

Privacy Protection of Tone-mapped HDR Images Using False Colors

 ISSN 1751-8644
 doi: 0000000000
 www.ietdl.org

 S. Çiftçi¹ ✉ A.O. Akyüz¹ A.M.G. Pinheiro² T. Ebrahimi³
¹ Department of Computer Engineering, Middle East Technical University, Ankara, Turkey

² Instituto de Telecomunicações, U.B.I., Portugal

³ Multimedia Signal Processing Group, École Polytechnique Fédérale de Lausanne, Switzerland

✉ E-mail: sciftci@ceng.metu.edu.tr

Abstract: High Dynamic Range (HDR) imaging has been developed for improved visual representation by capturing a wide range of luminance values. Because of its properties, HDR content might lead to a larger privacy intrusion, requiring new methods for privacy protection. Previously, false colors were proven to be effective for assuring privacy protection for low dynamic range (LDR) images. In this work, the reliability of false colors when used for privacy protection of HDR images represented by tone mapping operators (TMOs) is studied. Two different TMO techniques are tested, a simple TMO based on the Gamma transform and a more complex local TMO. Moreover, two false color palettes are also tested, and are applied to images that result from both TMOs and also to an LDR image that represents the center exposure in the image sequence used to create the HDR image. The degree of privacy protection is analyzed through both a subjective test using crowdsourcing and an objective test using face recognition algorithms (FRAs). It is concluded that the application of the two studied false color palettes reduces the recognition accuracy with respect to both tests.

1 Introduction

Throughout the world, visual surveillance has become a crucial aspect of everyday life. It is typically used for enhancing security, especially in critical areas, such as airports, train stations, and important buildings. It is also used to help individuals in difficult conditions such as in hospitals and elderly care centers.

Regardless of its application area, visual surveillance brings about the question of “how to protect the privacy of recorded individuals?”. While many studies tried to find an answer, as described in two comprehensive surveys [1, 2], a panacea for all “security-vs-privacy” issues appears elusive to be found.

The problem is further exacerbated due to the varied nature of image and video capture technologies. With the current trend shifting toward High Dynamic Range (HDR) capture systems [3, 4], privacy protection solutions proposed for non-HDR systems are becoming outdated.

A problem in HDR data used for surveillance is that, because typical monitors are Low Dynamic Range (LDR), the captured data must be tone-mapped first prior to display. Given the large number of tone-mapping operators (TMOs) [5], which operator is ideal for this task is currently unknown. However, previous work established that using HDR data, even when visualized by tone-mapping, improves face recognition accuracy of both human observers and face recognition algorithms (FRAs) [6]. Despite this important finding, it is currently unknown how tone-mapping operators interact with privacy protection solutions. The current study aims to shed light on this question.

To this end, we selected two TMOs representing simple and sophisticated solutions to tone-mapping as well as an LDR condition. We then applied a false color based privacy protection algorithm, previously found to be effective for LDR images [7–9], to these images using two different color palettes. The effectiveness of this result in preserving visual privacy is then assessed using both subjective and objective experiments. In doing so, we sought to answer the following research questions:

- Does false coloring reduce intelligibility of faces?
- Do different false color palettes have a different impact on intelligibility?

- How does false coloring interact with tone-mapping?
- Does false coloring have a different effect on face recognition algorithms than on human observers?

The rest of this paper is organized as follows. In the next section, we provide a brief overview of visual privacy protection especially in the context of HDR images. We then describe our dataset and evaluation methodologies. This is followed by our experimental results, discussions, and conclusions.

2 Related Work

2.1 Visual Privacy Protection

Privacy protection of multimedia content is gaining a special attention of the research community. The importance of multimedia content protection and the creation of multiple levels of accessing was considered in the earlier work of Eskicioglu *et al.* [10]. Anonymization of the visual content is of major importance for Visual Privacy Protection (VPP). Two main groups of methods can be considered, namely local and global methods [7]. Local methods consist of identifying the Region of Interests (ROIs) that require protection followed by a de-identification step [2, 11]. Global methods, on the other hand, apply privacy protection to the entire image.

