Backwards Compatible JPEG Stereoscopic High Dynamic Range Imaging

E. Selmanović, K. Debattista, T. Bashford-Rogers and A. Chalmers

WMG, University of Warwick

Abstract

In this paper we introduce Stereoscopic High Dynamic Range (SHDR) Imagery which is a novel tecnique that combines high dynamic range imaging and stereoscopy. Stereoscopic imaging captures two images representing the views of both eyes and allows for better depth perception. High dynamic range (HDR) imaging is an emerging technology which allows the capture, storage and display of real world lighting as opposed to traditional imagery which only captures a restricted range of light due to limitation in hardware capture and displays. HDR provides better contrast and more natural looking scenes. One of the main challenges that needs to be overcome for SHDR to be successful is an efficient storage format that compresses the very large sizes obtained by SHDR if left uncompressed; stereoscopic imaging requires the storage of two images and uncompressed HDR requires the storage of a floating point value per colour channel per pixel. In this paper we present a number of SHDR compression methods that are backward compatible with traditional JPEG, stereo JPEG and JPEG-HDR. The proposed methods can encode SHDR content to little more than that of a traditional LDR image and the backward compatibility property encourages early adopters to adopt the format since their content will still be viewable by any of the legacy viewers.

Categories and Subject Descriptors (according to ACM CCS): E.4 [Data]: Coding and Information Theory—Data compaction and compression

1. Introduction

Stereoscopic imagery and high dynamic range (HDR) imagery are relatively new imaging methods that both provide advantages over traditional monocular low dynamic range (LDR) imagery. Possibly due to their novelty and perhaps due to technical complications, these methods have not been combined together as yet; however, there is no reason to envisage that they cannot be complimentary. Stereoscopic high dynamic range (SHDR) imagery, as we term this combination, has the potential of providing richer fidelity when compared to the traditional LDR methods by combining the advantages of HDR and stereoscopic methods. Many challenges remain for this coupling to prove fruitful and in this paper we tackle the task of storage, in particular in creating a compression method that is not only efficient but also provides backward compatibility with the traditional JPEG LDR and JPEG-HDR [War06] to improve the possibilities of early adoption.

Stereoscopic (frequently called 3D) imaging is a technol-ogy that captures and displays images representing the two human eyes. The advantages of such a method relies on enabling or improving the illusion of depth. It can also improve task performance [WM08] and provide a strong cue for distance judgements [SGJ92]. Stereoscopic imagery recently became more widely available as it has now become popular in the consumer market. Consumer products that both capture and display stereoscopic images can be purchased at reasonable price and are slowly becoming standard. The entertainment industry is also adopting this techniques with more movies being released in 3D format and video games allowing 3D visualisation. Even some of the latest smartphones implement this technology.

HDR imagery is a novel technology [BADC11] that has seen several advances over the past few years. HDR enables the capture, storage and display of real world luminance, thus producing images that are more representative of the real world. While a traditional LDR image is limited to around eight exposures, HDR does not have such limi-

tations and is able to handle any luminance values that the human eye can see and more. HDR capture of static images has now become widespread, and is available on standard smartphones; HDR video is however still relatively unknown and only a few prototype cameras exist [CBB*09, TKTS11], hence we focus on still images in this paper. There have been a number of ways of storing HDR content, particularly for still images and we shall discuss these in Section 2. Display of HDR images has been mostly limited to tone mapping methods [BADC11] that compress the luminance for traditional displays. There are a few HDR displays available but not for the consumer market [SHS*04].

While certain benefits of HDR are straightforward such as the ability to capture and visualise content that would be otherwise over or under exposed, HDR can also benefit stereoscopic content in one of its stronger areas, depth perception, as it can provide further depth information via contrast [RHM11].

SHDR content has the potential of achieving the best of both worlds, resulting in images with higher fidelity than current imaging methods. One draw back of SHDR may be the size. SHDR images can be quite large if left uncompressed. A raw SHDR image at an HD resolution of $1,920 \times 1,080$ can be just short of 48MB. With the emergence of superHD's $7,680 \times 4,320$ resolution this would balloon to 759MB. The major contribution of this work is the ability to efficiently store SHDR images. We propose a number of novel methods which would make it possible to store SHDR images which are only marginally larger than traditional LDR images. Furthermore, four of the five methods presented provide the ability of opening any SHDR image in a traditional LDR viewer and in a traditional HDR viewer while the two of the methods are backwards compatible with LDR stereo viewers.

