High Dynamic Range Image Formats

Bernhard Holzer
Matr.Nr. 0326825

Institute of Computer Graphics & Algorithms
TU Vienna

Abstract

HDR-image formats are able to encode a much greater range of colors and light intensities than standard 24-bit formats do. The latter were designed to be conveniently displayed with 20 year old monitor technology, i.e. CRT monitors. Since high-fidelity digital imaging has become more and more important in the last years, so has the need for the use of HDRI-technology. This document gives an overview of available HDRI-formats, their history and applications. The most relevant formats are compared and their advantages, disadvantages, capabilities and features are discussed.

1. Introduction

Naturally, real-world scenes can include both very low and very high light intensities at the same time. This difference between the brightest and the darkest parts of an image is called *contrast ratio* or *dynamic range*. The perceived brightness of a scene is also referred as luminance, measured in candela per square meters [1]. The sun for example has a luminance level of a about 10^8 cd/m². An outside scene, only lit by starlight, has an intensity of about 10^{-3} cd/m² [2].

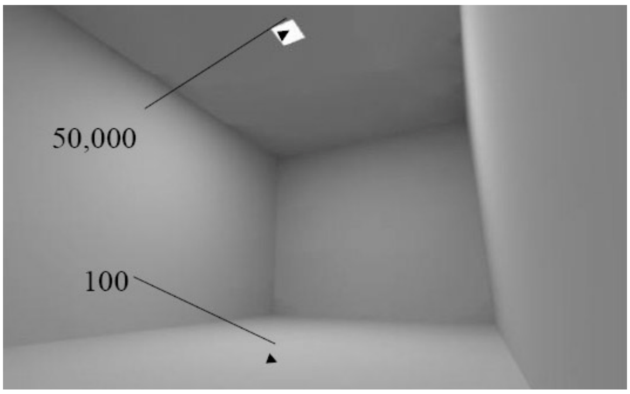

Figure 1: Concentrated light sources usually generate high dynamic range. The picture shows an example of light intensities in a rendered room [3].

The term is not clearly defined, but speaking of *high dynamic range* (HDR) usually means to cover more than about 4 orders of magnitude. The range of luminance levels which conventional 24 bit image formats can handle, i.e. 256, is called *low dynamic range*, where a slightly better performance is often referred as *medium dynamic range*.

1.1 Human perception

The human visual system is capable of perceiving roughly 4 orders of magnitude of intensities at any one moment. It can adjust another 6 orders up and down through a process called adaptation. This process does not work instantaneous and may take some minutes, for instance, in the case of entering a dark environment [4].

However, only covering that enormous range of intensities does not mean that we can see equally well at all levels, since there are two kinds of photoreceptors which have different sensitivities.

In low-light-environments, rods are capable of detecting small luminance differences, but their ability to distinguish color is poor as their visual acuity. This is called *scotopic vision*.

The three types of cones on the other hand, are responsible for perceiving a sharp vision of colors under well-lit conditions, which is called *photopic vision*. The overlapping region between the scotopic and photopic range is called *mesopic*. The cones have a handicap in differing luminance levels, which have to be rather large to be detectible [5].

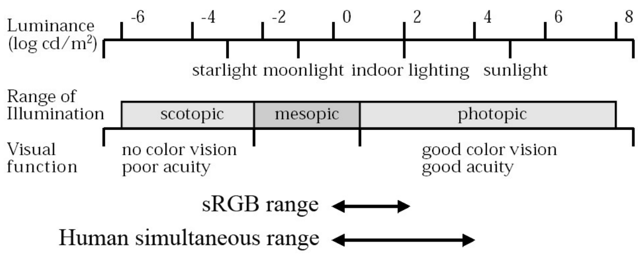

Figure 2: The range of luminances a the natural environment and associated visual parameters [6].

These differences of distinguishable intensities were evaluated in psychophysical studies. They have shown that over a large range of intensity, the minimum distinguishable difference is constant at about 1% [1]. This is called the *visible threshold* or *just-noticeable difference*. It was further evaluated, that the human vision has a response to light intensities that follows a logarithmic function.

1.2 Color

The three types of cones mentioned before are each responsible for the perception of either red, green or blue. A color impression results from the stimulation of the cones, according to the red, green and blue fractions of the received light.

This means, that by independently controlling these three color primaries, the eye can be fooled to see nearly any spectral distribution. However, not all perceivable colors can be mixed this way, since some are outside of the range that can be mixed with red, green and blue (The range of realizable colors with a given set of primaries is generally called *gamut*). This is compensated by using negative values of one or more primaries.