The local method presents the advantage that only the ROIs are changed by applying a local de-identification process, while the remaining part of the image can be left untouched. However, this method requires accurate detection of the regions that are possible of providing an identification of the relevant information intended to be protected. These methods include simple techniques such as masking, blurring and pixelation, or more elaborated methods like warping [12], morphing [13], or scrambling [14]. However, most of these methods are not fully reversible, rendering it impossible to recover the original image in case of necessity by authorized users. Furthermore, usually computer vision methods are required for ROI detection, which do not guarantee robustness especially under non-ideal capture conditions. Finally, local methods work under the assumption that ensuring privacy of sensitive regions is sufficient to prevent privacy violations, which in the general case is not true




Fig. 1: False color palettes used in this study. The dynamic range of the color maps are 5.74 for *Blues* and 14.59 for *DEF*.

as the context information around these regions may also lead to privacy intrusions [15].

In contrast, global methods do not require ROI definition. The image anonymization process is applied to the full image. However, the resulting image should still be able to provide some generic content definition. Moreover, the image transformation should be somehow reversible, allowing the recognition by authorized users. This is the domain we are exploring in this work by using false colors for privacy protection.

2.2 VPP on HDR images

Typically, HDR images are displayed using 8-bit per color channel on LDR display devices. For this purpose, the displayed images are obtained by using a TMO that maps the wide input range to the reduced range of the target display in a controlled manner. A large number of TMOs are available [5], and they are also typically categorized as global and local. Global TMOs apply a global tone curve that maps the HDR colors into the typical 8-bit representation. However, these models often fail to preserve high-contrast details in local regions. Local TMOs try to overcome this limitation by applying local transformations. However, these TMOs are often computationally expensive and in many cases may compromise the global contrast in favor of maximizing visibility in local regions.

Independently of the created artifacts, TMOs are fundamental to provide a visualization of HDR images in the legacy 8-bit displays. The use of different TMOs is important in the context of this study as different TMO lead to different perceptual quality [16] and also to different recognition accuracy [6]. In [6] it was concluded that HDR images represented by TMOs create a larger privacy intrusion than LDR images. The conclusion of this work was in accord with the preliminary studies presented in [17, 18]. In [17] automatic face recognition using sparse representation was tested with tone-mapped HDR images. Eight different TMOs were tested and it was observed that some operators might lead to better automatic recognition than the use of common LDR images. In [18] a crowdsourcing subjective test was administered and was shown that the use of TMOs lead to better human recognition than the use of LDR images.

Moreover, in [6] a database with people present in multiple lighting conditions was created. Then the effect of five TMOs on face recognition was tested using both crowdsourcing and objective evaluations. It was concluded that the accuracy of human recognition can increase by more than 10% by using tone-mapped HDR images. Moreover, face recognition algorithms can also improve their accuracy. This gave rise to the conclusion that HDR images lead to higher privacy intrusion. In this work, we propose to augment this approach by employing false coloring after tone-mapping and investigate how it affects visual privacy with respect to both human observers and face recognition algorithms.

3 Dataset and Experiments

3.1 Dataset

In their earlier work, Korshunov *et al.* [6] shared a dataset comprised of HDR images containing people under various lighting conditions. These images were first created by combining a sequence of multiple exposures and then tone-mapped using 5 TMOs. The center exposure in the sequence was used to represent an LDR control condition. Additionally, portrait photographs of people involved in these images were also captured. Faces cropped from the tone-mapped and LDR images as well as the portrait photographs were then used in a recognition experiment.

In the current study, we created false color versions of 16 faces from the same dataset. These faces were selected so that various lighting levels were represented (i.e. we selected faces from shadowed regions, regions receiving direct sunlight, etc.) For false coloring, we used two color palettes abbreviated as *Blues* and *DEF* palettes (Figure 1). The *Blues* palette was taken from the National Library of Medicine Insight Segmentation and Registration Toolkit (ITK) [19]. It exhibits a color variation from white to blue such that low input luminances map to more white colors while high input luminances map to more blue colors. The second color palette, *DEF*, is used as the default color palette in the Radiance global illumination software for false color visualization of HDR images [20]. It was designed to maximize the number of named colors while still depicting a progression from cold to hot. Both color palettes were used in earlier visual privacy protection experiments and were found to be effective in protecting privacy [9].

