

Saliency-driven Black Point Compensation

Albrecht Lindner^{a,b}, Nicolas Bonnier^b and Sabine Süsstrunk^a

^aEPFL, School of Computer and Communication Sciences, Lausanne, Switzerland;

^bOcé – Print Logic Technologies, Créteil, France

ABSTRACT

We present a novel framework for automatically determining whether or not to apply black point compensation (BPC) in image reproduction. Visually salient objects have a larger influence on determining image quality than the number of dark pixels in an image, and thus should drive the use of BPC. We propose a simple and efficient algorithmic implementation to determine when to apply BPC based on low-level saliency estimation. We evaluate our algorithm with a psychophysical experiment on an image data set printed with or without BPC on a Canon printer. We find that our algorithm is correctly able to predict the observers’ preferences in all cases when the saliency maps are unambiguous and accurate.

Keywords: black point compensation, saliency, color re-rendering, ICC, gamut mapping

1. INTRODUCTION

The conversion of an image from a device’s color encoding to another device’s color encoding is a difficult task if the target color gamut is significantly larger or smaller than the source color gamut. In the latter case, it can occur that an image contains pixels that are darker than the black point of the target gamut. There are two strategies to deal with this conflict: clipping and black point compensation (BPC). Clipping maps the too dark pixels to a nearby color value on the target gamut (e.g., the black point of the target gamut). This leads to a loss of details and tonal variations in the respective dark regions. BPC preserves these details by mapping most or all of these values inside the target gamut, thereby maintaining tonal variations in the dark areas but reducing contrast in the mid-tones and/or highlights of the image.

BPC significantly changes the visual appearance of an image, as illustrated in Figure 1. The birds and the wire images are rendered for a *Canon iPF700* inkjet printer with BPC enabled and disabled, respectively. The effect on dark regions can be observed in the two version of the birds image. The version with BPC (Fig. 1(a)) clearly shows the structures in the birds’ feathers. The version without BPC (Fig. 1(b)) lacks these structures due to clipping. But BPC also affects the mid-tones and highlights, as shown in the two versions of the wire image. The wire is rendered with more contrast in the version with disabled (Fig. 1(d)) than with enabled BPC (Fig. 1(c)).




Figure 1. Images of birds and a wire. Both images are rendered for a *Canon iPF700* inkjet printer in presentation print mode on standard red label paper. People prefer the birds with enabled (a) and the wire with disabled (d) BPC despite the fact that the wire image contains proportionally much more dark pixels than the birds image. For display on a screen, the images have been converted back to sRGB encoding. Depending on the display or printer used to reproduce these images, the differences might be more or less visible.

BPC can be seen as a trade-off between sacrificing the details in the dark regions in order to allow more contrast in the mid-tones and/or highlights versus retaining the details in the dark regions while forfeiting some contrast in the other image regions. A global mapping cannot simultaneously retain both the details in the dark regions and the contrast in the other regions. Thus, for a given image, a decision has to be made whether to enable or disable BPC.

Our hypothesis is that only the visually important regions of an image need to be considered in order to decide when to apply BPC or not, and not the whole image. This implies that the decision is driven not only by the difference in source and target gamut, but also by an image’s individual content. For the images in Fig. 1, the visually important image regions are those that show the birds or the wire, respectively. Even though in both images the backgrounds consist of more pixels than the objects, these regions are not visually important enough to drive the decision to enable or disable BPC.

In this paper, we propose a novel approach to automatically decide whether to apply BPC or not based on individual image content. According to our hypothesis it incorporates the concept of visual importance, which we estimate with saliency maps. The more a pixel is visually important, the higher its saliency value. We use a detector for low level saliency, which highlights regions that stand out to an observer if the image is presented for a short duration.¹ There are other types of saliency estimators that are of higher level or depend on a given task.² However, in this article we want to reflect an observer’s first impression and thus use low level saliency. We define a simple algorithm that uses saliency as a decision criterion: BPC is applied if and only if the dark regions are more salient than the rest of the image.

