

Compressed Exposure Sequences for HDR Imaging

Selin Sekmen
ASELSAN A.Ş.
Department of Computer Engineering
Middle East Technical University
06800, Ankara, TURKEY
ssekmen@aselsan.com.tr

Ahmet Oğuz Akyüz
Department of Computer Engineering
Middle East Technical University
06800, Ankara, TURKEY
akyuz@ceng.metu.edu.tr

ABSTRACT

High dynamic range (HDR) imaging techniques allow photographers to capture the luminance distribution in the real-world as it is, freeing them from the limitations of capture and display devices. One common approach for creating HDR images is the multiple exposures technique (MET). This technique is preferred by many photographers as multiple exposures can be captured with off-the-shelf digital cameras and later combined into an HDR image. In this study, we propose a storage scheme in order to simplify the maintenance and usability of such sequences. In our scheme, multiple exposures are stored inside a single JPEG file with the main image representing a user-selected reference exposure. Other exposures are not directly stored, but rather their differences with each other and the reference is stored in a compressed manner in the metadata section of the same file. This allows a significant reduction in file size without impacting quality. If necessary the original exposures can be reconstructed from this single JPEG file, which in turn can be used in a standard HDR workflow.

Keywords

HDR imaging, multiple exposures, JPEG, image metadata

1 INTRODUCTION

HDR imaging has been a rapidly emerging research field now finding applications in consumer photography and electronics [CNEa]. There are many approaches to create HDR images and videos including those involving dedicated HDR capture hardware [AlliC16a]. However, such devices are not only rare, but are also not preferred by most users due to their high cost. This typically leaves the photographers with the option of creating HDR images from multiple exposures. In this technique, a bracketed sequence of exposures are captured with a standard digital camera, with each exposure taken with a different exposure time. In such a sequence, short exposures contain detail in low-luminance regions whereas long exposures contain detail in high-luminance regions. By merging them into a single HDR image, the entire luminance range of the scene can be represented [Rei10a].

Arguably, the most problematic approach of this technique is the requirement of storing multiple images (typically 3 to 9) for each captured scene, together with

the HDR image itself. While the individual exposures can be purged after the HDR image is created, keeping them around may be preferable for several reasons. First, as typical display devices are low dynamic range (LDR), storing these LDR images allow one to rapidly view the captured scene, without resorting to tone mapping [Rei07a]. Second, the HDR image may often contain ghosting artifacts due to camera and object motion during the capture process. As new and improved algorithms are proposed continually, keeping the individual images allow one to create higher quality HDR images using these improved algorithms.

However, storing a large number of images for each scene (i.e. individual exposures, the HDR image, tone mapped image) is problematic for different reasons. It not only incurs extra storage costs but also makes photo-album maintenance a tedious task due to image repetition. These storage and maintenance problems may even discourage photographers from taking multiple exposures and shy away from HDR photography.

In this paper, we propose a method known as compressed exposure sequences (CES), which allows embedding the entirety of the information in a bracketed sequence into a single JPEG file. The main image, which will be shown by a standard image viewer, is typically selected as the middle exposure although another exposure or the tone mapped HDR image can be used as well. Differences between the subsequent exposures are stored in a compressed manner in the metadata sec-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

tion of the JPEG file. As demonstrated by several experiments, this scheme not only yields a smaller total file size but significantly simplifies the maintenance of exposure sequences while having a minimal impact on image quality.

2 RELATED WORK

There are many approaches to compress HDR images. JPEG-2000 [Sko01a] and JPEG-XR [Duf09a] have been developed to overcome the bit-depth limitation of the original JPEG format. Although the new standards offer superior compression performance to the original JPEG, they failed to achieve the desired level of adoption arguably due to their lack of backward compatibility with the original JPEG format.

All JPEG standards provide additional marker segments to keep an application specific data inside metadata sections [Pen92a, Wal92a]. This ability is used to create backward compatible HDR compression methods for images. JPEG-HDR is an extension to the standard JPEG format to store HDR images in a backward compatible manner [War06a]. The same ideas are extended to the video domain to allow for backward compatible HDR video storage [Man06a]. Beside these, there are dedicated HDR file formats such as RGBE, LogLuv, and OpenEXR [War94a, Lar98a, Kai02a]. However none of these file formats are backward compatible with existing software.

The most recent JPEG format, known as JPEG-Xt, also uses JPEG application markers to store HDR images. JPEG-Xt is backward-compatible with widely know JPEG format and provides lossy and lossless compression for high dynamic range images [Art15a].

The method proposed here can be seen as an extension of these backward-compatible storage schemes for HDR imaging. But unlike the existing methods that store a tone mapped image as the main image and auxiliary data for reconstructing the HDR image, we propose to store the entirety of information present in an exposure sequence. This simplifies the maintenance and usability of bracketed exposure sequences. It also allows for creating an HDR image at any time in future, using the original sequence stored in an efficient manner.

