dot displays [81] and laser projectors [82] with their extended gamut capabilities are becoming popular, in the near future the common case will be to have large color differences between the standard gamut and the gamut of the display. Therefore, this setup is created to investigate how different GEAs will perform when the difference between source and target gamuts is large. To this end, we map the source images from the small ‘Toy’ gamut (slightly smaller than the BT.709 gamut) to the large DCI-P3 gamut. On the chromaticity diagram, the difference in gamuts for this setup is nearly equal to the difference between BT.709 and BT.2020. This represents the future scenario where we need to show on a wide-gamut display some content that was mastered for TV.
- 2) *Setup 2: Mapping from BT.709 to DCI-P3 gamut.* In this setup we mimic the practical situation where the source material has BT.709 gamut and we map the source colors to the colors of the DCI-P3 gamut.

5.3.2 Competing GEAs

For each set-up, we compare the proposed method with the top-ranked GEAs in a recent work [56]). These GEAs are briefly explained as follows.

- **Same Drive Signal (SDS)** method linearly maps the RGB primaries of the source gamut to the RGB primaries of the destination device gamut, therefore making the full use of the gamut of the target display.
- **Hybrid Color Mapping (HCM)** is a combination of the SDS algorithm and the true-color algorithm. The true-color algorithm represents the input image in the target gamut without applying any extension. The HCM algorithm [50] analyzes the saturation of the input image and then linearly combines the output of the true-color method and the SDS method:

TABLE 1: Primaries of gamuts.

Gamuts	Red Primaries		Green Primaries		Blue Primaries	
	x	y	x	y	x	y
BT.2020	0.708	0.292	0.170	0.797	0.131	0.046
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
Toy	0.570	0.320	0.300	0.530	0.190	0.130




Fig. 11: Accuracy scores of competing GEAs: 15 observers took part in each experiment and 30 images were used.

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix}_{HCM} = (1-\kappa) \begin{bmatrix} R \\ G \\ B \end{bmatrix}_{true-color} + \kappa \begin{bmatrix} R \\ G \\ B \end{bmatrix}_{SDS}^{(7)}$$

where κ is a mixing factor that works as a function of saturation:

$$\kappa(S) = \begin{cases} 0, & \text{if } S \leq S_L \\ \frac{S-S_L}{S_H-S_L}, & \text{if } S_L < S < S_H \\ 1, & \text{if } S \geq S_H \end{cases} \quad (8)$$

S_L and S_H are constants to define the ranges of saturation for the mixing function κ , and their values that we used in our experiments are $S_L = 0.8$ and $S_H = 1$ as defined in [56].

The method of HCM aims at preserving natural colors by leaving unchanged the low-saturated colors such as flesh tones, while mapping the high-saturated colors using the SDS method.

- **GEA of Zamir et al.** [56] is a spatially-variant GEA, implemented as a PDE-based optimization procedure, that performs gamut extension in CIELAB color space by taking into account the analysis of distortions in hue, chroma and saturation.

5.3.3 Results of GEAs Under Experimental Setup 1 and Setup 2

Once the reproductions were obtained by applying GEAs on the input images of both setups, we conducted a psychophysical evaluation separately for each setup using the 15 observers mentioned in Section 5.1.

Fig. 11a presents the accuracy scores computed by analyzing the psychophysical data of the setup 1 where it can be seen that, when the difference between the source gamut and the destination gamut is large, the proposed GEA yields images that are perceptually more faithful to the reference

images than the other competing algorithms. The observers declared SDS [50] as the least accurate method, whereas the algorithm of [56] ranked second.

In Fig. 11b we present results for the setup 2 where it can be seen that, when the color difference between the source-destination gamut pair is small, our algorithm ranks first, followed by the HCM algorithm [50] and the method of [56].

5.4 Experiment 2: Evaluation of GRAs

This section is devoted to examining the image reproduction quality of competing GRAs. To obtain the reduced-gamut images, we apply the proposed GRA on the saturation channel of the input images by using the proposed GRA in an iterative manner (over the γ parameter) as described in Section 4.3.2. The other parameter for the kernel in Eq. (3) is the standard deviation for ω , which is equal to one-twentieth of the number of rows or columns of the input image (whichever is larger). In the formula of brightness modification in step 5 of the basic GR (see Section 4.2), we set $\rho = 0.20$.

5.4.1 Experimental Setups for GR

All the competing GRAs receive as input the wide-gamut DCI-P3 images and generate reproductions for the following two different experimental setups.