2. Related Work

When stored as a raw data, HDR images are considered to be composed of three floating point values, one for each of the red, green and blue channels, for a total of 96-bits per pixel (bpp); this results in 24MB per frame at HD resolution. Due to such prohibitively large sizes, a number of methods for storing HDR images exist. Ward [War91] introduced the RGBE format for use with the Radiance [War94] software. This format stores RGB values as an 8-bit mantissa per pixel and furthermore stores an 8-bit exponent for a total of 32-bpp. The LogLuv [Lar98] format also proposed by Ward stores images by storing colour and luminance separately. LogLuv supports two formats, a 24-bpp format and a 32-bpp format. The 24-bpp stores 10 bits for the luminance channel and 14 bits for the colour channels, but only supports just short of 5 orders of luminance magnitude. The 32-bpp format can achieve 38 orders of magnitude with 16 bits for chroma and 16 for luminance. The OpenEXR format [KBH04] is probably the most common HDR format

used by the entertainment industry, it stores each of the channels as 16-bit half precision float point values for a total of 48-bpp. It is frequently further compressed using lossless methods.

The above methods are considered storage methods for representing HDR. Another set of HDR imaging method formats are more akin to traditional compression. In particular, there are methods which are compatible with JPEG and JPEG2000. Ward and Simmons [WS04, War06] presented JPEG-HDR, an HDR compression method that is backward compatible with the JPEG format. This method stores a tone mapped version of the HDR image which is encoded using JPEG and the extra information termed the ratio image is stored in a sub-band of the JPEG format. The ratio image can be downsampled significantly as the human visual system is not very sensitive to low frequency changes in luminance. When opened by a traditional JPEG viewer the sub-band is ignored and the tone mapped image can be seen. The decoding process recovers the missing information from the tone mapped image from the sub-band image (which is upsampled back to the original size) to recover the HDR image. Xu et al. presented HDR-JPEG2000 [XPH05]. HDR-JPEG2000 transforms the HDR data into the 16-bit unsigned shorts supported natively by JPEG2000. Okuda and Adami [OA07] presented a method similar to JPEG-HDR which includes wavelet compression for the residuals and minimisation for the tone mapping parameters.

Stereoscopic images, in their most basic form are stored as two separate images. Perhaps, the most popular stereo image format is the JPEG stereo (JPS) [SS97] which supports a number of formats, storing images side by side, potentially at half the resolution or interleaved. A good overview of different techniques for stereoscopic content, focusing on stereoscopic video is provided by Smolic et al. [SMS07].

3. JPEG SHDR Methods

Our backward compatible JPEG SHDR method is inspired by Ward and Simmons' [War06] approach. As with this method we use the JFIF wrapper as a format for storing all the extra data that our SHDR methods produce. The main entry is used to store a tone mapped version of one of the stereo pairs (and for the HSBS instance a side by side version). This format provides storage channels for metadata. This is used to store the additional information such as ratio images, disparity maps and motion compensation information, depending on the technique proposed in order to restore the full SHDR content, as shall be discussed below. While, the number (16) of metadata channels and size (64KBs) are limited, they are sufficient for our proposed methods; furthermore, if required, this limitation can be overcome by using more storage channels which have the same identifier. For SHDR we also strive to achieve compatibility with current LDR stereo standards. We aim to achieve a balance between image size, quality and backward compatibility resulting in five methods

each of which is explained further in the following subsections.

3.1. Side-by-side (SBS) Method

The first technique we suggest strives to preserve quality of the original image and minimise data loss at the expense of larger sizes compared to the other methods. It is also a good foundation for further examination as it provides an initial reference point in terms of quality which other approaches should try to attain. Both of the images in this case are coded using JPEG-HDR so that the quality of the resorted image is kept high and further file size reduction is not considered.

This straightforward method starts by appending the right HDR image of the stereo pair to the left one. The result is a single side-by-side HDR image. This image is then compressed using JPEG-HDR. The resulting size is almost equivalent to compressing each image separately, and the fact that there is a large amount of correspondence between the stereo pair is not taken advantage of.