Based on these facts, the CIE (Commission Internationale de l´Eclairage) developed their XYZ color space in 1931. It contains any visible color, coded in strictly positive values, and is used as basis for translating colors from one color space into another.

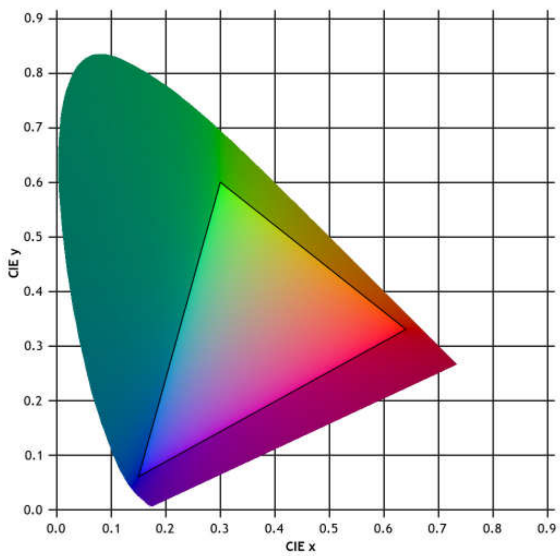

Figure 3: The CIE xy chromaticity diagram.

Instead of representing colors with 3 axes for each primary, the CIE XYZ color space uses a different concept: The x and y values define the chromaticity of a color, the z value stands for the luminance. The two-dimensional CIE xy chromaticity diagram is shown in figure 3, its horseshoe-shape represents the gamut of human vision. The triangle inside defines the boundary of the sRGB color space, which is used by CRT monitors. In this picture, the colors outside the triangle are therefore only stand-in colors: they cannot be displayed on a CRT or be stored in an image using sRGB.

The color distances in the CIE xy chromaticity diagram are not perceptually uniform. This is corrected in the CIE (u',v') color diagram, where equal distances between color coordinates now appear as equal color differences to a human observer.

1.3 Brightness encoding

There is a non-linear relationship between the voltages applied to a CRT electron gun and the resulting visible radiant energy (cd/m²). This can be expressed with a power law function, where the numerical value of the exponent is colloquially known as *gamma* (γ). Although the sRGB standard intentionally deviates from this value to a target γ of 2.2, typical CRT devices follow a power relation corresponding to a γ value between 2.4 and 2.8, which results in a slight contrast for displayed images [8].

$$I_{out} = K \cdot v^{\gamma}$$

Formula 1: output intensity is equal to some constant factor *K*, times the input voltage *v*, raised to a constant power γ.

By coincidence, the non-linearity of a CRT is effectively the inverse of the lightness sensitivity of human vision: if a transfer function is applied to code a signal to take advantage of the properties of lightness perception, the coding will be inverted by a CRT [7].

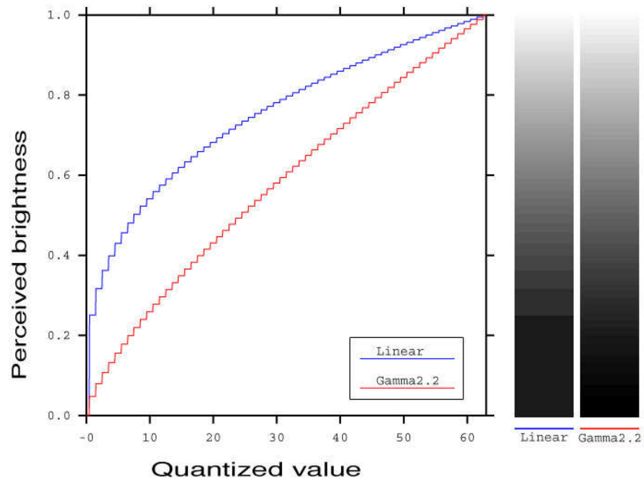

Figure 4: This image shows the perceived brightness using gamma and linear encoding [8]. A 6-bit resolution was used here to make the quantization steps more visible.

It is in the nature of the process of color digitalisation, that there is no escape from quantization errors. So one would like to keep them as low a possible or at least not noticeable in terms of human perception.

Like shown in figure 4, a linear quantization of luminance comes in hand with more visible steps in darker regions, as brightness is perceived non-linear. The encoding with a gamma function seems to be a good solution in this case. But dealing with high dynamic range usually means luminance ratios of several thousands or millions to one, and only improving the bit-resolution of the gamma-encoding does not provide evenly distributed steps. This is because it cannot be assumed that the viewer is adapted to a certain level of luminance, and the quantization error continues to increase with the decrease of luminance (figure 5).