We evaluate our algorithm on a set of images that all contain dark regions. For the psychophysical experiment, we printed all images with and without black point compensation, presented them pairwise to different observers, and collected which version they prefer. We find that salient regions indeed play a major role in the judgement of the observers. When the saliency estimator correctly predicts visual importance, our algorithm correctly predicts observers’ preferences. Images with dark salient objects on non-dark background are preferred with BPC and images with non-dark objects on dark background are preferred without BPC (see Figure 1).

This paper is organized as follows. Section 2 reviews the state of the art. In Section 3 we give a qualitative explanation of our novel approach and explain the implementation. Section 4 presents the image set used for the evaluation. The experimental details and results follow in Section 5. The results are discussed in Section 6 and the article is concluded in Section 7.

2. STATE OF THE ART

Black point compensation (BPC) can be applied when the dynamic range of an image is larger or smaller than that of the target device. We limit the discussion and experiments here to the cases where the image’s dynamic range is larger. In this case, the dynamic range needs to be compressed, as is the case in many color re-rendering (e.g., printing) applications. Yet, the idea of using saliency in color rendering can also be used in other applications.

The simplest way to deal with pixel values that are lower than the target black point is clipping. However, if details in the dark regions need to be preserved, BPC has to be applied. There is a large amount of literature on how to compress an image’s dynamic range. Most of the literature considers local and/or global color rendering from scene to output-referred image encoding,^{3,4} or, in other words, how to best represent a scene’s dynamic range in a limited image encoding range, such as defined for popular color image encoding formats like sRGB or Adobe RGB. Here, we specifically address global black point compensation (BPC) that is applied within the framework of the International Color Consortium (ICC)⁵ when re-rendering from output-referred image encoding (sRGB, Adobe RGB, etc.) to print or display device encoding.

Adobe Systems’ implementation of BPC⁶ first converts to a normalized XYZ encoding with white point [1,1,1]. The values are then linearly scaled to fit in the target’s range and are converted from XYZ to the target gamut in XYZ or CIELAB. Alternative methods have been proposed, such as lightness compression algorithms (LCA),⁷ where the pixel values are re-mapped on the lightness axis in CIELAB color space. LCA and BPC have in common that they reduce the contrast in the highlights but also in the mid-tones. Sigmoidal LCAs aim

at compensating for this by enhancing the contrast in the mid-tones and therefore lowering the contrast in the shadows and highlights. See Morovic and Luo⁸ or Bonnier⁹ for an overview.

The preservation of details in the dark regions at the expense of reducing contrast in the rest of the image is not always preferable and depends on the image (see examples in Figure 1). Thus the question arises how one can decide whether BPC is advantageous for an image. To our knowledge, there are three different methods how that is handled today:

1. BPC is applied by default in order to preserve as much information as possible.¹⁰
2. Decision whether to en-/disable BPC is done by a human expert.
3. BPC is only applied if the image contains a significant amount of dark pixels.^{9,11}

The first approach is automatic but does not take the image characteristics into account. For the example of the wire image (Fig. 1 (c), (d)), this leads to a suboptimal result. The second approach probably leads to the best result, but it is subjective and expensive for a large number of images. The third method is automatic and is image dependent. However, in reality it is difficult to determine a valid threshold above which there is a *significant amount of dark pixels*. Also the examples in Figure 1 show there is no universal threshold. The image of the wire in Figure 1(c) contains proportionally more dark pixels than the image of the birds in Fig. 1(a), but still the birds image is better reproduced with BPC, unlike the wire image.

3. AUTOMATIC BPC USING VISUAL IMPORTANCE

In this section we give a qualitative explanation of the proposed algorithm (Section 3.1) and then a technical explanation of the implementation (Section 3.2).

3.1 Qualitative explanation of the approach

BPC can be understood as a trade-off between preserving details in the dark shadows and high contrast in the mid-tones and highlights. Since they are mutually exclusive when applying a global mapping, it has to be determined which of the two gives a perceptually better result. The key idea of our approach is to sacrifice the one that is visually less important for the given image. This means that BPC is applied if the dark regions are visually important and BPC is not applied if the dark regions are not visually important.