There are two ways of formatting the output depending on viewer the image is aimed for. Data from the second image can be put into JPEG subbands making it available for traditional monocular image viewers which would show only one of the image pair. Alternatively, the tone-mapped images can be left side-by-side and saved as a stereo JPEG (JPS) which can be then viewable in LDR stereo viewers. This variant can also be opened in traditional viewers but both the left and right views would be displayed. While such behaviour provides at least some insight into the content of the file, it may not be desirable for the user.

3.2. Half Side-by-side (HSBS) Method

Another standard way of saving LDR stereo images puts the pair side by side but halves the horizontal resolution of each such that both can fit in the space of a single image. This is not dissimilar to having them interleaved. In this method, the coding process starts by resizing the images of the HDR pair as described and proceeds in the same manner as SBS. When the image is being decompressed images are resized back up again. Compared to the SBS method the image size is roughly halved but so is its resolution affecting the resulting quality. This method is included primarily because it is backward compatible with LDR stereo viewers.

3.3. Image Plus Disparity (IPD) Method

In IPD method we exploit the correlation between the left and right views and the correspondence between most of the pixels (exceptions include occlusions, transparencies and specular highlights). The image that represents these correspondences is disparity map. The precise maps can be obtained by using specialist equipment while taking stereoscopic photographs or they can be provided by the rendering

software in case of CG images. However, for the majority of stereo images, disparity data is not available and needs to be calculated using the stereo pair. This calculation is an ill-posed problem and the correct solution cannot be found for the majority of the scenes. Nevertheless, the problem has been widely researched by previous publications. Scharstein et al. [SSZ01] provide a taxonomy of stereo correspondence algorithms. They also created a testbed which evaluates the performance of more than one hundred of these algorithms.

Any disparity map can be used with this method (calculated or captured), however the smooth low frequency ones are preferred because of better compression rates. For our results the disparity maps used were obtained employing technique suggested by Mei et al. [MCS*11] which is named AD-census (ADC) and it is considered one of the best. This method utilises the GPU during calculations leading to fast performance, and also the technique produces the least number of errors according to aforementioned testbed. For further details and exact implementation steps please refer to author's article [MCS*11].

The encoding process using the IPD method is shown in Figure 1. It starts by generating a disparity map from HDR pair, if it is not already present. This map is then encoded using lossless LZW and Huffman coding. The output file size is not fixed but it is a fraction of the original image and can easily fit in JPEG subbands. One of the HDR images is compressed using HDR JPEG and thus stores the ratio image in the subband and a tone-mapped JPEG image. Decoding inverses the process (see Figure 2) by extracting and recovering the disparity map and restoring one of the images using JPEG-HDR decoding. This HDR image is warped using disparity to obtain the missing view. In the experiment we encoded the right image using HDR JPEG and restored the left through warping but this can be switched.


Figure 1: IPD encoding

Figure 2: IPD decoding

3.4. Image Plus Disparity with Corrections (IPDC) Method

During testing of the IPD method it was observed that some occluded regions did not restore well during image warping stage. Here edges were misplaced and a number of cases had major offset issues. The cause of such problems was that disparity maps were smoothed and background pixels that the restored image required, were occluded by the foreground objects in the available image. This resulted in those foreground objects being warped to wrong positions and some of the objects being perceived at incorrect depths.

The alternative to the method used for IPD is to use a disparity map that maps only the closeness of RGB values, such as the sum of absolute differences (SAD) method [CS09]. Such disparity maps do without the smoothness condition that maps such as ADC use and are therefore higher frequency. The high frequency means that these disparity maps would not compress well if used by themselves. Our proposed method avoids the problems of both these methods by combining them and using SAD in regions with large differences only. This process is shown in Figure 3 and is used instead of the step highlighted in the dashed square in Figure 1. Once the low frequency map is obtained using the ADC method, the image is warped and divided by the original. Differences above an empirically obtained threshold are identified. Those pixels on the ADC disparity map are updated with ones from high frequency map (which is obtained using SAD). The rest of the coding process is identical to IPD one. Decoding is the same as in the IPD case.
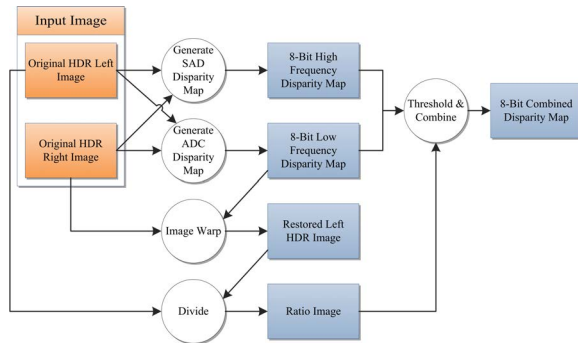