A better way to encode large ranges of luminances is to apply a logarithmic encoding. This method provides an constant relative error, but comes at the price of an maximum and minimum representable intensity value.

$$I_{out} = I_{min} \cdot \left[\frac{I_{max}}{I_{min}} \right]^{v}$$

Formula 2: Log encoding uses this formula.

Another alternative is the floating point representation. This encoding has no constant step sizes either, but its relative error follows a slight sawtooth pattern [8].

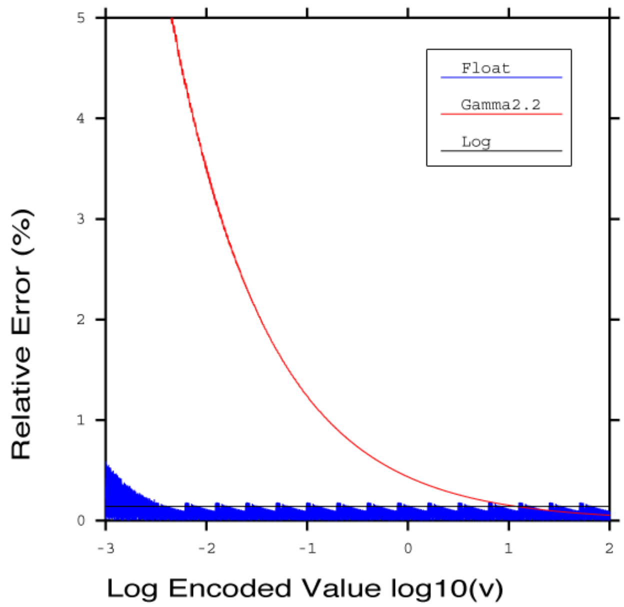

Figure 5: The relative error is plotted against the $\log_{10}$ of luminance values for three encoding methods [8].

1.4 Tone mapping

"*Tone mapping is a technique used to approximate the appearance of high dynamic range images in media with a more limited dynamic range. Print-outs, CRT or LCD monitors, and projectors all have a very limited dynamic range. Essentially, tone mapping addresses the problem of strong contrast reduction from the scene values (radiance) to the displayable range while preserving the image details and color appearance important to appreciate the original scene content.*"[9]

There are many techniques to achieve this aim, mainly divided into four categories of tone-mapping operators:

- **Local operators**
  The adaptation level is computed for each pixel individually, considering the surrounding neighbourhood pixels
- **Global operators**
  The entire image is used as neighbourhood for each pixel. This approach is computationally more efficient and generally suitable for images with medium dynamic range.
- **Frequency domain operators**
  The reduction of the dynamic range is based on the spatial frequency of image regions.
- **Gradient domain operators**
  A derivative of the image is modified.

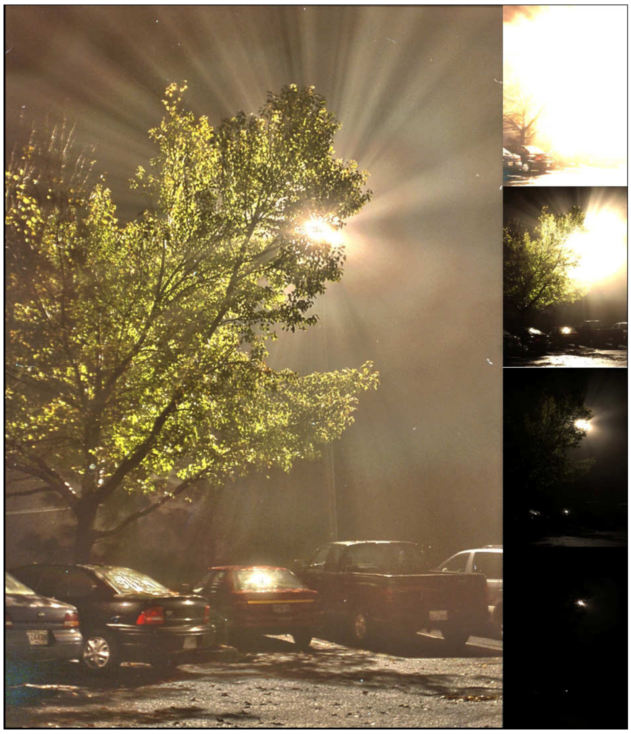

Figure 6: Tone-mapped image. On the right: linear mappings at different exposures [10].