There are many methods to detect visually important regions in a given image, e.g., saliency methods,^{1,12,13} object detectors for specific classes,¹⁴ or for general purpose.¹⁵ We chose to use the saliency detector of Achanta et al.¹ since it is easy and fast to compute. It outputs saliency maps, which highlight regions that stand out to an observer at first impression.

A saliency map is usually encoded as a gray level image where the pixel values correlate with estimated visual importance. In an ideal case, a saliency map highlights the entire salient object uniformly and has a sharp outline that follows the objects' or regions' boundaries. An example saliency map can be seen in Figure 2.




Figure 2. Image of the two birds (a) and its saliency map (b). The brighter a pixel in the saliency map, the more it is estimated to be visually important in the original image.

3.2 Implementation

Given an input image I , we compute its saliency map S using the approach from Achanta et al.¹ A pixel's value at position (x, y) in the saliency map is:

$$S(x, y) = \|I_{avg} - I_{blur}(x, y)\|^2 \quad (1)$$

where I_{avg} is a 3-dimensional vector that is the average value over the whole image I in the CIELAB color space, $I_{blur}(x, y)$ is the CIELAB value at the pixel position (x, y) of a slightly blurred version of the original image (blurring removes noise and image artifacts) and $\|\cdot\|$ is the Euclidean distance.

We calculate the average saliency of the dark and the non-dark regions, respectively:

$$S_{\text{dark}} = \frac{1}{N_{\text{dark}}} \sum_{L(x,y) < L_{\text{bk}}} S(x, y), \quad S_{\text{non-dark}} = \frac{1}{N_{\text{non-dark}}} \sum_{L(x,y) \geq L_{\text{bk}}} S(x, y) \quad (2)$$

where $L(x, y)$ is the CIELAB L^* value at a pixel position (x, y) , L_{bk} is the CIELAB L^* value of the target black point and N_{dark} and $N_{\text{non-dark}}$ are the number of pixels in the dark and non-dark regions, respectively.

The binary decision whether to apply BPC or not is done with a simple comparison of the average saliencies:

$$\text{Decision} = \begin{cases} \text{enable BPC:} & S_{\text{dark}} > S_{\text{non-dark}} \\ \text{disable BPC:} & \text{otherwise} \end{cases}$$

A relative measure between two saliency values within an image is better than an absolute threshold since an observer looks at only one image at a time. In this case the different regions of the given image compete for the user's attention and the one with the highest saliency values is visually more important.