- 1) *Setup 1: Mapping from DCI-P3 gamut to a small gamut.* We created this particular setup with a large difference between source and target gamuts, nearly as large as it is between BT.2020 and BT.709 gamuts. An experimental setup with such large difference in gamuts allows us to not only evaluate the performance of competing GRAs reliably but also provides us an indication of how these GRAs might perform when BT.2020 content becomes commonly available and needs to be mapped to BT.709 displays or DCI-P3 cinema projectors. To compute the results using the competing GRAs, we map the colors of the 15 DCI-P3 test images shown in Fig. 9 to the challenging smaller ‘Toy’ gamut.
- 2) *Setup 2: Mapping from DCI-P3 to BT.709 gamut.* Colorists perform this gamut reduction procedure by using 3D LUTs (as we mentioned in more detail in the introduction.) Therefore, we engaged a professional colorist from a post-production company to use their own in-house 3D LUTs and apply them on our DCI-P3 test images in order to create reduced-gamut BT.709 images. We also perform GR using the following competing GRAs.

5.4.2 Competing GRAs

- **LCLIP** [15] clips the chroma of the out-of-gamut colors to the destination gamut boundary along lines of constant hue and lightness.
- **Hue Preserving Minimum ΔE (HPMINDE)** [16] involves clipping of the out-of-gamut color to the closest color, in terms of ΔE error, on the boundary of the destination gamut along lines of constant hue.
- **Alsam and Farup** [29] proposed an iterative GRA that at iteration level zero behaves as a gamut clipping




Fig. 12: Accuracy scores of competing GRAs: 15 observers took part in each experiment and 15 images were used.




Fig. 13: Accuracy scores of competing GRAs for the skilled technicians case. The experiment was performed by 6 experts, and 15 images were used.

algorithm, and as the number of iterations increases the solution approaches spatial gamut mapping.

- **Schweiger et al.** [25] make use of a compression function that squeezes colors near the destination gamut boundary in order to accommodate the out-of-gamut colors. This is a method proposed and used by the British Broadcasting Corporation (BBC).

5.4.3 Results of GRAs Under Experimental Setup 1 and Setup 2

The 15 observers that took part in the GR experiment for setup 1 were the same observers that participated in the evaluation of the GE algorithms. The analysis of psychophysical data gathered for GRAs is presented in Fig. 12a. It can be seen in this figure that the proposed GRA produces images that are perceptually more faithful to the original images than any other competing method. It is evident from Fig. 12a that observers did not prefer the HPMINDE algorithm in most of the tests images, and therefore rated it as the least accurate method. The algorithms of Schweiger et al. [25], Alsam and Farup [25] and LCLIP [15] are, respectively, ranked second, third and fourth by the observers.

For the experimental setup 2 we also ran the psychophysical tests with 15 observers, of which 9 had experience in image processing and the other 6 were *skilled* technicians (colorists and editors) from a post-production company.

In order to reduce the number of pair comparisons, in this particular setup we opted to use the reproduced images of the top three ranked methods from setup 1 and the reduced-gamut images created by using the custom LUT of the same post-production company. Figure 12b shows the result for all the observers. Observers preferred the in-house LUT results over the other methods, with our GRA being ranked second.

More specifically, we can focus our attention on the result for this experiment when considering only the 6 skilled

	GE_S1			GE_S2			HPMINDE	LCLIP	GR_S1			GR_S2			3D LUT
	SDS	Zamir et al.	Ours	Zamir et al.	HCM	Ours			Alsam et al.	Schweiger et al.	Ours	Alsam et al.	Schweiger et al.	Ours	
CID	-0,209	-0,372	0,582	-0,465	0,279	0,186	0,102	-0,686	-0,279	0,610	0,254	-0,627	0,593	0,356	-0,322
iCID	-0,279	-0,279	0,559	-0,419	0,140	0,279	-0,178	-0,610	-0,254	0,584	0,457	-0,559	0,322	0,457	-0,220
Delta E00	-0,349	-0,279	0,628	-0,326	0,163	0,163	0,559	-0,279	-0,635	-0,152	0,508	-0,491	0,254	0,728	-0,491
LMSE	-0,652	0,303	0,349	0,465	-0,396	-0,070	0,127	-0,711	-0,152	0,686	0,051	-0,423	0,627	-0,119	-0,085
SC	-0,349	0,628	-0,279	0,582	-0,047	-0,535	-0,356	0,686	-0,711	0,102	0,279	-0,288	0,254	0,728	-0,694
NAE	-0,372	-0,116	0,489	-0,116	0,070	0,047	0,406	-0,610	-0,356	0,406	0,152	-0,593	0,356	0,525	-0,288
PSNR	0,396	0,023	-0,419	0,070	-0,093	0,023	-0,279	0,737	0,025	-0,559	0,076	0,593	-0,525	-0,186	0,119
AD	-0,559	0,675	-0,116	0,628	-0,326	-0,303	-0,356	0,711	-0,584	-0,025	0,254	-0,085	0,051	0,728	-0,694
PieApp	0,140	-0,303	0,163	-0,396	0,326	0,070	0,051	0,229	-0,711	0,051	0,381	-0,728	0,288	0,627	-0,186
LPIPS	-0,209	-0,233	0,442	-0,116	0,163	-0,047	-0,051	-0,610	-0,152	0,584	0,229	-0,322	0,457	0,152	-0,288
Observers' data	-0,525	-0,121	0,647	-0,107	-0,081	0,178	-0,812	-0,343	-0,073	0,516	0,712	-0,733	0,162	0,213	0,358