2. HDRI-Formats

Conventional 24-bit sRGB image formats [11], were developed with their output device in mind. This concept is called *output referred standard*, since such images only contain colors, which their target output device is capable to show. This is an advantage insofar, that there is no manipulation required before the image can be displayed. Hence, the gamut covers all it needs, no resources are wasted for undisplayable colors.

On the other side, this may be valid for today's display technology, but future possibilities are not considered. Of course no one can say what display technology may come in future, but its much likely that it will target on human visual abilities. Following this thought, a second approach aims to store at least as much information, the human visual system can resolve. This philosophy is called *scene referred standard*, because it intends to save the original captured scene values as closely as possible.

Generally HDR images are encoded with 16-bit/channel or higher, to preserve as much accuracy as possible and to avoid visible quantization, i.e. banding artefacts.

The inverse is often not the case: images with higher bit depth do not necessarily have a high dynamic range. Often the higher bit depth is simply used to lower image noise and increase accuracy [12].

2.1 Radiance RGBE (.hdr)

This format is probably the most widespread in the HDR-imaging community. It was first used as part of *Radiance*, a physically-based rendering system developed by Greg Ward at the Lawrence Berkeley National Laboratory in the late 1980s.

Because Radiance was designed to compute photometric quantities, it needed an adequate output format for saving all of its computed results.

The RGBE format uses 8 bit for each RGB component and another 8 bit for a shared exponent and uses run-length encoding as compression algorithm.

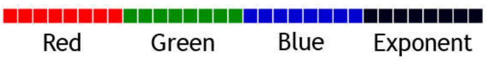

Figure 7: Bit distribution of Radiance RGBE/XYZE.

These 32 bit per pixel cover luminance values between 10^{38} and 10^{-38} cd/m², which is much more than anyone could perceive (or even use for color representation). This vast dynamic range of 76 orders of magnitude comes at expense of accuracy. The precision of about 1%, is just sufficient for surpassing human perception.

The next shortcoming of this encoding is, that it does only support positive RGB-values. Therefore it cannot represent some colors of the visible gamut [13].

This is fixed in a second variant of the format, called XYZE. Instead of the "real" primaries of RGB it uses "imaginary" primaries of the CIE XYZ-color space, extending the range of color to the entire visible gamut [8].

2.2 Pixar Log (.tif)

As a result of their need, to save rendered images for the use in film recording, Pixar (formerly a part of Lucasfilm) also developed one of the first HDR-image formats.

Since film has a greater dynamic range than a standard 24-bit/pixel image and a log-response instead of a gamma function, a logarithmic encoding for RGB values was implemented. With its 11 bits per channel, this format is able to cover a dynamic range of about 3.8 orders of magnitude with a precision of 0.4%, which is marginal for image processing. Like the Radiance RGBE format, the Pixar Log encoding suffers from a restricted color gamut. It supports ZIP lossless entropy compression and is implemented in Sam Leffler's freely distributed TIFF library.

2.3 SGI LogLuv (.tif)

The LogLuv tiff format was developed by Greg Ward at SGI, to create a more efficient, in other words perceptual-based, encoding than the Radiance RGBE.

There are 2 different variants of the LogLuv format, both store luminance logarithmically and CIE (u', v') chrominance values linear in separate channels.

The first variant uses 24 bits/pixel, where 10 bits are used to encode luminance logarithmically. The remaining 14 bits represent the chromaticity, based on a lookup of CIE (u',v') values (which is shown in figure 9). This variant is able to cover a dynamic range of 4.8 orders of magnitude in uniform 1.1% steps, but causes visible artefacts due to the limited range of luminances that it can represent [15] .

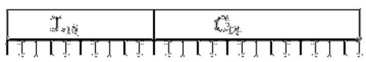

Figure 8: Bit distribution of the 24-bit LogLuv encoding [13].

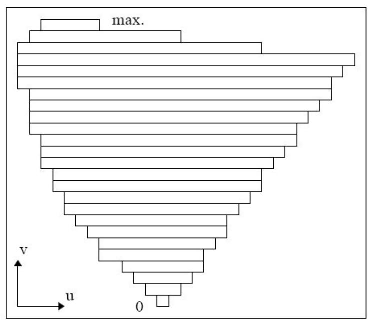

Figure 9: Lookup of CIE (u,v) values [13].