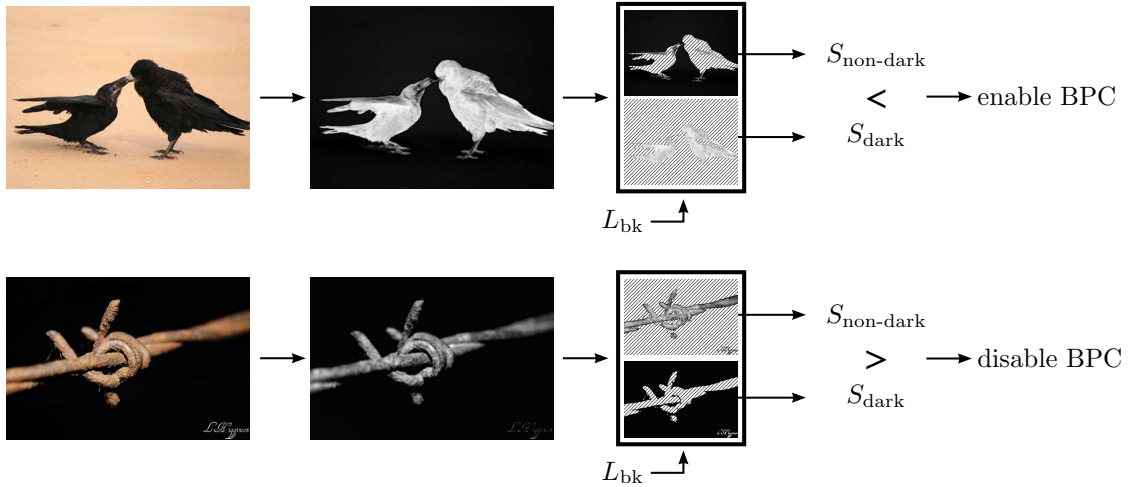


Figure 3. Diagram of the algorithm to decide whether or not to apply BPC. First a saliency map of the given image is calculated. Then the average saliencies (S_{dark} , $S_{\text{non-dark}}$) for the dark and the non-dark regions are calculated, respectively (e.g., for $S_{\text{non-dark}}$ the saliency values of the pixels that are darker than L_{bk} are not taken into account, which is depicted by overlying these regions with diagonal stripes). A comparison between the average saliencies leads to the decision to apply BPC or not. The top row shows an example where BPC is enabled ($S_{\text{dark}} > S_{\text{non-dark}}$) and the bottom row an example with disabled BPC ($S_{\text{dark}} \leq S_{\text{non-dark}}$).

4. IMAGE SET

To evaluate our saliency based decision criterion, we need a set of images that contain a significant amount of dark pixels. Further, the images need to be of sufficient quality for a high quality print, which we use for the psychophysical evaluation, and need to cover different scenes, such as dark main object on non-dark background and vice versa. To our knowledge, there exists no such publicly available database. There are public databases for research on saliency estimation¹⁶ but these collections generally contain low resolution images with JPEG compression. Therefore, we decided to create a database for our specific purpose. We downloaded images with a creative commons license from flickr* that contain a significant amount of dark pixels and that are suitable for high quality photo print. Most of the images had an embedded ICC profile and for those that did not, we assumed sRGB encoding. An overview of the 20 images is given in Figure 4 and the corresponding saliency maps are reproduced in Figure 5. The images and the saliency maps can be retrieved from our laboratory’s webpage: http://ivrg.epfl.ch/supplementary_material/AL_BPC