3 THE PROPOSED METHOD

The proposed method is termed as compressed exposure sequences (CES) as it aims to store multiple exposures inside a single JPEG file where the main image represents a user-selected reference exposure. CES is comprised of two main pipelines, one for compression and the other for decompression. In the compression pipeline, we use the multiple exposures of a scene and create a new JPEG image that stores all of the necessary data to recreate all of the original exposures. The

decompression pipeline reads the main image from the JPEG file together with its metadata and reconstructs the original exposures. These pipelines are illustrated in Figures 1 and 2 with the details elaborated below.

Compression Pipeline

The compression pipeline is demonstrated in Figure 1. Initially, all exposures are loaded and sorted according to their exposure times. Thus, we get the image set $\langle I_1, I_2, \dots, I_m, \dots, I_n \rangle$ from the shortest to the longest exposure, where I_m refers to the middle exposure. All of these exposures are comprised of 24-bits per pixel (i.e., 8-bits per color channel). We also extract the exposure times from the EXIF data during exposure loading, and label them as $\langle E_1, E_2, \dots, E_m, \dots, E_n \rangle$.

During the compression pipeline, the main aim is to store the minimum necessary data for each exposure inside the JPEG application markers, so that we can recreate each exposure using this auxiliary information. To save space, exposures are not directly stored inside the metadata, instead for each exposure the difference with respect to a reference image is stored. Therefore, at the initial state of the pipeline, we choose a reference image for each exposure according to Equation 1:

$$I_{ref_x} = \begin{cases} I_{x+1} & \text{if } x < mid, \\ I_{x-1} & \text{if } x > mid. \end{cases} \quad (1)$$

In other words, the reference exposure for a given exposure is an adjacent exposure in terms of exposure time. Choosing such adjacent exposures as reference maximizes the coherency between them.

The next stage in the pipeline is to align each exposure with its reference to compensate for the camera movement during the image capture process. The median threshold bitmap (MTB) algorithm [War03a] was used to determine misalignments. This algorithm was selected due to its simplicity, efficiency, and robustness against exposure differences. After the alignment process, we apply the inverse camera response function (iCRF) on aligned nonlinear exposures to linearize them (Equation 2). The resulting linear images are symbolized as L_x . Without loss of generality, we used the radiometric self calibration approach [Mit99a] to recover the camera response function. At this point we move from the integer domain to the floating point domain:

$$L_x = 255 f^{-1} \left(\frac{I_x}{255} \right). \quad (2)$$

The difference between an exposure and its reference is computed as defined in Equation 3:

$$L_{dif_x} = |L_x - k_x L_{ref_x}|. \quad (3)$$

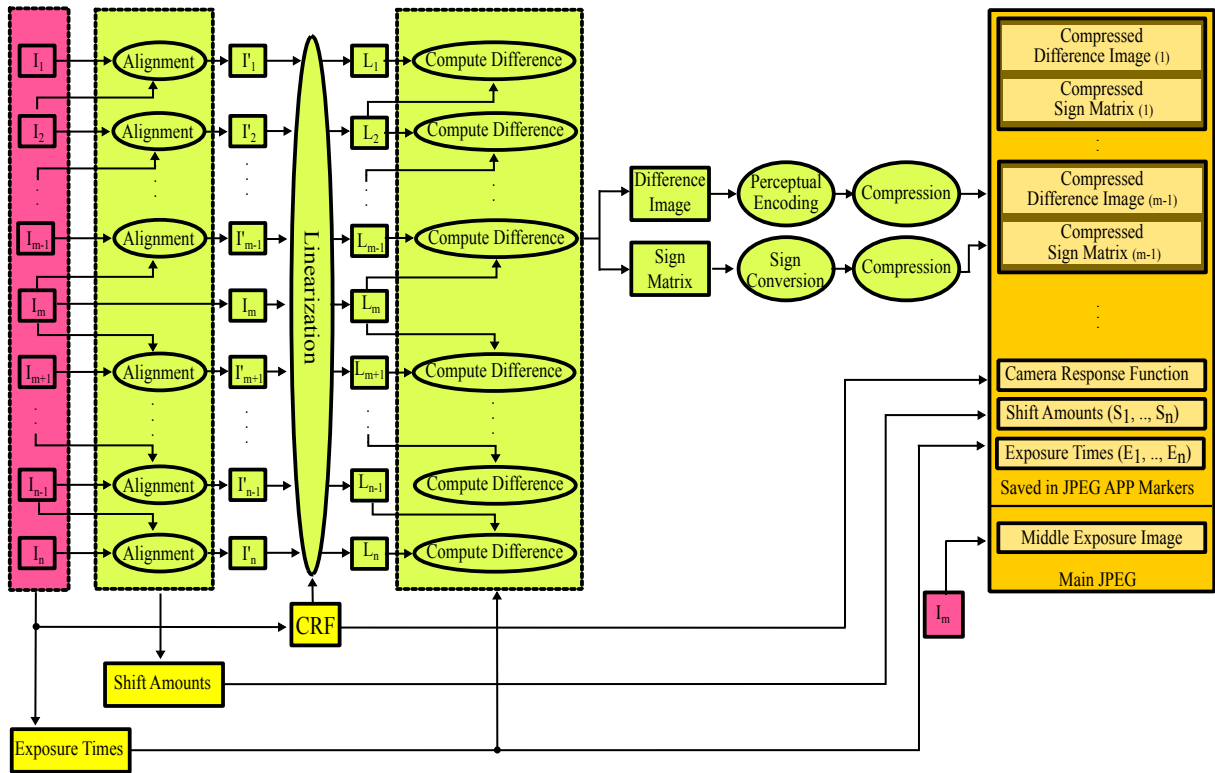


Figure 1: The compression pipeline of the CES algorithm.