The second variant requires a total of 32 bits/pixel. 8 bits are used each for the CIE u and v coordinates. This is sufficiently accurate, so a lookup is not needed. The luminance channel in this variant has a sign bit, so it can represent positive and negative log encoded values with the remaining 15 bits. This results in a dynamic range of 38 orders of magnitude in impercievable 0.3% steps and covers the entire visible gamut.

In contrast to its 24 bpp counterpart, the 32bit LogLuv encoding submits very well to run-length compression and therefore often produces smaller files [13]. Both variants are also a part of Lefflers's TIFF library.

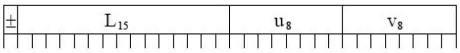

Figure 10: Bit distribution of the 32-bit LogLuv encoding.

2.4 ILM OpenEXR (.exr)

The Extended Range Format, developed by Industrial Light & Magic, was created for special effects rendering and compositing. It was made freely available to the public in 2002.

It supports 32-bit floating point or unsigned integer values per component, but its primary form is a 16-bit floating point per RGB-primary encoding. This so-called *half*-format is compatible with the *half* data type of NVIDIA's CG graphics language, which is widely supported in today's graphics hardware [15].

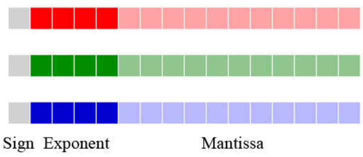

Figure 11: Bit breakdown of the OpenEXR Half encoding.

The representation of each color primary is divided into 1 sign bit, 5 exponent- and 10 mantissa-bits, so this format is also called *S5E10*. In this main variant, the OpenEXR format is capable of encoding 10.7 orders of magnitude with a quantisation step of about 0.1%. Due to the fact, that it is not restricted in representing only positive RGB values, it covers the entire perceivable color range.

The OpenEXR format currently supports a number of compression methods, and the possibility for the extension with compression codecs. In addition, it has wide range of features like multichannel support, annotation with arbitrary number of attributes (like color balance information from a camera) [16] and the unique feature of a definition of an active region of the image, called *display window*. This is very useful for image post processing, with wide filter kernels for example.

2.5 96-bit floating point TIFF / PFM

The TIFF specification and the similar Portable Floatmap format, (a floating-point version of the Portable Pixmap format), both provide a 32 bit per RGB channel floating point encoding, which results in impressive values of dynamic range and precision.

However, the downside of covering 79 orders of magnitude in 0.000003% steps is the space it requires for storage: Due to the extreme precision, about the last half of the mantissa stores basically just random noise, as image sources usually don’t generate data with this high accuracy. Not only the included noise won’t contain useful information, floating point data in general does not apply very well to compression [1, 8].

But nonetheless, this 96 Bit /pixel still have a right of existence. As lossless intermediate representation in image-processing for example, where high precision data is needed between steps of calculations or as non-color pixels such as remote sensing data.

Figure 12: Bit breakdown of the 96-bit floating point TIFF / PFM format.

2.6 Microsoft/HP scRGB

A new IEC standard, extending the sRGB, was recently developed by Microsoft and Hewlett-Packard. This currently not really wide supported format is called scRGB and comes in 3 different variants:

The first uses 16 linear encoded bits per RGB channel. Due to the log response of human eyes mentioned before, a linear encoding spends a lot of its precision unnecessarily at high levels of luminance, whereas better distinguishable darker levels are lacking of precision. A another improvement over the sRGB is that the scRGB generally allows negative primaries and therefore covers the visible color gamut. With its 48 bits per pixel and variable precision, this variant achieves about 3.5 orders of magnitude.

The other two substandards of scRGB employ 12 bits per primary either as RGB or YCC, both gamma encoded. These variants cover 3.2 orders of magnitude, also in variable steps. At highest luminance values, the representable colors of the scYCC version are not white, which leads to the strange results shown in figure 13 [8].

Figure 13: The scYCC-nl encoding shows strange color behavior at maximum luminance values in the right.

2.7 RAW/ DNG

Higher-level digital cameras often offer an alternative to saving photos as JPEG. When saving as JPEG, the color information from the CCD-sensor has to be filtered, converted to RGB, and gamma-corrected. White balancing, sharpening and other image corrections are also applied [17].

As the name suggests, the RAW-format contains unprocessed or minimally processed data direct from the sensor. Additionally, RAW files also contain some metadata like camera model, shutter speed, etc. Because a CCD captures 10 or 12 bits per pixel, a RAW has a greater dynamic range than an 8 bit image. But the main advantage of shooting photos in RAW format is, that the filtering and other processes mentioned above can be done later by the user. This requires special software, but it provides much more flexibility. It and allows the user to tweak his photos after the shoot as he likes, less worrying about the camera settings.