Figure 3: IPDC disparity generation

3.5. Motion Compensation (MC) Method

The final method we propose is based on the observation that two frames differ only by the camera position which is not dissimilar from the temporal motion between subsequent frames in videos. This fact motivated investigation in using standard video coders to compress the SHDR image. Two views of a stereo pair are treated as two consecutive frames in the video which is then processed. Any video coder or camera motion compensation can be used here. For our tests we used H.264 codec.

The pipeline for this method is presented in Figure 4. Encoding starts by compressing left and right images separately using JPEG-HDR. The tone-mapped images are merged to create a two frame video, which is processed by the video encoder. Due to backwards compatibility we are only interested in the second frame data of the compressed video which gets extracted. In the H.264 example the second frame corresponds to the predicted frame (p-frame) data. The first frame is left stored in the original JPEG-HDR format and is not MC encoded. The extracted data depends on the compression quality used but for moderate compression values it is rather small (see Table 2) and can fit in additional JPEG subbands together with the JPEG-HDR ratio images for both frames.

The decoding process is the inverse of the coding pipeline (refer to Figure 5). A video file consisting of two frames is generated using the tone-mapped image which is appended to the second frame data. This video is then decoded which provides the second image for the reconstruction using JPEG-HDR. A standard viewer would only open the stored JPEG and a JPEG-HDR viewer will open only the HDR image of the stored view.

Figure 4: MC encoding

Figure 5: MC decoding

4. Results and Analysis

We used 19 SHDR images to evaluate and test proposed methods all captured at an HD resolution of 1920×1080. It was desired to have variety of scenes which differ in dynamic range, depth, frequency, amount of noise and contrast. Four scenes were computer generated. Tone-mapped versions of the left view together with dynamic range and disparity range are shown in Figure 7. When coding JPEG-HDR images we wanted to preserve the quality so it was set

to 95. The disparity maps for IPD method were generated using default settings suggested by authors. SAD method used in IPDC had window set to size 3.

A compatibility table of all methods is represented in Table 1. Please note that SBS method is either compatible with traditional LDR viewer or LDR stereo viewer depending on how the second frame data is formatted.

The difference between the output of each method for one of the tested images is shown in Figure 6. Only one image of the pair is shown. For IPD, IPDC and MC methods this is the disparity or motion reconstructed one, as the other is equivalent to the image obtained using SBS method. We chose a small segment (highlighted with the yellow rectangle) from Scene 19 to examine the differences in more detail. All images are tonemapped so some of the artifacts might be as a result of the tonemapping. As expected, the SBS method is the one which is the most similar to the original image as there are no prominent artifacts. The loss of resolution is visible in the HSBS image which is blurrier than the rest. Mistakes made during image warping due to the occlusion are present in IPD method especially in the case of steering wheel. IPDC manages to fix this problem to an extent but small mistakes are noticeable towards the top of the steering wheel which has gray values instead of the original yellowish ones. MC method looks very similar to the original but on closer inspection some "blocky" artifacts due to MPEG compression are visible.

In order to evaluate the performance of the presented methods the following quantities were measured: file size of the compressed SHDR image (in kilobytes), normalised root-mean-square error (NRMSE) and peak signal-to-noise ratio (PSNR) between the decoded images and the original HDR images. NRMSE and PSNR were measured for left and right views separately and than averaged. File sizes are shown in Table 2. The last table column contains the sizes of a single LDR image of the same scene which is JPEG encoded using the same quality settings as the proposed methods. NRMSE, expressed as a percentage, is presented in Table 3 and PSNR is shown in Table 4. Average values for all of the scenes are shown in the last row.