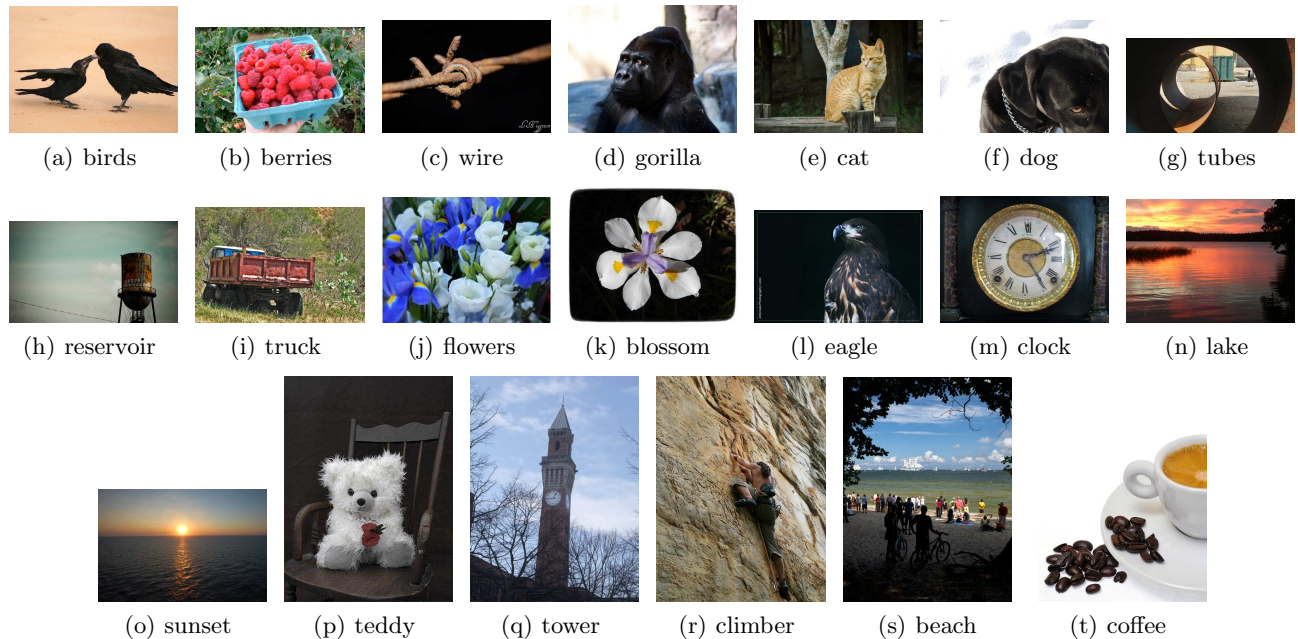


Figure 4. Dataset of images. Images have been chosen when they contain a significant amount of black pixels and are suitable for a high quality photo print.

5. EVALUATION OF PROPOSED METHOD

Section 5.1 explains the printing process and Section 5.2 presents the psychophysical experiment and the results.

5.1 Print setup

The images were scaled so that the longer edge was 5 inches long at 600 dpi resolution. They were printed on a *Canon iPF700* inkjet printer on coated matte paper in high quality print mode. Each image was printed using the ICC relative colorimetric rendering intent⁵ with enabled and disabled BPC, respectively. The printer ICC version 2 profile was calculated with an Océ software. For the used printer, the black point in Equation 2 is $L_{bk} = 16.80$. The conversion from source to target gamut was done using the Onyx ProductionHouse[†] Color Management Module (CMM). The CMM is an Onyx internally optimized version of the IccProfLib library found in the SampleICC project, which is publicly available.[‡] IccProfLib provides an example implementation of Adobe’s BPC,⁶ which is used to perform black point compensation in Onyx ProductionHouse.