As for TMOs, we first selected a simple operator known as Gamma mapping. This operator scales the luminance values such that the log-average luminance is 0.18 [21], followed by clamping to the $[0, 1]$ range. Finally, gamma correction is applied with an exponent of 0.45 closely approximating the sRGB gamma [22]. The second operator represents a local and more sophisticated approach to tone mapping, and is known as the *Mantiuk* operator [23]. We also used an LDR image, center exposure in the sequence, to compare the tone-mapped HDR pipeline with the standard LDR pipeline. The resulting face images before and after false coloring are shown in Figures 2 and 3.

With 16 faces, 2 TMOs, an LDR image (center exposure in the sequence), and 2 false color palettes, we obtained a dataset comprised of 96 different face images. These images are used in a visual recognition experiment as explained in the following section.

3.2 Subjective Experiment

For our experiment, we designed a web-based interface that starts with an entry page explaining the experimental task and collecting anonymous data from the participants such as their age, gender, and




Fig. 2: Face images in the dataset before being converted to false color.

familiarity with computer graphics and image processing. Following this page ensued a total of 96 trials. In each trial, a “probe” face was shown together with 9 “gallery” faces. The participants’ task was to select which of the gallery faces correspond to the probe face. The participants indicated their response by selecting a gallery face and pressing the “Next” button to continue with the next trial. Figure 4 illustrates a sample trial page.

The order of trials as well as the ordering of the gallery faces in each trial was randomized. We also inserted several fake trials to break the regularity of the experiment. This was made to render

it more difficult to make decisions for a trial based on earlier trials. With 96 real and 4 fake trials, the experiment could be finished within 15 minutes at a normal pace.

While more than 70 participants attended the experiment, only 48 of them completed it in entirety (20 females and 28 males). We only analyzed their data to obtain a uniform block design.

3.3 Objective Experiment

In addition to subjective evaluation, we also performed an objective evaluation experiment by using three face recognition algorithms (FRAs). For compatibility with the earlier work, we used Korshunov




Fig. 3: False color visualizations for some tone-mapping and color palette combinations.

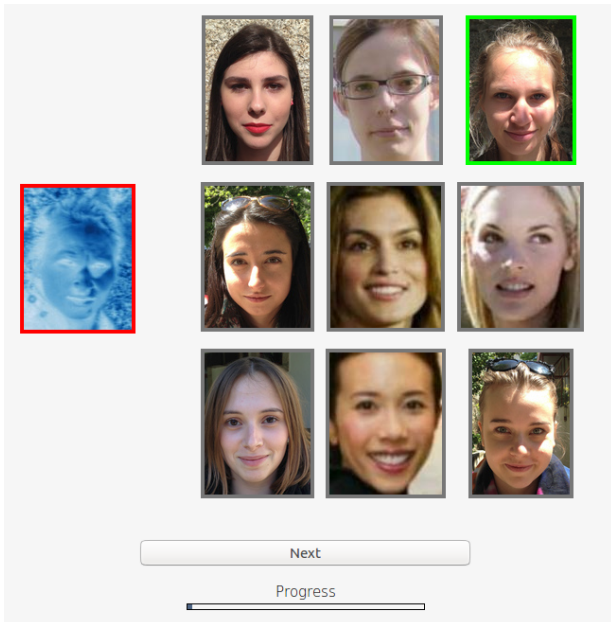


Fig. 4: A sample trial page with a probe face shown on the left and the gallery faces shown on the right. The participants' task was to select the correct face from the gallery that matches the probe face.

et al.'s framework [24] that automates training and testing three FRAs namely EigenFaces [25], FisherFaces [26], and local binary patterns histograms (LBPH) [27].