Fig. 14: Comparison between the results of different image metrics and the results from psychophysical evaluation. Metrics were considered as observers in a pair comparison experiment. Each experiment is color coded individually. Color codes are green for the best result and red for the worst one.

technicians. This result is shown in Fig. 13. In this case, we can see that the trend is very similar to the one obtained by the 15 observers (the ranking of the algorithms is not modified), but also that the skilled technicians are more inclined to select the in-house LUT of their post-production company, probably because they might be more inclined to select the solution they are used to working with. Also, the use of a LUT is well suited for the case of DCI-P3 to BT.709 reduction, where the blue primary is essentially the same for both gamuts, and the differences in the other two primaries are rather small: however, for larger gamut differences, the LUT approach might be hard to generalize.

5.5 Video Results

In order to test the temporal coherence we apply the proposed gamut reduction and gamut extension methods to all frames of videos independently. We confirm that the results produced by our algorithms are free from artifacts. The videos are available at <http://ip4ec.upf.edu/GamutMappingSubmission>.

6 DOES ANY ERROR METRIC APPROXIMATE OUR PSYCHOPHYSICAL RESULTS?

In this section we evaluate if there exists any image metric able to predict the result of our psychophysical test, following the same strategy we used for the GE case in [56]. To this end we consider a total of 10 metrics: a perceptually-based color image difference (CID) metric [83] particularly tailored for the gamut mapping problem, its more recent extension iCID [36], CIE Δ E00 [84], the metrics presented in [85] such as Laplacian mean square error (LMSE), structural content (SC), normalized absolute error (NAE), peak signal-to-noise ratio (PSNR), and absolute difference (AD), and finally two very recent, state-of-the-art deep learning metrics for perceived appearance, PieAPP [86] and LPIPS [87], learned from human judgements on large image databases. All these metrics are full-reference, and therefore, they have access to the reference images, as do the observers in our experiments.

In order to perform a fair comparison to our experiment, we consider the metrics as if they were observers in our pair comparison test. This means that, for each metric, we will run all the possible comparisons, and in each comparison we will give a 1 to the image with better metric value and a 0 to

the image with worse metric value. Later, we will apply the Thurstone Case V analysis [78] to each of the image metrics to end up with the preference values for each of the methods. These preference values will therefore be comparable to the ones shown for the psychophysical analysis in our previous section. For readers interested in the exact numerical values of the metrics (e.g. the mean value for each method, etc.), we provide them in the supplementary material.

Figure 14 shows the result of the aforementioned analysis. Each of the experimental setups is individually colored with a color code where the hue goes from pure red for the lowest value to pure green for the highest one. Therefore, for any metric to be able to predict the psychophysical results, its color code should match that of the results of the observers, shown in the last row. We can see that there is only one specific case where we could argue that this is happening, the NAE metric in the first setup of gamut extension. However, the same metric is not able to predict the observers' response in any of the other three cases.

It is interesting to mention that the CID and iCID metrics, which were specifically developed for gamut mapping, do not match the observer data; one possible reason is that these metrics were designed for input images in the BT.709 gamut, while in this paper the input images are in DCI-P3, which may explain the limitations of CID and iCID in the context of our problem.

Another significant result is that two state-of-the-art deep learning metrics, PieAPP [86] and LPIPS [87], designed to predict perceptual image error like human observers and based on large scale datasets (20K images in one case, 160K images in the other) labeled with pair-comparison preferences and, in the case of [87], using close to 500K human judgements, are not able to predict the observers' preference in any of the experimental set-ups that we have tested. This result is important, as it suggests that current deep learning approaches are not accurate enough for validating (and therefore developing) GM methods for cinema applications, although it's not a surprising result in the sense that several very recent works have also shown how large or very large scale image databases (250,000+ images) can be used to train deep neural networks to predict user preference but whose performance decays remarkably when used on images that belong to some other dataset different from the one used for training [87], [88].