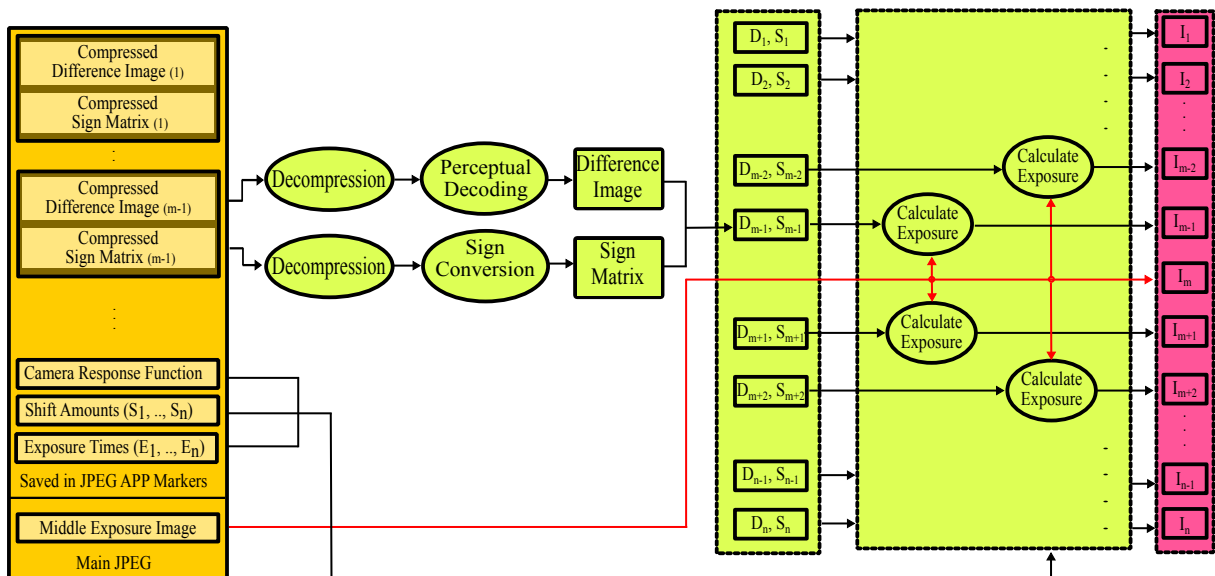


Figure 2: The decompression pipeline of the CES algorithm.

Before computation of this difference, we find the best exposure ratio k_x that minimizes the mean value of this difference using the gradient descent algorithm:¹

$$k_x = \arg \min_k |L_x - kL_{ref_x}|. \quad (4)$$

This is done to bring the difference values as close as possible to zero to obtain higher compression rates.

The sign matrix stores the sign of the difference matrix for each pixel. The values for the sign matrix are selected as either 0 or 10, with the former indicating a positive difference and the latter a negative one. The reason for selecting these values instead of 0 and 1 is that if we apply lossy compression for the sign matrices, we can interpret values less than 5 as positive and the others as negative. This provides robustness against the lossy compression artifacts:

$$L_{sign_x} = \begin{cases} 10 & \text{if } L_x - k_x L_{ref_x} < 0 \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

As explained above, the result of the computed differences yield two data set for each exposure, namely the difference and sign matrices. Before applying compression to difference matrices, we apply gamma correction to them to simulate a simple form of perceptual encoding as shown in Equation 6. This provides a larger pixel value range for dark pixels for which the quantization artifacts would be more noticeable. In this step, we move back to the integer domain which enables better compression at the cost of introducing quantization errors:

$$P_{dif_x} = 255 \left(\frac{L_{dif_x}}{255} \right)^{1.0/2.2} \quad (6)$$

After perceptual encoding, we apply compression to the difference and sign matrices to decrease the data storage size. Our method allows for both lossless and lossy compression. Lossless compression is done using the ZLIB's Huffman coding based compression algorithm [Deu96a], whereas for lossy compression we use JPEG compression. The results of compressions are saved inside the JPEG metadata sections. In addition to the difference and sign matrices, we also store the exposure time ratios, alignment data, and the camera response function coefficients inside the JPEG application markers as well. The middle exposure, I_m , is stored as the main JPEG image. It is important to note that this middle exposure is not re-encoded, but copied as a binary data stream to prevent re-encoding artifacts.

¹ Note that the actual exposure time ratios are not always the best choice for k due to the presence of under- and over-exposure regions.