The Eigenfaces algorithm is based on an eigenvector decomposition of a large number of faces. Any face image can be reconstructed with a weighted combination of these face eigenvectors. Fisherfaces is proposed as an improvement over Eigenfaces, which uses Linear Discriminant Analysis (LDA) instead of Principle Component Analysis (PCA). LBPH classifies face descriptions using spatial histograms of micro-patterns [27, 28]. The basic local binary operator works on 3×3 pixel blocks of an image. The pixels in each block are coded by comparing the center pixel with the edge pixels of the block. The size of this block can be varied to support multi-scale analysis. The spatial histogram of this coded image is then used as the main feature for face recognition.

For training, we tone-mapped each HDR image using 5 different TMOs used in earlier work [6], namely *Gamma*, *Reinhard* [21],

Drago [29], *Mantiuk* [23], and *Mai* [30]. We also included the corresponding faces from the LDR image giving rise to 6 training images for each face. During testing, each false color face is evaluated by the three FRAs and the best matching face is returned. The results of both subjective and objective experiments are explained in the next section.

4 Results and Analysis

4.1 Subjective Experiment

The overall results of the subjective experiment is shown in Table 1. This table indicates that faces tone-mapped with the *Mantiuk* operator, when averaged with respect to all participants, faces, and color palettes, could be recognized with 0.69 accuracy. *Gamma* has a slightly higher accuracy of 0.71 while the *LDR* results are less likely to be recognizable (0.58). As for the color palette, the *Blues* palette, on average, produces a recognition rate of 0.63 while the *DEF* palette a rate of 0.69. We note that as trials were independent from each other and there were 9 gallery faces per trial, the chance recognition rate was $1/9 \approx 0.11$.

After performing this descriptive analysis, we conducted a three-way within-subjects ANOVA to identify which conditions are actually different from each other. In this analysis, the independent variables were *Face*, *TMO*, and *Palette*. Their interaction produced four extra variables, namely *Face* \times *TMO*, *Face* \times *Palette*, *TMO* \times *Palette*, and *Face* \times *TMO* \times *Palette*. The ANOVA results indicated that there is a significant difference between levels of all variables (all p values were less than 0.001), allowing us to reject the null hypothesis that there are no differences.

This allowed us to conduct multiple t-tests with Bonferroni correction as post-hoc to identify which levels are actually different from each other. In the following, we report the results for the *TMO* and *Palette* variables as they were the main focus of this study. Figure 5 shows the means and confidence intervals for different levels of each variable.

Consistent with these plots, the post-hoc tests indicated that *Gamma* and *Mantiuk* are indistinguishable from each other ($p = 0.48$) while the *LDR* has a statistically different and lower score. As for the color palettes, *Blues* and *DEF* were found to be statistically different ($p < 0.001$). For the interaction between *TMO* and *Palette*, we found that *LDR-Blues* is statistically less recognizable from all other combinations. The combination involving the *DEF* palette was always found to have a higher score than the combination with the *Blues* palette, suggesting that the latter better preserves privacy. This relationship appears to be reversed for the *Mantiuk* TMO, which is

Table 1 Results of the subjective experiment. Numbers represent the mean accuracy in $[0, 1]$ range.

	Blues	DEF	Mean
Mantiuk	0.70	0.68	0.69
Gamma	0.68	0.74	0.71
LDR	0.51	0.65	0.58
Mean	0.63	0.69	




Fig. 5: Means and confidence intervals for each independent variable.

an issue that we discuss in the following section. Statistical similarity groups for all levels of this interaction variable are shown in Figure 6.

4.2 Objective Experiment

The objective evaluation results using the three FRAs are reported in Table 2. Here it can be observed that the *Blues* palette, irrespective of the tone-mapping type, yields a recognition rate of almost zero (here the chance rate is $1/16 \approx 6\%$). As for the *DEF* palette, the highest recognition rates are obtained with the *Gamma* TMO, followed by *LDR* and then *Mantiuk*.