Table 5 provides a summary of the results. The compression ratio gives the average compression ratio compared to a raw HD stereo image. The "quality × compression" measure is a multiplication of the average image size for each method multiplied by the average NRMSE. This value is presented to give an idea of what gives the best "bang for the buck" but should not be taken as a definitive measure as the different methods have distinct qualities. For "quality × compression" the smaller values are considered better.

On the whole it is the MC method that achieves the best overall "quality × compression", it also has the added advantage that it is backward compatible with JPEG-HDR and JPEG. The HSBS achieves the second best overall score, this is primarily due to it having the highest compression

ratio but this comes at the expense of quality, for which it is second from the bottom. It is backward compatible with LDR stereo JPEG but not fully compatible with JPEG-HDR and traditional JPEG; the image shown on a traditional JPEG viewer would be both images side by side. The SBS method is third overall and is backward compatible with all possible formats, however it is the largest in size, which may be to much of a prohibitive obstacle, potentially hampering the uptake of SHDR. The disparity methods IPDC and IPD are second from last and last respectively. However, the results depend on the quality of the disparity maps produced, and as these may become better, as this is an active research area the scope of such disparity based methods is likely to improve. In addition these methods are backward compatible with traditional JPEG and HDR JPEG and produce relatively small image sizes not much larger than JPEG-HDR images.

Table 1: Compatibility. * depending on storage either Mono or Stereo LDR - not both

	SBS	HSBS	IPD	IPDC	MC
Mono LDR	✓*	✗	✓	✓	✓
Stereo LDR	✓*	✓	✗	✗	✗
Mono HDR	✓	✓	✓	✓	✓
Stereo HDR	✓	✓	✓	✓	✓

Table 2: Sizes of compressed images (in KB) and single LDR image of the same scene shown for comparison.

Img	SBS	HSBS	IPD	IPDC	MC	LDR
1	1087	561	601	717	683	476
2	1088	558	613	728	671	482
3	868	459	488	616	544	370
4	897	493	494	589	557	384
5	891	506	474	514	549	383
6	1697	899	956	1392	1104	794
7	1021	512	572	704	616	449
8	784	417	436	458	484	328
9	1648	837	890	1119	977	769
10	1215	617	629	682	707	544
11	1087	549	563	577	639	487
12	436	267	245	263	285	168
13	1326	689	728	877	803	601
14	1004	530	549	589	617	439
15	888	455	549	656	543	362
16	1457	744	786	1024	913	685
17	1456	793	849	1132	883	676
18	954	491	529	603	590	412
19	1026	537	569	617	640	449
Avg	1096	574	606	729	674	487

(1) Original scene 19　　　　　(2) SBS　　　　　(3) HSBS

(4) IPD　　　　　(5) IPDC　　　　　(6) MC

Figure 6: Decoded image for each method

Table 3: NRMSE measurements (%)

Img	SBS	HSBS	IPD	IPDC	MC
1	0.12	0.23	0.40	0.30	0.13
2	0.11	0.24	0.25	0.18	0.12
3	0.08	0.11	0.20	0.19	0.09
4	0.08	0.17	0.11	0.11	0.09
5	0.37	0.41	0.56	0.48	0.60
6	0.32	0.59	0.77	0.50	0.44
7	0.08	0.26	0.23	0.16	0.10
8	0.17	0.40	0.30	0.25	0.21
9	0.43	0.71	0.76	0.61	0.50
10	0.56	0.89	0.97	0.67	0.73
11	0.94	1.23	1.06	1.04	1.37
12	0.18	0.29	0.21	0.20	0.30
13	0.21	0.42	0.44	0.38	0.26
14	0.18	0.25	0.35	0.24	0.18
15	0.30	0.34	0.57	0.39	0.32
16	0.26	0.36	0.49	0.42	0.31
17	0.37	0.75	1.27	0.72	0.53
18	0.08	0.15	0.13	0.12	0.09
19	0.20	0.30	0.32	0.28	0.21
Average	0.27	0.43	0.49	0.38	0.35