In summary, there does not seem to be an adequate metric that is able to predict the preference of observers for GM results. This has two important consequences. First, that in order to evaluate gamut mapping methods, we still need to rely only on psychophysical studies. However, conducting subjective studies is operationally difficult, hard to replicate from experiment to experiment, economically costly and may require special equipment (cinema projector, large screen, etc.). Second, that we cannot develop GM methods or optimize GM results simply by maximizing an image appearance metric or minimizing an error metric (as it is done in other contexts, e.g. [89]), which of course would be extremely practical. Therefore, the results presented in this section point out the importance of working towards defining better metrics for the gamut mapping problem that are able to predict the observers' preference, which we strongly believe would be of great importance for the color imaging community.

7 CONCLUSION

We have introduced a GM framework based on some basic vision science facts and models. The algorithms for GE and GR are simple, of low computational complexity, produce results without artifacts and outperform state-of-the-art methods according to psychophysical experiments.

We also tested a number of established and state-of-the-art objective metrics on the results produced by the GM methods we have compared, and we have observed that these metrics correlate poorly with the choices made by the observers. Therefore there is a need for developing an image quality metric for GM: this would be a very significant contribution, as it would greatly simplify the validation of new methods and would also allow to optimize GM results by optimizing said metric.

APPENDIX A

In this Appendix we will give the motivation for the shape of the kernels that we use for gamut reduction and extension, presented in Section 4.3.1.

In order to map colors from one gamut to another, Zamir *et al.* [38] made use of an image energy functional related with visual perception and adapted from the work of Bertalmío *et al.* [90]

$$E(I) = \frac{\alpha}{2} \sum_x (I(x) - \mu)^2 + \frac{\beta}{2} \sum_x (I(x) - I_0(x))^2 - \frac{\gamma}{2} \sum_x \sum_y w(x, y) |I(x) - I(y)|, \quad (9)$$

where α , β and γ are constant weights, I is a color channel (R , G or B) in the range $[0, 1]$, I_0 is the original image channel, μ is the average value of an image channel, $w(x, y)$ is a normalized Gaussian kernel, and x and y are pixel locations. The minimization of the image energy (9) yields either an extended-gamut image or a reduced-gamut image depending on the sign of the parameter γ : a positive value of γ leads to contrast enhancement and gamut extension, and a negative value of γ performs contrast reduction and also gamut reduction. However, minimizing Eq. (2) requires to

perform an iterative procedure like gradient descent to obtain the final image, thus being a time-consuming approach.

Recently, Kim *et al.* [69] proposed to replace the absolute value function in the third term of Eq. (9) with a squared function

$$E(I) = \frac{\alpha}{2} \sum_x (I(x) - \mu)^2 + \frac{\beta}{2} \sum_x (I(x) - I_0(x))^2 - \frac{\gamma}{2} \sum_x \sum_y w(x, y) (I(x) - I(y))^2. \quad (10)$$

With this modification the Euler-Lagrange equation can be solved directly:

$$\delta E(I) = \alpha(I(x) - \mu) + \beta(I(x) - I_0(x)) - \gamma \sum_x w(x, y) (I(x) - I(y)) = 0. \quad (11)$$

We can see that the last term is equal to the convolution of the image I and the kernel $w(x, y)$. Rearranging, and with an abuse of notation to avoid writing the pixel location x we reach to

$$(\alpha + \beta - \gamma)I + \gamma w * I = \alpha\mu + \beta I_0 \quad (12)$$

Applying the Fourier transform we obtain

$$(\alpha + \beta - \gamma)\mathcal{F}(I) + \gamma\mathcal{F}(w)\mathcal{F}(I) = \mathcal{F}(\alpha\mu + \beta I_0) \quad (13)$$

Rearranging and applying the inverse Fourier transform we obtain the solution

$$I = \mathcal{F}^{-1} \left(\frac{\mathcal{F}(\alpha\mu + \beta I_0)}{(\alpha + \beta - \gamma) + \gamma\mathcal{F}(w)} \right) \quad (14)$$

which can be rewritten in terms of convolving a kernel with the input image I_0 as

$$I = \beta I_0 * K + C \quad (15)$$

where C is a flat, single-value image with all pixels equal to the product of α , the mean of the input image channel, and the mean of the kernel K . The kernel K is computed as

$$K = \mathcal{F}^{-1} \left(\frac{1}{(\alpha + \beta - \gamma) + \gamma\mathcal{F}(w)} \right), \quad (16)$$

The sign of γ defines the type of GM and the shape of the kernel (see Fig. 3): in the case of gamut extension γ is a positive constant and in the case of gamut reduction it is negative.

From the above, convolving with K the three channels of an RGB image would yield a gamut-mapped result. Perhaps counter-intuitively, we also obtain a gamut-mapped result if we convolve K with the saturation channel S of an image in HSV colorspace, and then rectify the output. For instance, for GE, the constant γ is positive, which means that convolution with K will increase contrast and therefore locally maximal values of S will go up, while some S values will go down, but then rectification ensures that the S values that decreased are brought back to their original value; the net result is that after convolution and rectification, S has increased, and therefore the gamut has been extended. The same reasoning can of course be applied to the gamut reduction case.