We also computed the structural similarity index (SSIM) [31] of the false color protected images with respect to their tone-mapped originals (Table 3). The results indicate that the images protected with the *Blues* palette have negative SSIM values, which is expected given the inverse luminance-pixel value relationship exhibited by this palette (see Figure 7). The SSIM values for the *DEF* palette are around zero suggesting that the images protected by this palette bear no negative or positive correlation to the original images. This is also expected given the up-and-down behavior of this palette for relating pixel value to luminance (also shown in Figure 7).

5 Discussion

The results of the current experiment especially when juxtaposed with the earlier works of Korshunov *et al.* [6] and Çiftçi *et al.* [7] reveal several interesting findings.

In [6], it was found that using HDR data, even when it is tone-mapped, improves face recognition accuracy with respect to both face recognition algorithms and human observers. It was also found that while FRAs are significantly affected by the choice of the TMO, human observers remain almost intact. In the current study, we determined what happens if the tone-mapped images are subjected to a false-color based privacy protection algorithm by observing that face recognition rates drop as reported in Table 4. According to this table, using the *Blues* color palette gives rise to a mean performance drop of about 20% (average of the first **Change** column). The *DEF* palette, on the other hand, produces a more meager drop of around 15% (average of the second **Change** column). Unfortunately, comparing these results through statistical confidence analysis is not feasible due to large differences between the number of participants in both experiments and a different selection procedure for gallery faces. However, large drops in recognition rates still allow us to answer our first and second research questions: false coloring

indeed reduces the intelligibility of faces and the amount of reduction depends on the used palette. This finding is also in line with the pioneering study of Rogowitz and Kalvin [32]. In that study, the authors established that when faces are visualized in false color, the degree to which the color palette is monotonically increasing in luminance is strongly correlated with the recognizability of a face. Palettes violating this property simulate a situation in which the dominant illumination is coming from the bottom, which rarely happens in nature. The *Blues* palette, as being monotonically decreasing as shown in Figure 7, violates this property more strongly than the *DEF* palette explaining its effectiveness in preserving privacy.

These findings also complement the earlier results of Çiftçi *et al.* [7]. In that work, the authors performed an experiment in which they asked participants' *opinion* about how well false colored faces are likely to be recognizable. The results had indicated that the *DEF* palette is the least likely to be recognizable among the four tested palettes (The *Blues* palette was not used in that study). Different from that work, the current study defined an objective recognition task and found that the recognition rates drop but not to the extent of being completely unrecognizable. However, given the sandboxed nature of the current experiment (all clean frontal view, one face compared with only nine faces), the actual recognition rate is likely to be much lower in a realistic surveillance context.

As for the effect of tone-mapping when combined with false-coloring, we found that there is no statistically significant difference between using a complex operator versus a simple operator answering our third research question. Perhaps unexpectedly, we found Mantiuk *et al.*'s operator [23] to produce marginally less recognizable results than a simple *Gamma* mapping. This could be attributed to the increased noise in that operator's outputs, especially for faces captured under a darker illumination (cf. Figure 2 (a) and (b)). This is typical of local operators as they tend to flatten the overall contrast in favor of local visibility. This noise when passed through a false color palette becomes decorrelated and, as a result, gets further exaggerated as exemplified in Figure 8.

Considering the objective experiment (Table 2), the fact that the *Blues* palette has a lower recognition rate than the *DEF* palette is in agreement with the subjective experiment. The almost zero recognition rates can be explained by the fact that these FRAs are very sensitive to relationships between luminance values of neighboring pixels. When this relationship is inverted, as it happens for the *Blues* palette as shown in Figure 7, the FRAs performance gets significantly hampered. It should, however, be noted that FRAs appear to be impacted much more significantly than human observers under this condition.

0.51	0.65	0.68	0.69	0.70	0.74
<u>LDR-Blues</u>	<u>LDR-DEF</u>	<u>Mantiuk-DEF</u>	<u>Gamma-Blues</u>	<u>Mantiuk-Blues</u>	<u>Gamma-DEF</u>

Fig. 6: Statistical significance groups for the interaction variable TMO×Palette. Items underlined by the same line are in the same group.