Table 4: PSNR measurements

Img	SBS	HSBS	IPD	IPDC	MC
1	58	53	51	53	58
2	59	52	54	56	59
3	62	59	55	56	61
4	62	55	59	59	61
5	49	48	46	47	45
6	50	45	44	47	47
7	62	52	55	57	60
8	56	48	52	53	54
9	47	43	43	45	46
10	45	41	41	44	43
11	41	38	40	40	38
12	55	51	54	54	51
13	54	48	49	49	52
14	55	52	50	53	55
15	51	49	46	48	50
16	52	49	47	48	50
17	49	42	41	44	46
18	62	56	59	59	61
19	54	51	51	51	54
Average	54	49	49	51	52

5. Conclusion and Future Work

In this paper we have introduced SHDR as a combination of stereoscopic and high dynamic range imaging. We have proposed a number of SHDR compression methods based on the JPEG standard and the JPEG HDR compression method, and provided an analysis of such methods against ground truth data. The MC, SBS, IPDC and IPD methods are all backward compatible with traditional JPEG and JPEG-HDR, and while our results demonstrate a better overall performance for MC, the disparity based methods should not be discounted as they may improve with time as disparity methods improve. The SBS and HSBS methods are backward compatible with stereo JPEG and could be used on occasions

Table 5: Results Summary. Compression Ratio (CR) and Quality × Compression (Q×C). NRMSE is used for quality and for Q×C lower figures are better.

	SBS	HSBS	IPD	IPDC	MC
CR	44	85	80	67	72
Q×C	291	245	300	278	234

where this is sufficient. We believe the compatibility aspect to be crucial for SHDR in its infancy as it may improve uptake and allow professional and amateur photographers to capture SHDR without worrying about potential audiences, while reaping the benefits for those who support the technology and for future generations.

Future work will look into investigating further potential compression gains that could be achieved, perhaps by ignoring the backward compatibility constraint that we have imposed; although this may be better to use once the technology is widely adopted and if great improvements are possible. Furthermore, we have evaluated our methods with NRMSE and PSNR, since no perceptual metric exists for evaluating SHDR, future work will include user studies; in particular it will be interesting to see if the disparity based methods would improve as a result of binocular fusion of the two images which is currently difficult to measure without human participation. The next frontier for SHDR storage would be to develop an efficient method for storing video SHDR and some form of compatibility with MPEG in the mould of MPEG-HDR [MEM06]. It is envisaged that similar methods to those presented above may also be used for multiview imaging techniques also.

SHDR is still an emerging imaging method and a great deal of research is required in order to make it fully functional. We hope that other researchers will contribute to improving all aspects of SHDR from capture to delivery so that SHDR may become a future imaging standard.

Acknowledgments

References

[BADC11] BANTERLE F., ARTUSI A., DEBATTISTA K., CHALMERS A.: *Advanced High Dynamic Range Imaging: Theory and Practice*. A K Peters/CRC Press, 2011. 1, 2

[CBB*09] CHALMERS A., BONNET G., BANTERLE F., DUBLA P., DEBATTISTA K., ARTUSI A., MOIR C.: High-dynamic-range video solution. In *ACM SIGGRAPH ASIA 2009 Art Gallery & Emerging Technologies: Adaptation on - SIGGRAPH ASIA '09* (New York, New York, USA, Dec. 2009), ACM Press, p. 71. 2

[CS09] CYGANEK B., SIEBERT J. P.: *An Introduction to 3D Computer Vision Techniques and Algorithms*. John Wiley & Sons, Ltd, Chichester, UK, Jan. 2009. 4

[KBH04] KAINZ F., BOGART R., HESS D.: The openexr image file format. *GPU Gems: Programming Techniques, Tips and Tricks for Real-Time Graphics, R. Fernando, Ed. Pearson Higher Education* (2004). 2

[Lar98] LARSON G.: Logluv encoding for full-gamut, high-dynamic range images. *Journal of Graphics Tools 3* (1998), 15–32. 2

[MCS*11] MEI X., CUI C., SUN X., ZHOU M., WANG Q., WANG H.: On Building an Accurate Stereo Matching System on Graphics Hardware. In *Computer Vision Workshops (ICCV Workshops)* (2011), pp. 467—-474. 3