In our application, for performing GM we fix in Eq. (16) the parameter values $\beta = 1$ and $\alpha = \frac{\lfloor \gamma \rfloor}{20}$, which leads to the kernels presented in equations 2 and 3.

APPENDIX B

In this Appendix we present the standard procedure of converting between RGB and HSV color spaces [91].

RGB to HSV:

Given R , G , and B components in the range $[0,1]$, we can calculate H , S , and V components as:

$$\begin{aligned} V &= \max(R, G, B) \\ m &= \min(R, G, B) \\ H' &= \begin{cases} \text{undefined,} & \text{if } V - m = 0 \\ \frac{G-B}{V-m} \bmod 6, & \text{if } V = R \\ \frac{B-R}{V-m} + 2, & \text{if } V = G \\ \frac{R-G}{V-m} + 4, & \text{if } V = B \end{cases} \\ H &= 60^\circ \times H' \\ S &= \begin{cases} 0, & \text{if } V = 0 \\ \frac{V-m}{V}, & \text{otherwise} \end{cases} \end{aligned}$$

HSV to RGB:

Given $H \in [0^\circ, 360^\circ]$, $S \in [0, 1]$, and $V \in [0, 1]$ components, we can calculate R , G , and B as:

$$\begin{aligned} C &= V \times S \\ H' &= \frac{H}{60^\circ} \\ X &= C \times (1 - |(H' \bmod 2) - 1|) \\ (R_1, G_1, B_1) &= \begin{cases} (0, 0, 0), & \text{if } H \text{ is undefined} \\ (C, X, 0), & \text{if } 0 \leq H' \leq 1 \\ (X, C, 0), & \text{if } 1 \leq H' \leq 2 \\ (0, C, X), & \text{if } 2 \leq H' \leq 3 \\ (0, X, C), & \text{if } 3 \leq H' \leq 4 \\ (X, 0, C), & \text{if } 4 \leq H' \leq 5 \\ (C, 0, X), & \text{if } 5 \leq H' \leq 6 \end{cases} \\ (R, G, B) &= (R_1 + V - C, G_1 + V - C, B_1 + V - C) \end{aligned}$$

ACKNOWLEDGEMENTS

The authors are grateful to all the participants of the psychophysical experiments. Special thanks to Dirk Maes from Barco N.V. and to Stephane Cattan from Deluxe-Spain for their invaluable help and unwavering support. This work has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement number 761544 (project HDR4EU) and under grant agreement number 780470 (project SAUCE), and by the Spanish government and FEDER Fund, grant ref. TIN2015-71537-P (MINECO/FEDER,UE). The work of J. Vazquez-Corral was supported by the Spanish government under Grant IJCI-2014-19516.