One drawback is, that RAW-files are usually uncompressed [18] and therefore require much more space in memory. Furthermore, there is no single RAW-format, but many manufacturer specific and even camera-specific implementations.

There have been some efforts to create a standard for RAW-files, including Portable Floatmap [18], the Kodak ERI-JPEG, the Adobe Digital Negative (see description below) and the OpenRAW Working Group (which advocate the open documentation of digital camera RAW files).

The Digital Negative (DNG) was published by Adobe in 2004, and is an attempt to unify the current proprietary RAW formats by providing a wrapper for the existing proprietary RAW implementations [19].

DNG is an extension of the TIFF 6.0 format, and is compatible with the TIFF-EP standard [20]. It supports 8, 16 and 32 bits per sample and provides either uncompressed, lossless Huffman JPEG or DCT Baseline JPEG compression. It provides several meta data tags such as EXIF but also proprietary or private tags/markers.

2.8 FITS

This is a file format, which has been in use in the astronomical community for over 2 decades. Although it can represent a wide range of non-image-data, the earliest and most commonly use is as an image. It supports several integer and floating point encodings, such as IEEE 754 single and double precision and is capable to store x-ray and infrared exposures in the same file [21].

2.9 DDS

DDS stands for DirectDraw Surface, and is the preferred file format for storing textures in the development of DirectX Games. It is capable of handling various encodings, including 16-bit (half) and 32-bit (single) floating point values. DDS supports alpha, cube maps both with and without mipmap levels, uncompressed and compressed pixel formats and DXTn Texture compression algorithms [22].

2.10 Kodak Cineon/DPX

Kodak's Cineon is also to be counted to the earlier HDR-formats. It was designed as format for storing single frames of motion picture film from film recorders and scanners, and has found a large distribution upon this field. The format provides 10-bits per channel with a logarithmic encoding (as described previously, film has a logarithmic response to luminance).

A newer modification of Cineon is the DPX format, meaning Digital Picture Exchange. It supports bit depths of 1, 8, 10, 12, 16, 32, and 64 bit [23], both RLE- or uncompressed values, logarithmic, linear and various other encodings. There are also useful headers for the use in Film and Video production.

Although DPX is a lot more flexible, the older Cineon format is still in widespread use [24].

2.11 JPEG-HDR/ERI-JPEG/JPEG 2000

With minimum loss of image quality even under repeated processing steps in mind, the HDR-standards described so far often use a lossless compression (if any) to reach this goal. But there are fields of application, where less storage cost is preferred over the possibility of multiple editing steps, such as HDR photography and HDR Video. In this case, a HDR format using a lossy compression would be more convenient, as it achieves much higher compression ratios and therefore smaller files.

One solution to this aim was developed by Ward and Simmons [29]. Normally, a tone-mapping process comes in hand with loss of information. In their proposed subband-method, a HDR image is tone-mapped (using conventional techniques) into the 8-bit JPEG-format, but retaining restorative HDR-data. This information, encoded as an 8-bit greyscale ratio image plus some additional conversion formulas, is stored as metadata into the JPEG-file. This information is enough to allow a later reversal of the mapping process and so restoring the original scene-referred data.

An important advantage of this approach is the full backward-compatibility with standard-JPEGs: a conventional jpeg decoder would ignore the unknown metadata and display the tone mapped LDR image. A decoder capable of the extension, could retrieve the hidden ratio image and obtain the entire HDR-image [15].

A similar method was introduced by Kodak's ERI-JPEG (meaning Extended Range Imaging JPEG) in 2002. It was meant as alternative to RAW images [18].

A residual image, which represents the arithmetic difference between the scene-referred input- ERIMM RGB color space and the recorded sRGB image, is also stored in a subband to retrieve the original color information [23].

Another approach of Xu et al make use of the JPEG 2000 encoding. First, pixels from original floating-point values in the logarithm domain are mapped into integer values, then the image is encoded using the JPEG 2000 codec. This method is able to apply a compression of an HDR image from very low bit rates to visually lossless [30].

3. Comparison of HDRI-encodings

In this section, HDR-formats are compared in terms of archiving them for future use. Since the capabilities of future displays are unknown, it would be reasonable to aim for the highest quality possible, but also storage costs should be kept in mind. Table 1 lists the most interesting candidates discussed in detail in the section above.