*<http://www.flickr.com/creativecommons/>

†<http://www.onyxgfx.com/>

‡<http://sourceforge.net/projects/sampleicc/>

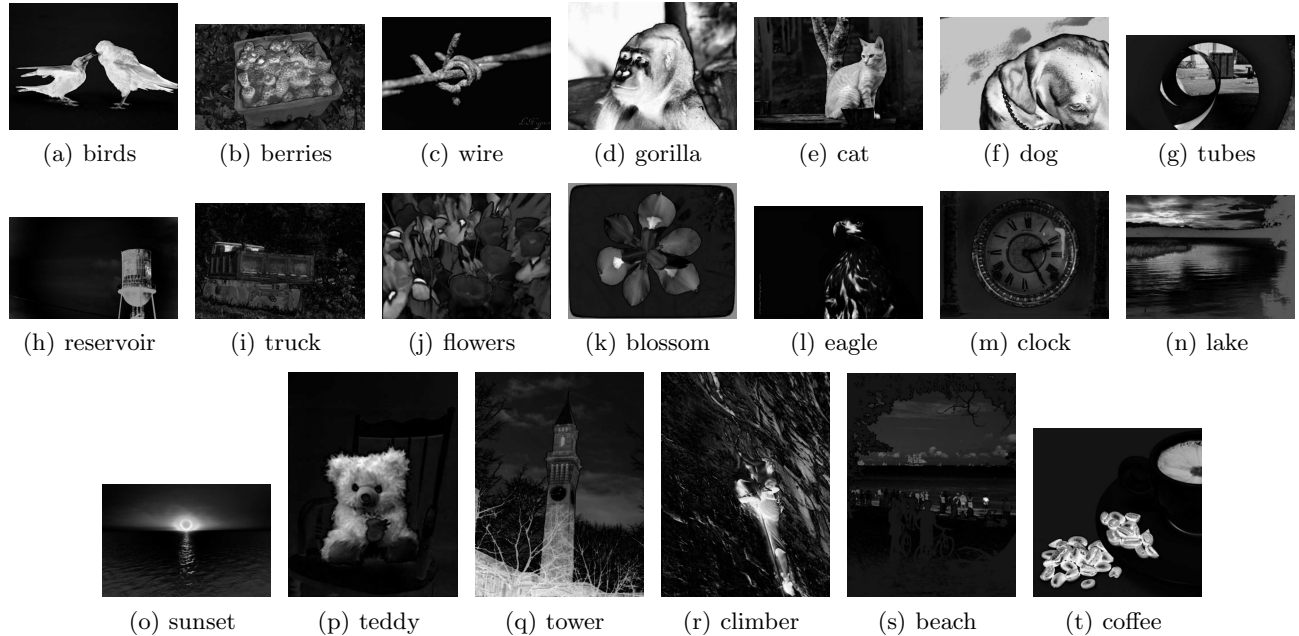


Figure 5. Saliency maps of all images from our dataset in Figure 4. Brighter pixels indicate more salient pixels.

5.2 Evaluation

The psychophysical experiment was conducted under controlled lighting conditions: a *Gretag Macbeth* viewing booth with daylight illuminant. The two versions of each image were shown in pairs to human observers. The given task was to decide which version they prefer. Table 1 shows the collected responses per image. We evaluated statistical significance for paired comparisons using the toolbox from Green and MacDonald.¹⁷

1. Correct predictions:

BPC	berries (b)	wire (c)	cat (e)	tubes (g)	blossom (k)	clock (m)	teddy (p)
enable	3	2	3	3	3	4	2
disable	<u>12</u>	<u>13</u>	<u>12</u>	<u>12</u>	<u>12</u>	<u>11</u>	<u>13</u>

BPC	birds (a)	gorilla (d)	dog (f)	reservoir (h)	coffee (t)
enable	<u>12</u>	<u>15</u>	<u>11</u>	<u>14</u>	<u>12</u>
disable	3	0	4	1	3

2. Ambiguous cases:

BPC	truck (i)	flowers (j)	eagle (l)	tower (q)
enable	<u>7</u>	7	<u>8</u>	<u>8</u>
disable	8	<u>8</u>	7	7

3. Erroneous predictions:

BPC	lake (n)	sunset (o)	climber (r)	beach (s)
enable	<u>5</u>	<u>3</u>	<u>5</u>	13
disable	10	12	10	<u>2</u>

Table 1. Observer’s responses. Significant tendencies towards enabled or disabled BPC are marked in bold font (95% confidence). The predictions of our algorithm are marked by underlining the number in the appropriate row. The 1st table lists images where our algorithm’s predictions match the observers’ opinions on disabled and enabled BPD, respectively. The 2nd table lists ambiguous cases, and the 3rd table lists erroneous predictions.

6. DISCUSSION

In this section we discuss the cases where the observers agree with our algorithm’s predictions (Section 6.1), ambiguous cases (Section 6.2) and the cases where our algorithm does not predict correctly (Section 6.3).

6.1 Correct predictions

Images with disabled BPC

The images for which our algorithm and the observers agreed on disabling BPC are (Fig. 4): berries (b), wire (c), cat (e), tubes (g), blossom (k), clock (m), and teddy (p). All these images have in common that they show a non-dark main object with a very dark non-salient surround. Many of the observers reported during the evaluation that they preferred disabled BPC since the contrast in the mid-tones and highlights is more important than the details in the dark regions. Since the saliency map correctly highlights the main object, the algorithm correctly decides to disable BPC.

Images with enabled BPC

The observers confirmed the estimation of our algorithm to enable BPC for the following images (Fig. 4): birds (a), gorilla (d), dog (f), reservoir (h), and coffee (t). In opposition to the previously discussed images, these images contain a dark main object on a non-dark background. In these cases, the details in the salient dark objects are judged visually important and thus BPC has to be enabled. If BPC is disabled the main object will be clipped and contain flattened regions that are not pleasing to a human observer (see Figure 1(b) for an example).