REFERENCES

- [1] SMPTE. D-cinema quality – reference projector and environment (report 431-2:2011), 2011.
- [2] ITU-R. Recommendation BT.709-5: parameter values for the HDTV standards for production and international programme exchange, 2002.
- [3] ITU-R. Recommendation BT.2020: parameter values for ultra high definition television systems for production and international programme exchange, 2012.
- [4] M. R. Pointer. The gamut of real surface colours. *Color Research & Application*, 5(3):145–155, 1980.
- [5] M. Bertalmio. *Image Processing for Cinema*, volume 4. CRC Press, Taylor & Francis, 2014.
- [6] D. Bankston. The color-space conundrum, part one. *American Cinematographer*, page 6, 2005.
- [7] G. Kennel. *Color and mastering for digital cinema: digital cinema industry handbook series*. Taylor & Francis US, 2007.
- [8] Y. Kusakabe, Y. Iwasaki, and Y. Nishida. Wide-color-gamut super hi-vision projector. In *Proc. ITE annual convention (in Japanese)*, 2013.
- [9] B. D. Silverstein, A. F. Kurtz, J. R. Bietry, and G. E. Nothhard. A laser-based digital cinema projector. *SID Symposium Digest of Technical Papers*, 42(1):326–329, 2011.
- [10] J. Morovič. *Color gamut mapping*, volume 10. Wiley, 2008.
- [11] B. A. Olshausen. 20 years of learning about vision: Questions answered, questions unanswered, and questions not yet asked. In *20 Years of Computational Neuroscience*, pages 243–270. Springer, 2013.
- [12] N. Katoh and M. Ito. Gamut mapping for computer generated images (ii). In *Proc. of 4th IS&T/SID Color Imaging Conference*, pages 126–129, 1996.
- [13] G. Marcu and S. Abe. Gamut mapping for color simulation on CRT devices. In *Proc. of Color Imaging: Device-Independent Color, Color Hard Copy, and Graphic Arts*, 1996.
- [14] E. D Montag and M. D. Fairchild. Gamut mapping: Evaluation of chroma clipping techniques for three destination gamuts. In *Color and Imaging Conference*, pages 57–61, 1998.
- [15] J. J. Sara. *The automated reproduction of pictures with nonreproducible colors*. PhD thesis, Massachusetts Institute of Technology (MIT), 1984.
- [16] 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.
- [17] K. Masaoka, Y. Kusakabe, T. Yamashita, Y. Nishida, T. Ikeda, and M. Sugawara. Algorithm design for gamut mapping from UHDTV to HDTV. *Journal of Display Technology*, 12(7):760–769, 2016.
- [18] R. S. Gentile, E. Walowitz, and J. P. Allebach. A comparison of techniques for color gamut mismatch compensation. *Journal of Imaging Technology*, 16:176–181, 1990.
- [19] P. G. Herzog and M. Müller. Gamut mapping using an analytical color gamut representation. In *Proc. of Color Imaging: Device-Independent Color, Color Hard Copy, and Graphic Arts*, pages 117–128, 1997.
- [20] A. J. Johnson. Perceptual requirements of digital picture processing. *Paper presented at IARAIGAI symposium and printed in part in Printing World*, 1979.
- [21] J. Morovič. *To Develop a Universal Gamut Mapping Algorithm*. PhD thesis, University of Derby, UK, 1998.
- [22] UGRA. UGRA GAMCOM version 1.1: Program for the color gamut compression and for the comparison of calculated and measured values. Technical report, UGRA, St. Gallen, 17 July 1995.
- [23] J. Yuan, J. Y. Hardeberg, and G. Chen. Development and evaluation of a hybrid point-wise gamut mapping framework. In *Colour and Visual Computing Symposium*, pages 1–4, 2015.
- [24] J.A. Pytlarz, K.D. Thurston, D. Brooks, P. Boon, and R. Atkins. Real time cross-mapping of high dynamic range images. *IET Conference Proceedings*, 2016.
- [25] F. Schweiger, T. Borer, and M. Pindoria. Luminance-preserving colour conversion. In *SMPTE Annual Technical Conference and Exhibition*, pages 1–9, 2016.
- [26] 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.
- [27] J. Meyer and B. Barth. Color gamut matching for hard copy. In *Proc. of SID Digest*, pages 86–89, 1989.
- [28] M. Zhu, J. Y. Hardeberg, N. Wang, and B. Sun. Spatial gamut mapping based on guided filter. In *Proc. of IS&T/SPIE Electronic Imaging*, pages 1–4, 2016.