Encoding	Covers Gamut	bpp	Dynamic Range ($\log_{10}$)	Accuracy
sRGB	No	24	1.6	Variable
Pixar Log	No	33	3.8	0.4%
RGBE	No	32	76	1%
XYZE	Yes	32	76	1%
LogLuv 24	Yes	24	4.8	1.1%
LogLuv 32	Yes	32	38	0.3%
OpenEXR	Yes	48	10.7	0.1%
scRGB	Yes	48	3.5	Variable
scRGB-nl	Yes	36	3.2	Variable
scYCC-nl	Yes	36	3.2	Variable
JPEG-HDR	Can	1-7	9.5	Variable
TIFF96	Yes	96	79	0.000003%

Table 1: This table is a summary of the HDR formats and their properties, mentioned in the previous section. The sRGB format is only shown as a baseline.

We are interested in formats that do surpass the human vision, i.e. cover the visible gamut and at least 4 orders of magnitude, and does have a decent uniform accuracy. For reasons of storage efficiency, the (for most tasks) un-necessarily precise 96 bpp TIFF is excluded.

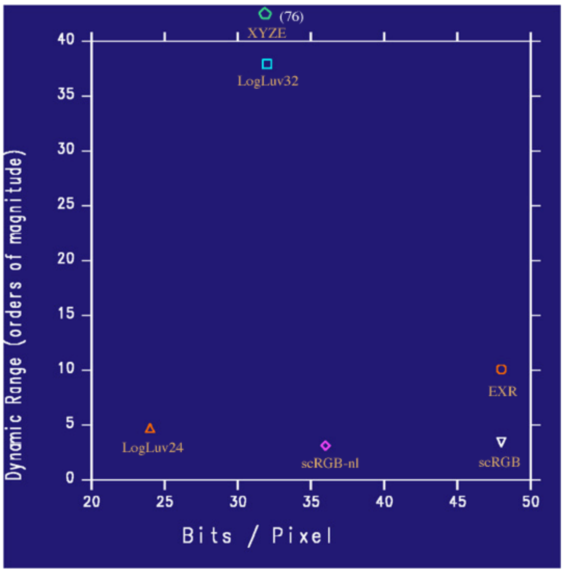

Figure 14: Cost (bits/pixel) vs. benefit (dynamic range) of full-gamut formats [8].

For the comparison itself the CIE ΔE* 1994 color-difference metric is used. Because this metric assumes a global white adaptation value, which does not make sense for HDR imagery, it was modified for this test by applying a *local reference white* to consider local adaptation. In this approach, the local reference white is defined in this comparison by the brightest value in a 50 pixel radius around the current tested pixel.

For an intuitive visualization of encoding shortcomings, a synthetic test pattern was converted from 96-bit TIFF to other encodings and then compared pixel-by-pixel. Each peak represents one revolution through the visible gamut and spans 1 order of magnitude, 8 orders in total. The grey regions contain random noise at each luminance level.

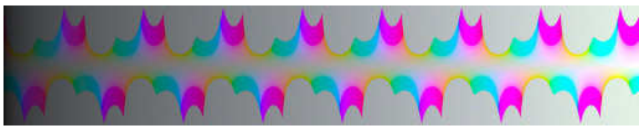

Figure 15: This test pattern shows a computer generated spiral slice through the visible gamut, covering 8 orders of magnitude.

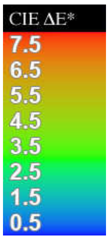

Figure 16: This scale is used by all the following false color plots of the test pattern (Figure 15). CIE ΔE* values over 2 are considered visible, values over 5 are evident.

A LDR image is also included in the test, to demonstrate the inferiority of sRGB images when it comes to the representation of scene referred data (Figure 17). Likewise, the Radiance RGBE format also shows its shortcomings in representing highly saturated colors, because it cannot contain colors out of the sRGB gamut (Figure 18).

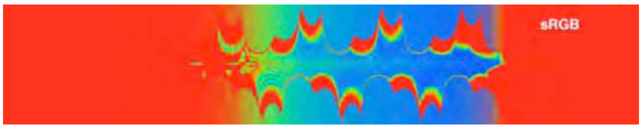

Figure 17: False color plot of the test pattern showing error behaviour of conventional LDR images (sRGB).