Decompression Pipeline

The decompression pipeline is the reverse of the compression pipeline. Initially, the main JPEG image and its metadata are loaded followed by decompressing the metadata to recover the sign and difference matrices (Figure 2).

This is followed by applying perceptual decoding to undo the effect of perceptual encoding that was performed during compression:

$$L_{dif_x} = 255 \left(\frac{P_{dif_x}}{255} \right)^{2.2}. \quad (7)$$

Next, each exposure is reconstructed one by one starting from the exposures closest to the middle exposure by using Equation 8:

$$L_x = k_x L_{ref} + L_{dif_x} L'_{sign_x} \quad (8)$$

where

$$L'_{sign_x} = \begin{cases} 1 & L_{sign_x} < 5 \\ -1 & \text{otherwise} \end{cases} \quad (9)$$

At this point, all of the computed data is still in linearized form and must be converted to the nonlinear camera response domain using the recorded camera response function:

$$I_x = 255 f \left(\frac{L_x}{255} \right) \quad (10)$$

Finally, we align the exposures using the stored shift amounts to reconstruct a close approximation of the original JPEG exposures.

4 RESULTS

In this section, we present the results of the experiments conducted using a dataset of 10 exposure sequences representing various scene conditions (Figure 3).

The quality measurements are done using two well-known objective image quality metrics, namely structural similarity index (SSIM) [Wan04a] and peak signal-to-noise ratio (PSNR) [Teo94a] using the original exposures as ground-truth for both metrics.

Initially, we show the visual quality of our method in lossless compression mode. Then we demonstrate the effect of lossy compression on both file size and recovered image quality.

Lossless Compression

Here we illustrate the lossless compression and decompression results of our algorithm via a test case containing five exposures of the *Mug* image. In Figure 4, the first line shows the original exposure set, the second line displays the recovered exposures by using our

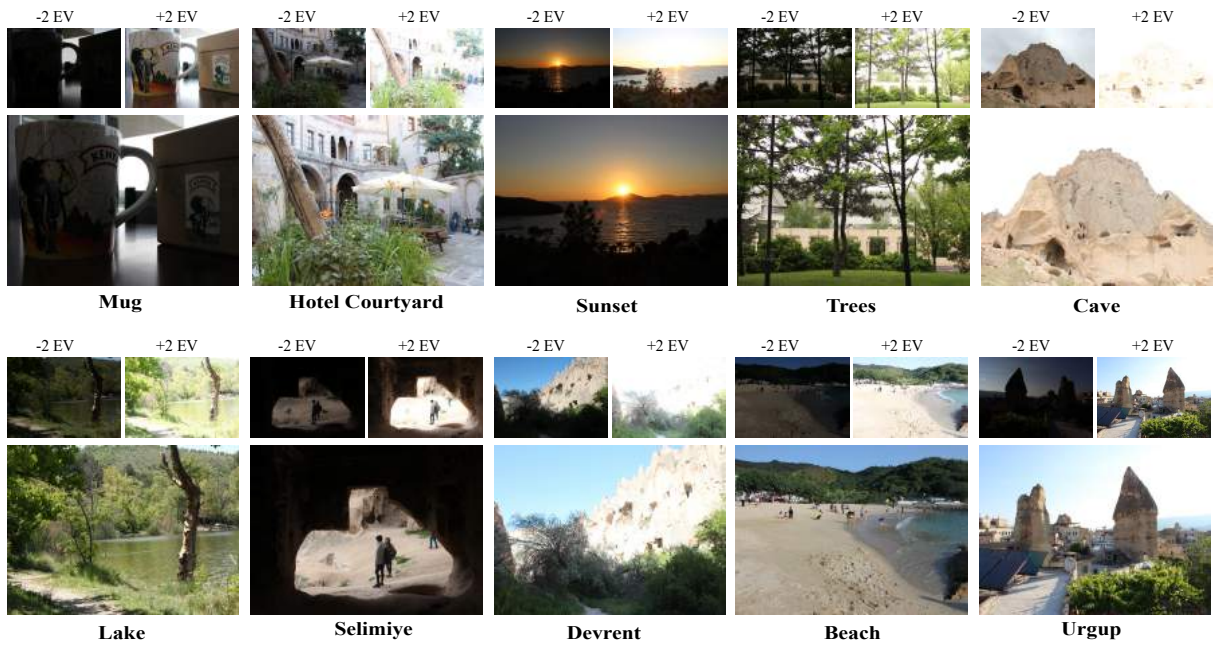


Figure 3: Our quality experiments are conducted using 10 exposure sequences whose middle exposures with two other exposures are shown in this figure.