- [29] 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.
- [30] I. Farup, C. Gatta, and A. Rizzi. A multiscale framework for spatial gamut mapping. *IEEE Transactions on Image Processing*, 16(10):2423–2435, 2007.
- [31] C. Gatta and I. Farup. Gamut mapping in RGB colour spaces with the iterative ratios diffusion algorithm. In *Proc. of IS&T/SPIE Electronic Imaging*, pages 12–20, 2017.
- [32] J. J. McCann. A spatial colour gamut calculation to optimize colour appearance. In *Colour Image Science: Exploiting Digital Media*, pages 213–233, 2002.
- [33] 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.
- [34] 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.
- [35] S. Nakauchi, S. Hatanaka, and S. Usui. Color gamut mapping based on a perceptual image difference measure. *Color Research & Application*, 24(4):280–291, 1999.
- [36] J. Preiss, F. Fernandes, and P. Urban. Color-image quality assessment: From prediction to optimization. *IEEE Transactions on Image Processing*, 23(3):1366–1378, 2014.
- [37] J. Preiss and P. Urban. Image-difference measure optimized gamut mapping. In *Proc. of IS&T/SID 20th Color Imaging Conference*, pages 230–235, 2012.
- [38] S. W. Zamir, J. Vazquez-Corral, and M. Bertalmío. Gamut mapping in cinematography through perceptually-based contrast modification. *IEEE Journal of Selected Topics in Signal Processing*, 8(3):490–503, 2014.
- [39] G. J. Braun. *A Paradigm for Color Gamut Mapping of Pictorial Images*. PhD thesis, Rochester Institute of Technology, USA, 1999.
- [40] Y. Ling. Investigation of a gamut extension algorithm. Master's thesis, University of Derby, UK, 2001.
- [41] X. Chen. Investigation of gamut extension algorithms. Master's thesis, University of Derby, UK, 2002.
- [42] 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.
- [43] T. Hoshino. Color estimation method for expanding a color image for reproduction in a different color gamut, May 1994. US Patent 5,317,426.
- [44] B. H. Kang, J. Morović, M. R. Luo, and M. S. Cho. Gamut compression and extension algorithms based on observer experimental data. *ETRI journal*, 25(3):156–170, 2003.
- [45] 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.
- [46] H. Pan and S. Daly. A gamut-mapping algorithm with separate skin and non-skin color preference controls for wide-color-gamut TV. *SID Symposium Digest of Technical Papers*, 39(1):1363–1366, 2008.
- [47] S. E. Casella, R. L. Heckaman, and M. D. Fairchild. Mapping standard image content to wide-gamut displays. In *Color and Imaging Conference*, pages 106–111, 2008.
- [48] R. L. Heckaman and J. Sullivan. Rendering digital cinema and broadcast TV content to wide gamut display media. *SID Symposium Digest of Technical Papers*, 42(1):225–228, 2011.
- [49] 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.
- [50] J. Laird, R. Muijs, and J. Kuang. Development and evaluation of gamut extension algorithms. *Color Research & Application*, 34(6):443–451, 2009.
- [51] X. Meng, G. Song, and H. Li. A human skin-color-preserving extension algorithm for wide gamut displays. In *Proc. of International Conference on Information Technology and Software Engineering*, Lecture Notes in Electrical Engineering, pages 705–713, 2013.
- [52] G. Song, H. Cao, and H. Huang. Hue preserving multi-level expansion method based on saturation for wide gamut displays. *Journal of Information & Computational Science*, 11(2):461–472, 2014.
- [53] G. Song, X. Meng, H. Li, and Y. Han. Skin color region protect algorithm for color gamut extension. *Journal of Information & Computational Science*, 11(6):1909–1916, 2014.
- [54] Y. Li, G. Song, and H. Li. A multilevel gamut extension method for wide gamut displays. In *Proc. of International Conference on Electric Information and Control Engineering (ICEICE)*, pages 1035–1038, 2011.
- [55] S. W. Zamir, J. Vazquez-Corral, and M. Bertalmío. Perceptually-based gamut extension algorithm for emerging wide color gamut display and projection technologies. In *SMPTE Annual Technical Conference & Exhibition*, pages 1–11, 2016.
- [56] S. W. Zamir, J. Vazquez-Corral, and M. Bertalmío. Gamut extension for cinema. *IEEE Transactions on Image Processing*, 26(4):1595–1606, 2017.
- [57] R. Shapley and C. Enroth-Cugell. Visual adaptation and retinal gain controls. *Progress in retinal research*, 3:263–346, 1984.
- [58] M. Carandini and D. J. Heeger. Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, 13(1):51, 2012.
- [59] P. Gouras. Colour vision. *eLS*, 2009.
- [60] Jihyun Yeonan-Kim and Marcelo Bertalmío. Retinal lateral inhibition provides the biological basis of long-range spatial induction. *PLoS one*, 11(12):e0168963, 2016.
- [61] B. Blakeslee and M. E. McCourt. A unified theory of brightness contrast and assimilation incorporating oriented multiscale spatial filtering and contrast normalization. *Vision research*, 44(21):2483–2503, 2004.
- [62] A. B. Watson and J. A. Solomon. Model of visual contrast gain control and pattern masking. *JOSA A*, 14(9):2379–2391, 1997.
- [63] A. Hanazawa, H. Komatsu, and I. Murakami. Neural selectivity for hue and saturation of colour in the primary visual cortex of the monkey. *European Journal of Neuroscience*, 12(5):1753–1763, 2000.
- [64] V. A. Billock and B. H. Tsou. Sensory recoding via neural synchronization: integrating hue and luminance into chromatic brightness and saturation. *JOSA A*, 22(10):2289–2298, 2005.
- [65] D. Xing, A. Ouni, S. Chen, H. Sahmoud, J. Gordon, and R. Shapley. Brightness-color interactions in human early visual cortex. *Journal of Neuroscience*, 35(5):2226–2232, 2015.
- [66] M. Ayama and M. Ikeda. Brightness-to-luminance ratio of colored light in the entire chromaticity diagram. *Color Research & Application*, 23(5):274–287, 1998.
- [67] R. W. Pridmore. Chroma, chromatic luminance, and luminous reflectance. part ii: Related models of chroma, colorfulness, and brightness. *Color Research & Application*, 34(1):55–67, 2009.
- [68] M. Withouck, K. A.G. Smet, W. R. Ryckaert, M.R. Pointer, G. Deconinck, J. Koenderink, and P. Hanselaer. Brightness perception of unrelated self-luminous colors. *JOSA A*, 30(6):1248–1255, 2013.
- [69] J. Kim, T. Batard, and M. Bertalmío. Retinal processing optimizes contrast coding. *Journal of Vision*, 16(12):1151–1151, 2016.
- [70] A. Grimaldi, D. Kane, and M. Bertalmío. Statistics of natural images as a function of dynamic range. *Journal of Vision*, 2018.
- [71] E. Reinhard and K. Devlin. Dynamic range reduction inspired by photoreceptor physiology. *IEEE Transactions on Visualization & Computer Graphics*, (1):13–24, 2005.
- [72] 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*, 33(10):2002–2012, 2011.
- [73] E. Sikudova, T. Pouli, A. Artusi, A. O. Akyuz, F. Banterle, Z. M. Mazlumoglu, and E. Reinhard. A gamut-mapping framework for color-accurate reproduction of hdr images. *arXiv preprint arXiv:1711.08925*, 2017.
- [74] C. Bist, R. Cozot, G. Madec, and X. Ducloux. Style aware tone expansion for hdr displays. In *Graphics Interface*, pages 57–63, 2016.
- [75] Barco. <http://www.barco.com/en/Products-Solutions/Projectors/Digital-cinema-projectors/2K-Digital-cinema-projector-for-screens-up-to-12m-40ft.aspx>, 2007.
- [76] S. Andrian, H. Brendel, T. Seybold, and J. Goldstone. Beyond the kodak image set: A new reference set of color image sequences. In *IEEE International Conference on Image Processing*, pages 2289–2293, 2013.
- [77] J. Froehlich, S. Grandinetti, B. Eberhardt, S. Walter, A. Schilling, and H. Brendel. Creating cinematic wide gamut HDR-video for the evaluation of tone mapping operators and HDR-displays. In *Proc. of IS&T/SPIE Electronic Imaging*, 2014.
- [78] L. L. Thurstone. A law of comparative judgment. *Psychological Review*, 34(4):273–286, 1927.
- [79] P.910 ITU-T RECOMMENDATION. Subjective video quality assessment methods for multimedia applications. 2008.
- [80] BT.500-13 ITU-R RECOMMENDATION. Methodology for the subjective assessment of the quality of television pictures. 2012.