[MEM06] MANTIUK R., EFREMOV A., MYSZKOWSKI K.: Backward compatible high dynamic range MPEG video compression. *ACM Transactions on 1*, 212 (2006), 713–723. 7

[OA07] OKUDA M., ADAMI N.: Two-layer coding algorithm for high dynamic range images based on luminance compensation. *Journal of Visual Communication and Image Representation 18*, 5 (Oct. 2007), 377–386. 2

[RHM11] REMPEL A. G., HEIDRICH W., MANTIUK R.: The role of contrast in the perceived depth of monocular imagery. In *Proceedings of the ACM SIGGRAPH Symposium on Applied Perception in Graphics and Visualization - APGV '11* (New York, New York, USA, Aug. 2011), ACM Press, p. 115. 2

[SGJ92] SERVOS P., GOODALE M., JAKOBSON L.: The role of binocular vision in prehension: a kinematic analysis. *Vision Research 32*, 8 (1992), 1513–1521. 1

[SHS*04] SEETZEN H., HEIDRICH W., STUERZLINGER W., WARD G., WHITEHEAD L., TRENTACOSTE M., GHOSH A., VOROZCOVS A.: High dynamic range display systems. *ACM Transactions on Graphics (TOG) 23*, 3 (2004), 760—-768. 2

[SMS07] SMOLIC A., MUELLER K., STEFANOSKI N.: Coding Algorithms for 3DTV - A Survey. *Circuits and Systems for Video Technology, IEEE Transactions on 17*, 11 (2007), 1606–1621. 2

[SS97] SIRAGUSA J., SWIFT D. C.: General purpose stereoscopic data descriptor. *VRex, Inc.* (1997). 2

[SSZ01] SCHARSTEIN D., SZELISKI R., ZABIH R.: A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *Proceedings IEEE Workshop on Stereo and Multi-Baseline Vision*, 1 (2001), 131–140. 3

[TKTS11] TOCCI M., KISER C., TOCCI N., SEN P.: A Versatile HDR Video Production System. *ACM Transactions on Graphics (TOG) 30*, 4 (2011), 41–49. 2

[War91] WARD G.: Real pixels. *Graphics Gems II* (1991), 80–83. 2

[War94] WARD G. J.: The RADIANCE lighting simulation and rendering system. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques - SIGGRAPH '94* (New York, New York, USA, July 1994), ACM Press, pp. 459–472. 2

[War06] WARD G.: JPEG-HDR: A backwards-compatible, high dynamic range extension to JPEG. *ACM SIGGRAPH 2006 Courses* (2006), 8. 1, 2

[WM08] WARE C., MITCHELL P.: Visualizing graphs in three dimensions. *ACM Transactions on Applied Perception 5*, 1 (Jan. 2008), 1–15. 1

(1) DR: 5.50 - Disparity: 58

(2) DR: 5.58 - Disparity: 36

(3) DR: 5.67 - Disparity: 28

(4) DR: 5.16 - Disparity: 31

(5) DR: 8.82 - Disparity: 40

(6) DR: 3.08 - Disparity: 72

(7) DR: 5.50 - Disparity: 30

(8) DR: 4.24 - Disparity: 34

(9) DR: 3.21 - Disparity: 20

(10) DR: 5.02 - Disparity: 28

(11) DR: 3.56 - Disparity: 22

(12) DR: 6.10 - Disparity: 22

(13) DR: 3.55 - Disparity: 38

(14) DR: 5.63 - Disparity: 34

(15) DR: 4.04 - Disparity: 64

(16) DR: 3.09 - Disparity: 34

(17) DR: 2.90 - Disparity: 20

(18) DR: 5.38 - Disparity: 36

Figure 7: SHDR Scenes (tone-mapped) with Dynamic Range (Weber Contrast) and Disparity Range

[WS04] WARD G., SIMMONS M.: Subband encoding of high dynamic range imagery. In *Proceedings of the 1st Symposium on Applied perception in graphics and visualization - APGV '04* (New York, New York, USA, Aug. 2004), ACM Press, p. 83. 2

[XPH05] XU R., PATTANAIK S., HUGHES C.: High-dynamic-range still-image encoding in jpeg 2000. *Computer Graphics and Applications, IEEE 25*, 6 (2005), 57–64. 2