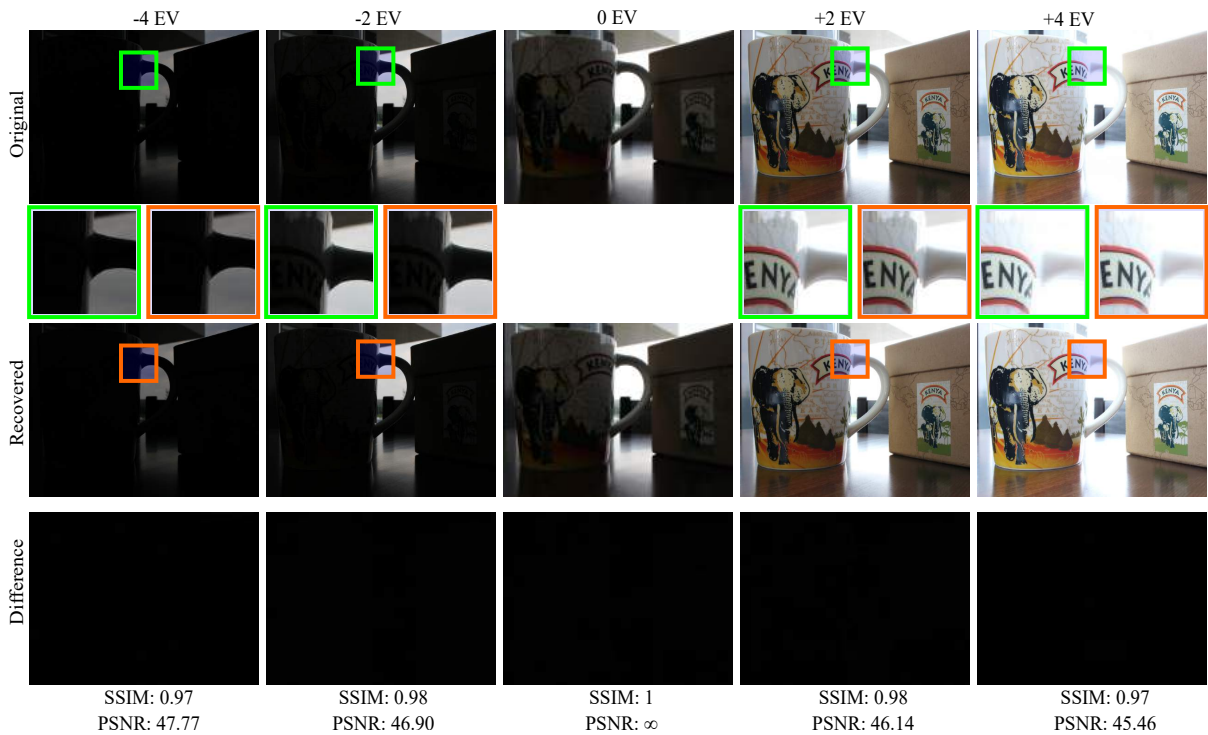


Figure 4: The results of lossless compression for a test image. The insets show the selected regions in the original resolution.

method, and finally the last line represents the difference between the original and the recovered images.

For each exposure we calculate the SSIM index and the PSNR value to show the perceptual and numerical similarity between the recovered and original images. The results show that the recovered exposures are very close to the original ones in the lossless compression case, as expected. The reason for not obtaining exact equality is due to the quantization of the image differences to whole numbers. If the storage size is not very critical and the main aim is to be able to reconstruct the original exposures, our lossless compression mode offers a plausible solution.

Lossy Compression

For many applications, reducing the image size can be important to minimize the storage and transmission costs, and therefore using the lossless mode may not be desirable. In Figure 5, we report the results obtained in case the image difference data is compressed in a lossy manner. Similar to Figure 4, the top row shows the originals, the middle row our reconstructions, and the bottom row their differences. Although the SSIM and PSNR values can be seen to have dropped compared to the lossless mode, the reconstructions can still be considered close to the originals.

The Effect of Perceptual Encoding

In this section, we demonstrate the effect of using perceptual encoding before compressing the image difference data. In Figure 6, (a) represents the original image, (b) the result obtained with perceptual encoding, (c) the result without perceptual encoding, and (d) the difference between the two. As can be seen both visually and numerically indicated by the SSIM and PSNR values, perceptual encoding has an important contribution to the quality of the reconstructed exposures.

The Effect of Exposure Ratio Optimization

We described in Section 3 that optimizing the exposure ratios to compute image differences leads to better results than using the actual exposure time values (see Equation 4). This can be explained using the following simplified example. Imagine that two pixels' values are 170 and 200 in an exposure captured with Δt exposure time. Due to sensor saturation, the same pixels' values could be both clamped at 255 in the next exposure captured with $2\Delta t$ exposure time. If one uses the actual exposure time ratio to normalize the second image one would obtain 127.5 for both pixels. The difference of these with 170 and 200 would yield 42.5 and 72.5 respectively. These are the values that would be compressed with the JPEG compression algorithm.

On the other hand, if one scales the second image with a ratio of $185/255$, both pixels in the difference image would attain the value of 15. Note that not only

this difference would compress better, the reconstruction error from encoding/decoding such a difference would actually be smaller. Although this example is an over-simplification, we observed in practice that it is the very reason that optimizing the exposure ratios improves the reconstruction quality over using the actual exposure times. As demonstrated in Figure 7, the benefit of this optimization is most visible in long exposures with large saturated regions.

Table 1 shows the effect of this optimization for 4 different image sets each comprised of 9 exposures. It not only leads to decrease in the compressed image size, but also produces higher SSIM and PSNR values.