- [81] C. Jian, V. Hardev, and J. Yurek. Quantum dot displays: Giving LCDs a competitive edge through color. *Nanotechnology Law & Business*, 11:4–13, 2014.
- [82] B. Beck. Lasers: Coming to a theater near you. <http://spectrum.ieee.org/consumer-electronics/audiovideo/lasers-coming-to-a-theater-near-you>, 2014.
- [83] 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.
- [84] M. R. Luo, G. Cui, and B. Rigg. The development of the CIE 2000 colour-difference formula: CIEDE2000. *Color Research & Application*, 26(5):340–350, 2001.
- [85] A. M. Eskicioglu and P. S. Fisher. Image quality measures and their performance. *IEEE Transactions on Communications*, 43(12):2959–2965, Dec 1995.
- [86] E. Prashnani, H. Cai, Y. Mostofi, and P. Sen. Pieapp: Perceptual image-error assessment through pairwise preference. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [87] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- [88] H. Talebi and P. Milanfar. Nima: Neural image assessment. *IEEE Transactions on Image Processing*, 27(8):3998–4011, 2018.
- [89] P. Cyriac, T. Batard, and M. Bertalmío. A nonlocal variational formulation for the improvement of tone mapped images. *SIAM Journal on Imaging Sciences*, 7(4):2340–2363, 2014.
- [90] 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.
- [91] A. R. Smith. Color gamut transform pairs. In *SIGGRAPH*, pages 12–19. ACM, 1978.



cinema professionals in order to improve vision models.

Marcelo Bertalmío received the Ph.D. degree in electrical and computer engineering from the University of Minnesota in 2001. He is a Full Professor at University Pompeu Fabra in Barcelona, Spain. Since 2012 he has been working on high dynamic range and wide color gamut imaging problems for cinema, and his favourite movies are the standard dynamic range and black and white films of Ernst Lubitsch and Luis Buñuel. His current interests are in the application of vision science to cinema and in using the know-how of



Syed Waqas Zamir received the Ph.D. degree from University Pompeu Fabra, Spain, in 2017. He is a Senior Research Associate at Inception Institute of Artificial Intelligence in United Arab Emirates. His research interests include color vision, computational photography and low-level vision.



Javier Vazquez-Corral received the Ph.D. degree in computer science from the Universitat Autnoma de Barcelona (UAB), Bellaterra, Spain, in 2011. He is a Postdoctoral fellow with the Department of Information and Communication Technologies at the Universitat Pompeu Fabra (UPF) in Barcelona. His research interests are related to the use of color in image processing and computer vision problems. He is also interested in bridging the gap between color in the human brain and its use in computer-vision applications.