These images support our hypothesis that the main salient region (or regions) is more important for BPC estimation than non-salient other regions in an image. Observers do not value preserved contrast in regions that are not (or less) visually important. If the contrast has to be reduced for output on a device with smaller gamut, the salient regions’ contrast has to be retained as much as possible.

6.2 Ambiguous cases

The observers’ judgements were ambiguous for four of the images (Fig. 4): truck (i), flowers (j), eagle (l) and tower (q). These images contain competing salient regions that need contradictory BPC settings. Let us consider the eagle image (Fig. 4(l)) to illustrate. The judgements of the observers are almost even (8 en- and 7 disable). Participants who preferred the image with enabled BPC referred to the eagle’s back. The entire back of the eagle is very dark and the plumes are much more distinct from the background if BPC is enabled. Observers who preferred the image with disabled BPC pointed out the increased contrast of the eagle’s white cheek and the spots on its chest. In terms of applying BPC, the two competing salient regions are contradictory and observers do not agree whether the eagle’s spots are more or less important than the eagle’s back. The same ambiguity is present in the truck image (i) where the two competing regions are the truck’s red loading space and the dark wheels. Observers who preferred the image without BPC argued with the red color that looked better in this version. Participants who preferred the image with BPC pointed out that the wheels and axis of the truck are better visible.

These images highlight the shortcomings of all automatic methods, as none can actually take into account *individual* observers’ judgments. Color rendering is not just a science, it is also an art. Personal preference, based on an individual’s experience and interest will always play a role in the judgment of image quality, and we are still very far from predicting these effects.

6.3 Erroneous predictions

There are four images where human observers statistically significantly do not agree with the prediction from our algorithm. The images are (Fig. 4): lake (n), sunset (o), climber (r), beach (s). Especially the beach image (s) is of particular interest. It demonstrates the dependence of our method on a good saliency estimation. All observers who chose the image with enabled BPC argued with the increased visibility of the bikers in the foreground where the silhouette outline is better defined against the dark beach. Even though it is obvious to an observer that the bikers are visually important, the saliency estimator fails with this difficult scene since it

is based on color and not on shape (see Equation 1). This image illustrates that our hypothesis—salient regions are important to predict observer preference—is correct. However, if the automatic saliency estimation method fails, then our algorithm fails.

The climber image (r) has two competing salient regions (the climber and the rock wall). In this image, the region showing the wall is bigger than the region of the climber. In this case, observers gave more weight to the wall and preferred disabled BPC. The saliency estimator highlights the climber more than the wall and thus the algorithm’s prediction is wrong. In the lake image (n), the dark grass and trees are estimated to be salient enough to enable BPC but the observers preferred the more saturated appearance of the red and orange clouds in the version with disabled BPC. They used the same reasoning for the sunset image (o).

7. CONCLUSIONS AND FUTURE WORK

We propose an algorithm that provides an automatic decision criterion when to apply black point compensation (BPC) in color re-rendering applications, such as mapping from output-referred image to print encoding within the ICC framework. The decision criterion is based on a saliency map, which highlights in an image the pixels that are visually most important. The algorithm computes the average saliency of the image regions that are darker than the target’s gamut black point and compares to the average saliency of the other regions. If the average from the dark regions is higher, BPC is enabled. Accordingly, BPC is disabled if the non-dark regions are more salient.

To demonstrate the validity of our approach, we selected a set of 20 images with dark regions and varying scene content. All images were printed on a *Canon iPF700* with BPC enabled and disabled, respectively. The 20 image pairs were judged by 15 observers in a controlled lighting environment. The observers were asked to indicate for each image which reproduction they prefer: with enabled or disabled BPC.

The results show that our algorithm is well able to predict observers’ preference when the computed saliency map accurately highlights visual importance. When the saliency detection is ambiguous or wrong, then the observers are either conflicted or prefer the other version.

Future work will explore spatially adaptive BPC where each region of the image is processed according to its local content. Such techniques are used in high dynamic range image rendering¹⁸ and could also be adopted for BPC.