The Effect of JPEG Compression

Lossy compression results are also affected by JPEG compression quality, which control the degree of quantization of the DCT coefficients. We therefore conducted experiments to determine the effect of JPEG compression on both image quality and size. The results are shown in Table 2. When we consider both image quality and the amount of compression achieved, the quality level of 85 yields a good compromise. At this level, the compressed image size was found to stay lower than the total size of the original exposures, while the SSIM and PSNR values were generally high.

5 CONCLUSION & FUTURE WORK

In this study, we presented a scheme called compressed exposure sequences that provides a simple solution for storage and maintenance of bracketed exposure sequences commonly used by HDR photographers. We used JPEG as a container format due to its wide availability and its support for metadata, which makes our solution possible. Our two modes, namely lossless and lossy, provide flexibility to the users. Those who prioritize image quality over storage size may choose to use the lossless mode. Conversely, those users who are mainly concerned with the storage size can opt for the lossy mode. In both modes, the middle exposure is stored exactly as it is produced by the camera so the users can be certain that it remains unaltered by our algorithm.

We have shown that plausible results are obtained by using a dataset of 10 exposure sequences. As one of the primary future works, we plan to extend this dataset with more images. Secondly, the effect of camera and object movement on both image quality and storage size must be better evaluated. Finally, as bracketed sequences resemble video frames that are captured in close succession, video compression algorithms can be used to further reduce the storage size without compromising quality.




Figure 5: The results of lossy compression with JPEG compression quality value of 85 for a test image.




Figure 6: a: Original image (Urgup), b: Reconstructed image with perceptual encoding, c: Reconstructed image without perceptual encoding, d: Difference between b and c (10 times enlarged).

CES Size (MB)	Mug (32.6 MB)				Urgup (41.5 MB)				Sunset (39.5 MB)				Selimiye (33.7 MB)			
	Computed		Actual		Computed		Actual		Computed		Actual		Computed		Actual	
	Exposure Rate	PSNR	Exposure Rate	PSNR	Exposure Rate	PSNR	Exposure Rate	PSNR	Exposure Rate	PSNR	Exposure Rate	PSNR	Exposure Rate	PSNR	Exposure Rate	PSNR
	23.4		27.8			31.4	33.7			31.7	33.3			29.8	32	
<i>Exp 1</i>	0.96	45.36	0.96	44.48	0.96	44.64	0.95	42.66	0.94	37.68	0.93	37.48	0.97	46.49	0.96	43.34
<i>Exp 2</i>	0.96	44.40	0.97	44.91	0.94	41.35	0.93	40.37	0.92	36.01	0.93	36.16	0.96	44.64	0.96	42.90
<i>Exp 3</i>	0.95	42.82	0.96	44.03	0.94	39.46	0.93	38.55	0.90	34.53	0.91	34.85	0.95	41.83	0.95	41.68
<i>Exp 4</i>	0.95	41.33	0.96	43.14	0.94	38.41	0.93	37.94	0.93	34.87	0.94	35.17	0.94	40.62	0.94	40.94
<i>Exp 5</i>	1.00	∞	1.00	∞	1.00	∞	1.00	∞	1.00	∞	1.00	∞	1.00	∞	1.00	∞
<i>Exp 6</i>	0.96	42.35	0.96	41.36	0.93	33.53	0.91	32.71	0.95	32.26	0.94	31.68	0.93	37.84	0.94	38.69
<i>Exp 7</i>	0.95	40.04	0.90	36.68	0.90	29.88	0.83	27.21	0.92	29.80	0.89	27.68	0.90	34.82	0.89	34.21
<i>Exp 8</i>	0.95	39.51	0.83	32.96	0.88	28.15	0.76	23.30	0.92	30.25	0.83	25.17	0.92	36.06	0.82	31.85
<i>Exp 9</i>	0.95	38.50	0.77	29.50	0.90	28.43	0.67	20.86	0.92	30.71	0.72	22.27	0.93	34.87	0.72	27.20

Table 1: The effect of exposure rate optimization (Equation 4) on total size and image quality. The *Computed Exposure Rate* column shows the results obtained by using Equation 4 to compute the exposure ratios. The *Actual Exposure Rate* column shows the results if the original exposure time ratios are used. Note that the optimization yields both smaller total size (*CES Size*) and better image quality (*SSIM* and *PSNR*).



Figure 7: a: Original image (Sunset, shortest exposure), b: Reconstructed image with exposure rate optimization, c: Reconstructed image with the original exposure time ratios, d: Difference between b and c (10 times enlarged). e: Original image (Sunset, longest exposure), f: Reconstructed image with exposure rate optimization, g: Reconstructed image with original the exposure time ratios, h: Difference between f and g (10 times enlarged).