Table 2 Objective evaluation results using face recognition algorithms. Numbers indicate the accuracy in [0, 1] range.

	LDR-Blues	LDR-DEF	Gamma-Blues	Gamma-DEF	Mantiuk-Blues	Mantiuk-DEF
EigenFaces	0.06	0.25	0.00	0.94	0.00	0.31
FisherFaces	0.00	0.69	0.00	0.56	0.00	0.38
LBPH	0.00	0.88	0.00	0.94	0.13	0.25

Table 3 SSIM values computed for 16 face images with different TMO–palette combinations.

Images	LDR-Blues	LDR-DEF	Gamma-Blues	Gamma-DEF	Mantiuk-Blues	Mantiuk-DEF
1	0.02	0.00	-0.12	-0.10	-0.27	-0.14
2	-0.23	-0.14	-0.36	0.30	-0.28	0.17
3	0.02	0.00	-0.09	-0.04	-0.26	-0.13
4	0.01	-0.01	-0.19	-0.11	-0.40	-0.14
5	-0.45	0.09	-0.35	0.17	-0.47	0.32
6	-0.40	0.00	-0.51	0.37	-0.56	0.38
7	-0.08	-0.07	-0.53	0.02	-0.48	-0.07
8	-0.57	0.07	-0.55	0.38	-0.61	0.44
9	-0.14	-0.15	-0.50	0.10	-0.52	0.17
10	-0.04	-0.06	-0.26	-0.10	-0.20	-0.10
11	-0.30	-0.04	-0.44	0.14	-0.42	0.12
12	0.02	0.00	-0.08	-0.07	-0.42	-0.12
13	-0.44	0.10	-0.43	0.11	-0.45	0.26
14	0.00	-0.01	-0.25	-0.11	-0.37	-0.01
15	0.02	-0.02	-0.12	-0.04	-0.31	-0.06
16	-0.50	-0.14	-0.46	0.37	-0.53	0.48
Mean	-0.19	-0.02	-0.33	0.09	-0.41	0.10

Considering the DEF palette, the Mantiuk operator yields the lowest average recognition rate. This can also be attributed to the increased noise in the outputs of this operator. However, in contrast to the subjective experiment the recognition rate of Mantiuk-DEF is lower than that of LDR-DEF. This could be because human observers are more tolerant to noise than FRAs. We believe that these observations shed light onto our last research question on whether false coloring tone-mapped content affects FRAs differently than human observers.

6 Conclusions and Future Work

In this study, we set out to understand how to protect visual privacy in an HDR capture setting. To this end, we experimented with a recent and promising technique that involves using false colors to apply a global protection on the entire image rather than on defined ROIs. We found that this technique improves privacy against both human observers and face recognition algorithms. We further observed that using a simple tone-mapping operator combined with one of the color palettes actually yielded a greater recognition rate than the local operator, a finding which we did not anticipate. However, this increment was not statistically significant in the subjective experiment. We attributed this to the amplified noise that was present in the local operator’s outputs. This noise, when transformed by false coloring, appeared to outweigh the extra detail reproduction induced by the local TMO.

As future work, we intend to investigate how to protect privacy directly in the HDR domain without applying tone-mapping. The subband encoding scheme used by earlier studies [33, 34] could be employed to store restorative information as encrypted metadata in a backward compatible HDR format. This would allow the recovery of the original HDR information by authorized users. Finally, there appears to be a need to create an HDR video dataset for benchmarking privacy protection algorithms. We intend to focus on this as another promising line of future work.

7 Acknowledgment

This work was supported by The Scientific and Technological Research Council of Turkey (TUBITAK) project 114E445, Portuguese Funding Agency FCT project UID/EEA/50008/2013, and by COST Action IC1206.