8. ACKNOWLEDGEMENTS

Many thanks to all people that make such research possible by sharing their images: Tatters:) (flower), René Ehrhardt (cat), Just chaos (gorilla), shuttermonkey (coffee), ell brown (tower), foxypar4 (birds), reservoir (jake-liefer), jinterwas (eagle), Repoort (clock), PeterJot (beach), Per Ola Wiberg ~ Powi (beach), windy_ (tubes), dok1 (truck), HiggySTFC (wire), UggBoy♡UggGirl [PHOTO : WORLD : SENSE] (flowers), edachsund (climber), tentwo.teneight (dog), redgoober4life (berries), Marius Pedersen (sunset), MSRA Salient Object Database¹⁶ (teddy).

REFERENCES

1. Achanta, R., Hemami, S., Estrada, F., and Süsstrunk, S., “Frequency-tuned salient region detection,” in [*Proc. IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*], (2009).
2. Yarbus, A., [*Eye Movements and Vision*], Plenum Press, New York (1967 (Originally published in Russian 1962)).
3. Holm, J., “Photographic tone and colour reproduction goals,” in [*Proc. CIE Expert Symposium 96 on Colour Standards for Image Technology*], 51–56 (March 1996).
4. Reinhard, E., Heidrich, W., Debevec, P., Pattanaik, S., Ward, G., and Myszkowski, K., [*High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*], Morgan Kaufmann, 2nd ed. (June 2010).
5. International Color Consortium, *Specification ICC.1:2004-10 (Profile version 4.2.0.0)* (2004).
6. Adobe Systems, “Adobe systems’ implementation of black point compensation.” <http://www.color.org/AdobeBPC.pdf> (downloaded 06/2010).

7. Braun, G. J. and Fairchild, M. D., “Image lightness rescaling using sigmoidal contrast enhancement functions,” *Journal of Electronic Imaging* **8**, 380–393 (October 1999).
8. Morovic, J. and Luo, M. R., “The fundamentals of gamut mapping: A survey,” *Journal of Imaging Science and Technology* **45**(3), 283–290 (2001).
9. Bonnier, N., *Contribution to Spatial Gamut Mapping Algorithms*, PhD thesis, Ecole Nationale Supérieure des Télécommunications (September 2008).
10. “Black-point compensation: theory and application,” White Paper 40, International Color Consortium (2010).
11. Uribe, J., “Understanding black point compensation.” http://cias.rit.edu/~gravure/tt/pdf/cm/TT5_Jorge01.pdf (downloaded 06/2010).
12. Itti, L., Koch, C., and Niebur, E., “A model of saliency-based visual attention for rapid scene analysis,” in [*IEEE Transactions on Pattern Analysis and Machine Intelligence*], **20**(11), 1254–1259 (1998).
13. Achanta, R. and Süsstrunk, S., “Saliency detection using maximum symmetric surround,” in [*Proc. IEEE International Conference on Image Processing (ICIP)*], (September 2010).
14. Sitjes, J. G., Bosch, X. B., de Weijer, J. V., Bagdanov, A., Gual, J. S., and Sabat, J. G., “Harmony potentials for joint classification and segmentation,” in [*Proc. IEEE Computer Vision and Pattern Recognition (CVPR)*], (June 2010).
15. Alexe, B., Deselaers, T., and Ferrari, V., “What is an object?,” in [*Proc. IEEE Computer Vision and Pattern Recognition (CVPR)*], (June 2010).
16. Liu, T., Sun, J., Zheng, N.-N., Tang, X., and Shum, H.-Y., “Learning to detect a salient object,” in [*Proc. IEEE Computer Vision and Pattern Recognition (CVPR)*], (2007).
17. Green, P. and MacDonald, L., eds., [*Colour Engineering: Achieving Device Independent Colour*], John Wiley & Sons (2002).
18. Meylan, L., Alleysson, D., and Süsstrunk, S., “A model of retinal local adaptation for the tone mapping of color filter array images,” *Journal of the Optical Society of America A (JOSA A)* **24**(9), 2807–2816 (2007).