6 REFERENCES

- [All16a] Peter Shirley Alan Chalmers, Patrizio Campisi and Igor Garcia Olaizola. *High Dynamic Range Video: Concepts, Technologies, and Applications*. Academic Press (Elsevier), London, 2016.
- [Art15a] Alessandro Artusi, Rafał K. Mantiuk, Thomas Richter, Philippe Hanhart, Pavel Korshunov, Massimiliano Agostinelli, Arkady Ten, and Touradj Ebrahimi. Overview and evaluation of the jpeg xt hdr image compression standard. *Journal of Real-Time Image Processing*, pages 1–16, 2015.
- [CNEa] CNET. What is HDR for TVs, and why should you care? <https://www.cnet.com/news/what-is-hdr-for-tvs-and-why-should-you-care/>. Accessed: 2017-11-25.
- [Deu96a] Peter Deutsch and Jean-Loup Gailly. Zlib compressed data format specification version 3.3. Technical report, 1996.
- [Duf09a] Frédéric Dufaux, Gary J Sullivan, and Touradj Ebrahimi. The jpeg xr image coding standard [standards in a nutshell]. *IEEE Signal Processing Magazine*, 26(6):195–204, 2009.
- [Kai02a] F. Kains, R. Bogart, D. Hess, P. Schneider, and B. Anderson. OpenEXR. <https://www.openexr.org/>, 2002. Accessed: 2017-11-26.
- [Lar98a] Gregory Ward Larson. Overcoming gamut and dynamic range limitations in digital images. *Color and Imaging Conference*, 1998(1):214–219, 1998.
- [Man06a] Rafał Mantiuk, Alexander Efremov, Karol Myszkowski, and Hans-Peter Seidel. Backward compatible high dynamic range mpeg video compression. *ACM TOG*, 25(3):713–723, 2006.
- [Mit99a] T. Mitsunaga and S. K. Nayar. Radiometric self calibration. In *Proceedings. 1999 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (Cat. No PR00149)*, volume 1, page 380 Vol. 1, 1999.
- [Pen92a] William B Pennebaker and Joan L Mitchell. *JPEG: Still image data compression standard*. Springer Science & Business Media, 1992.
- [Rei07a] Erik Reinhard, Timo Kunkel, Yoann Marion, Jonathan Brouillat, Rémi Cozot, and Kadi Bouatouch. Image display algorithms for high- and low-dynamic-range display devices. *Journal of the Society for Information Display*, 15(12):997–1014, 2007.
- [Rei10a] Erik Reinhard, Greg Ward, Sumanta Pattanaik, and Paul Debevec. *High Dynamic Range Imaging: Acquisition, Display and Image-Based Lighting*. Morgan Kaufmann, San Fran-