8 References

- José Ramón Padilla-López, Alexandros Andre Chaaaroui, and Francisco Flórez-Revuelta. Visual privacy protection methods: A survey. *Expert Systems with Applications*, 42(9):4177–4195, 2015.
- Slobodan Ribaric, Aladdin Ariyaeinia, and Nikola Pavesic. De-identification for privacy protection in multimedia content: A survey. *Signal Processing: Image Communication*, 47:131–151, 2016.
- Bernd Hoefflinger. *High-dynamic-range (HDR) vision*. Springer, 2007.
- Peter Shirley Igor Garcia Olaizola Alan Chalmers, Patrizio Campisi. *High Dynamic Range Video: Concepts, Technologies, and Applications*. Academic Press (Elsevier), London, 2016.
- Erik Reinhard, Greg Ward, Sumanta Pattanaik, and Paul Debevec. *High Dynamic Range Imaging: Acquisition, Display and Image-Based Lighting*. Morgan Kaufmann, San Francisco, second edition edition, 2010.
- Pavel Korshunov, Marco V Bernardo, Antonio MG Pinheiro, and Touradj Ebrahimi. Impact of tone-mapping algorithms on subjective and objective face recognition in hdr images. In *Proceedings of the Fourth International Workshop on Crowdsourcing for Multimedia*, pages 39–44. ACM, 2015.
- Serdar Çiftçi, Pavel Korshunov, Ahmet Oğuz Akyüz, and Touradj Ebrahimi. Using false colors to protect visual privacy of sensitive content. In *IS&T/SPIE Electronic Imaging*, pages 93941L–93941L. Intl. Society for Optics and Photonics, 2015.
- Serdar Çiftçi, Pavel Korshunov, Ahmet Oğuz Akyüz, and Touradj Ebrahimi. Mediaeval 2015 drone protect task: Privacy protection in surveillance systems using false coloring. In *MediaEval Benchmarking Initiative for Multimedia Evaluation*, 2015.
- Serdar Çiftçi, Ahmet Oğuz Akyüz, and Touradj Ebrahimi. A reliable and reversible image privacy protection based on false colors. *IEEE Transactions on Multimedia*, PP(99):1–1, 2017. ISSN 1520-9210. doi: 10.1109/TMM.2017.2728479.
- Ahmet M Eskicioglu and Edward J Delp. An overview of multimedia content protection in consumer electronics devices. *Signal Processing: Image Communication*, 16(7):681 – 699, 2001.
- S. Ribaric and N. Pavesic. An overview of face de-identification in still images and videos. In *2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, volume 04, pages 1–6, May 2015.
- Pavel Korshunov and Touradj Ebrahimi. Using warping for privacy protection in video surveillance. In *Digital Signal Processing (DSP), 2013 18th Intl. Conf. on*, pages 1–6. IEEE, 2013.
- Pavel Korshunov and Touradj Ebrahimi. Using face morphing to protect privacy. In *Advanced Video and Signal Based Surveillance (AVSS), 2013 10th IEEE Intl. Conf. on*, pages 208–213. IEEE, 2013.

Table 4 Comparison of the face recognition rates of the current work by the earlier work of Korshunov *et al.*[6]. The first change column represents the change w.r.t the Blues palette whereas the second change column represents the same for the DEF palette.

	Korshunov et al. [6]	Blues	Change	DEF	Change
Mantiuk	0.869	0.701	16.8%	0.675	19.4%
Gamma	0.873	0.678	19.5%	0.742	13.1%
LDR	0.766	0.513	25.3%	0.650	11.6%




Fig. 7: The curves show the variation of relative luminance with respect to palette index.

14 Frederic Dufaux and Touradj Ebrahimi. Scrambling for privacy protection in video surveillance systems. *Circuits and systems for video technology, IEEE Trans. on*, 18(8):1168–1174, 2008.

15 Mukesh Saini, Pradeep K Atrey, Sharad Mehrotra, and Mohan Kankanhalli. W3-privacy: understanding what, when, and where inference channels in multi-camera surveillance video. *Multimedia Tools and Applications*, 68(1):135–158, 2014.

16 Patrick Ledda, Alan Chalmers, Tom Troscianko, and Helge Seetzen. Evaluation of tone mapping operators using a high dynamic range display. In *ACM Transactions on Graphics (TOG)*, volume 24, pages 640–648. ACM, 2005.