		JPEG QUALITY															
		40		50		60		70		80		85		90			
Original Exposure Size (MB)	CES Size (MB)	%		%		%		%		%		%		%			
		SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR		
Mug	32.6	10.1	30.98	11.2	34.36	12.5	38.34	14.9	45.71	19.3	59.20	23.4	71.78	31	95.09		
	Min	0.92	36.02	0.93	36.85	0.93	37.48	0.94	38.37	0.94	39.24	0.95	39.89	0.95	40.56		
	Avg	0.94	39.06	0.94	39.87	0.94	40.10	0.95	41	0.96	41.72	0.96	42.32	0.96	42.93		
	Max	0.96	44.07	0.96	44.90	0.97	44.89	0.97	45.71	0.97	46.29	0.97	46.74	0.97	47.15		
Hotel Courtyard	50.9	21.4	42.04	23.5	46.17	25.9	50.88	29.8	58.55	36.5	83.30	42.4	83.30	53.2	104.52		
	Min	0.74	24.74	0.76	25.57	0.77	26.12	0.79	26.78	0.81	27.87	0.83	28.56	0.84	29.6		
	Avg	0.80	28.20	0.82	29.27	0.83	29.71	0.84	30.41	0.86	31.61	0.87	32.38	0.89	33.29		
	Max	0.93	38.83	0.94	39.91	0.94	40.31	0.95	40.96	0.96	41.54	0.96	42.27	0.96	42.59		
Sunset	39.5	14.40	36.46	16.00	40.51	17.90	45.32	21.20	53.67	26.70	67.59	31.80	80.51	40.70	103.04		
	Min	0.77	26.06	0.78	26.71	0.78	27.29	0.79	28.10	0.80	29.10	0.81	29.74	0.82	30.66		
	Avg	0.83	29.27	0.84	29.97	0.85	30.48	0.85	31.17	0.86	32.06	0.87	32.63	0.88	33.44		
	Max	0.88	33.05	0.89	33.53	0.89	33.80	0.90	34.12	0.91	34.71	0.92	35.19	0.93	35.83		
Trees	64.8	27.1	41.82	29.8	45.99	32.8	50.62	37.6	58.02	45.8	70.68	53	81.79	65.9	101.70		
	Min	0.66	20.20	0.68	20.84	0.70	21.26	0.72	21.8	0.74	22.58	0.76	23.16	0.78	24.04		
	Avg	0.72	25.12	0.75	25.84	0.76	26.38	0.78	27.08	0.8	28.05	0.82	28.71	0.84	29.61		
	Max	0.83	32.65	0.85	33.45	0.86	33.92	0.88	34.73	0.89	35.6	0.9	36.19	0.91	36.88		
Cave	30.60	14.00	45.75	15.60	50.98	17.40	56.86	20.30	66.34	25.10	82.03	29.40	96.08	37.50	122.55		
	Min	0.71	27.90	0.73	28.81	0.74	29.36	0.76	30.07	0.79	31.15	0.80	31.84	0.83	32.94		
	Avg	0.86	34.24	0.87	34.87	0.88	35.61	0.89	36.62	0.90	37.85	0.91	38.65	0.92	39.72		
	Max	0.98	39.84	0.98	39.48	0.98	40.88	0.99	42.81	0.99	44.26	0.99	45.38	0.99	46.60		
Lake	65.6	28.6	43.60	31.7	48.32	35	53.35	40.3	61.43	49.4	75.30	57.5	87.65	71.8	109.45		
	Min	0.52	17.98	0.55	18.57	0.58	19.07	0.61	19.69	0.64	20.6	0.67	21.27	0.7	22.31		
	Avg	0.67	24.67	0.69	25.47	0.71	26.05	0.74	26.74	0.77	27.84	0.79	28.57	0.81	29.52		
	Max	0.84	35.35	0.86	36.30	0.87	36.78	0.89	37.51	0.91	38.55	0.92	39.2	0.93	39.95		
Selimiye	33.7	12.60	0.37	14.20	0.42	16.00	0.47	19.20	0.57	24.80	0.74	29.90	0.89	39.30	1.17		
	Min	0.72	30.78	0.73	31.51	0.73	31.94	0.74	32.48	0.75	33.42	0.75	33.96	0.76	34.66		
	Avg	0.82	34.49	0.82	35.24	0.83	35.55	0.84	36.07	0.85	36.83	0.85	37.28	0.86	37.79		
	Max	0.88	39.17	0.88	39.85	0.89	39.91	0.90	40.31	0.91	40.86	0.92	41.14	0.93	41.40		
Devrent	57.2	23.4	40.91	25.9	45.28	28.7	50.17	33.2	58.04	40.8	71.33	47.5	83.04	59.3	103.67		
	Min	0.73	23.63	0.75	24.52	0.76	25.05	0.78	25.84	0.81	26.95	0.83	27.73	0.85	28.92		
	Avg	0.79	25.51	0.81	26.47	0.83	27.13	0.85	28.02	0.87	29.23	0.88	30.07	0.9	31.25		
	Max	0.86	27.67	0.88	28.51	0.89	28.97	0.91	30.01	0.93	31.47	0.94	32.55	0.95	33.43		
Beach	55.7	24.20	43.45	26.90	48.29	29.80	53.50	34.60	62.12	42.70	76.66	50.00	89.77	62.90	112.93		
	Min	0.68	22.61	0.71	23.12	0.73	23.66	0.75	24.41	0.76	25.27	0.76	25.91	0.77	26.88		
	Avg	0.73	27.41	0.76	28.41	0.77	28.86	0.79	29.54	0.81	30.31	0.83	30.84	0.84	31.59		
	Max	0.76	34.63	0.79	35.27	0.81	35.36	0.83	35.54	0.86	36.15	0.88	36.52	0.90	36.88		
Urgup	41.5	14.70	35.42	16.30	39.28	18.10	43.61	21.20	51.08	26.60	64.10	31.60	76.14	40.50	97.59		
	Min	0.77	27.74	0.77	25.51	0.77	25.99	0.77	26.59	0.78	27.51	0.78	28.23	0.78	29.24		
	Avg	0.81	30.92	0.82	31.54	0.83	31.89	0.84	32.41	0.85	33.11	0.85	33.60	0.86	34.23		
	Max	0.86	37.26	0.87	37.57	0.88	37.67	0.89	37.98	0.90	38.21	0.91	38.38	0.92	38.55		

Table 2: Lossy compression results with respect to different JPEG qualities (40 – 90); each JPEG quality column gives the size of the compressed sequence of images (*CES Size*), the relative size as a percentage of the original total exposure size (%), and the quality metric results (*SSIM* & *PSNR*).

cisco, second edition edition, 2010.

[Sko01a] Athanasios Skodras, Charilaos Christopoulos, and Touradj Ebrahimi. The jpeg 2000 still image compression standard. *IEEE Signal processing magazine*, 18(5):36–58, 2001.

[Teo94a] P. C. Teo and D. J. Heeger. Perceptual image distortion. In *Proceedings of 1st International Conference on Image Processing*, volume 2, pages 982–986 vol.2, Nov 1994.

[Wal92a] Gregory K Wallace. The jpeg still picture compression standard. *IEEE Trans. on Consumer Electronics*, 38(1):xviii–xxxiv, 1992.

[Wan04a] 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.

[War94a] Gregory J. Ward. The radiance lighting simulation and rendering system. In *Proceedings*

of the 21st Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '94, pages 459–472, New York, NY, USA, 1994. ACM.

[War03a] Greg Ward. Fast, robust image registration for compositing high dynamic range photographs from hand-held exposures. *Journal of graphics tools*, 8(2):17–30, 2003.

[War06a] 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.