17 Manuela Pereira, Juan-Carlos Moreno, Hugo Proença, and António MG Pinheiro. Automatic face recognition in hdr imaging. In *Proc. SPIE*, volume 9138, pages 913804–913804, 2014.

18 Pavel Korshunov, Hiromi Nemoto, Athanassios Skodras, and Touradj Ebrahimi. Crowdsourcing-based evaluation of privacy in hdr images. In *Optics, Photonics, And Digital Technologies For Multimedia Applications Iii*, volume 9138. Spie-Int Soc Optical Engineering, 2014.

19 KitwarePublics. Colormaps. <http://www.itk.org/Wiki/Colormaps>. [Online; accessed 21-May-2016].

20 G.W. Larson and R. Shakespeare. *Rendering with Radiance: The Art and Science of Lighting Visualization*. Computer Graphics and Geometric Modeling Series. Morgan Kaufmann, 1998.

21 Erik Reinhard, Michael Stark, Peter Shirley, and James Ferwerda. Photographic tone reproduction for digital images. *ACM transactions on graphics (TOG)*, 21(3): 267–276, 2002.

22 ITU. *ITU-R Recommendation BT.709-5, Parameter Values for the HDTV Standards for Production and for International Programme Exchange*. ITU (International Telecommunication Union), Geneva, 2002.

23 Rafal Mantiuk, Alexander Efremov, Karol Myszkowski, and Hans-Peter Seidel. Backward compatible high dynamic range mpeg video compression. *ACM TOG*, 25(3):713–723, 2006.

24 Pavel Korshunov, Andrea Melle, Jean-Luc Dugelay, and Touradj Ebrahimi. A framework for objective evaluation of privacy filters in video surveillance. In *SPIE Applications of Digital Image Processing XXXVI*, volume 8856, San Diego, California, USA, August 2013. Spie-Int Soc Optical Engineering.

25 Matthew A Turk and Alex P Pentland. Face recognition using eigenfaces. In *Proc. Computer Vision and Pattern Recognition*, pages 586–591. IEEE, 1991.

26 Peter N Belhumeur, João P Hespanha, and David J Kriegman. Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. *Patt. Analy. and Mach. Intel., IEEE Trans. on*, 19(7):711–720, 1997.

27 Timo Ahonen, Abdenour Hadid, and Matti Pietikainen. Face description with local binary patterns: Application to face recognition. *Patt. Analy. and Mach. Intel., IEEE Trans. on*, 28(12):2037–2041, 2006.

28 Matti Pietikäinen, Abdenour Hadid, Guoying Zhao, and Timo Ahonen. Local binary patterns for still images. In *Computer vision using local binary patterns*, pages 13–47. Springer, 2011.

29 Frédéric Drago, Karol Myszkowski, Thomas Annen, and Norishige Chiba. Adaptive logarithmic mapping for displaying high contrast scenes. In *Computer Graphics Forum*, volume 22, pages 419–426. Wiley Online Library, 2003.

30 Zicong Mai, Colin Doutre, Panos Nasiopoulos, and Rabab K Ward. Subjective evaluation of tone-mapping methods on 3d images. In *Digital Signal Processing (DSP), 2011 17th International Conference on*, pages 1–6. IEEE, 2011.

31 Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

32 Bernice E Rogowitz and Alan D Kalvin. The " which blair project": a quick visual method for evaluating perceptual color maps. In *Visualization, 2001. VIS'01. Proc.*, pages 183–556. IEEE, 2001.

33 Greg Ward and Maryann Simmons. JPEG-HDR: A backwards-compatible, high dynamic range extension to JPEG. In *ACM SIGGRAPH 2006 Courses*, page 3. ACM, 2006.

34 Lin Yuan, Pavel Korshunov, and Touradj Ebrahimi. Privacy-preserving photo sharing based on a secure jpeg. In *2015 IEEE Conf. on Computer Communications Workshops*, pages 185–190. IEEE, 2015.




Fig. 8: Note the pronounced noise on Mantiuk *et al.*'s [23] after false color visualization.