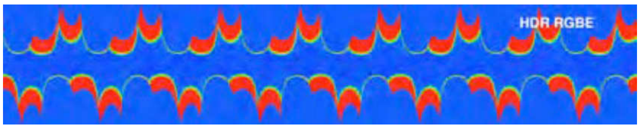

Figure 18: The RGBE encoding covers the full dynamic range, but fails to represent highly saturated colors.

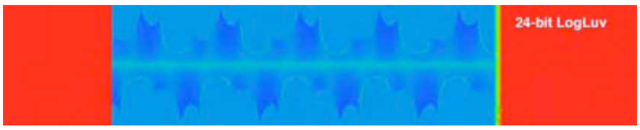

Figure 19: The 24-bit LogLuv encoding covers the visible gamut, but spans only 4.8 orders of magnitude.

As shown in Figure 20**,** the remaining candidates have performed very well, making them an excellent choice for HDR image archiving.

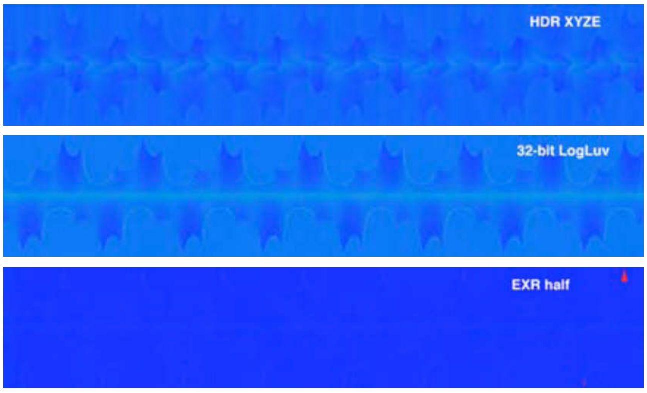

Figure 20: The XYZE, the 32-bit LogLuv and the OpenEXR encoding keep errors mostly under the visible threshold over the entire dynamic range. Only the OpenEXR encoding shows some difficulties in representing a few high saturated colors at high luminance values.

For a practical interest, compression time and file size were also measured over a test image set [8] and are shown in figure 21.

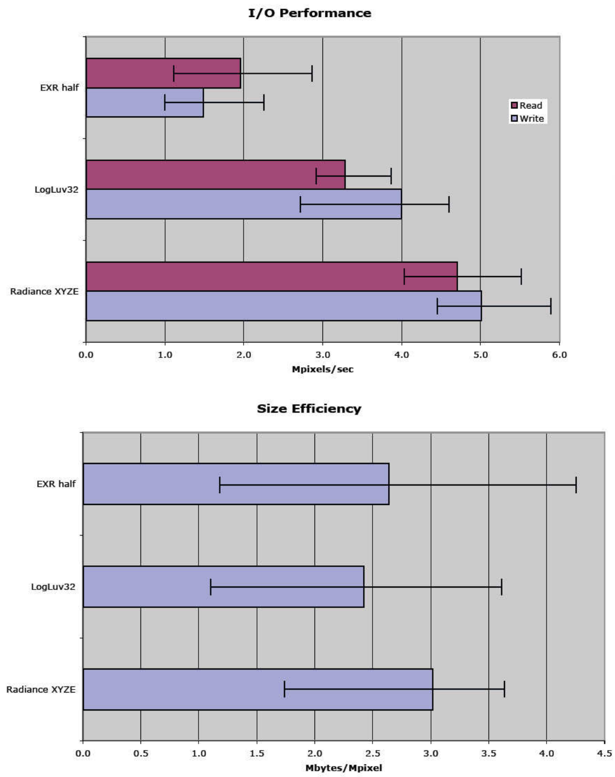

Figure 21: I/O performance and file size comparison. The lines show the range of values occurred in the test [8].

4. Applications of HDRI

4.1.1 Digital Photography

Classical analogue photography has still has the advantage over its digital version, that it can capture a greater dynamic range. But digital cameras are getting better, and are about to make use of HDR imaging technology. Various RAW formats are currently in use already, which tends in the direction of HDR and its only a matter of time, before cameras will write their images in formats that are more capable.

4.1.2 Rendering

As one of the first applications for HDRI, physically based rendering is still a field with one of the highest demands on HDRI formats. There is the need for a sufficient color representation, high accuracy and dynamic range, additional channels for alpha or depth information to name a few.

A popular use of HDR images are light probes, which are used for image based lighting. Figure 23 shows a mixed reality rendering, using a photo, rendered objects and a light probe to illuminate the rendered objects with the captured light of the real scene.

a)

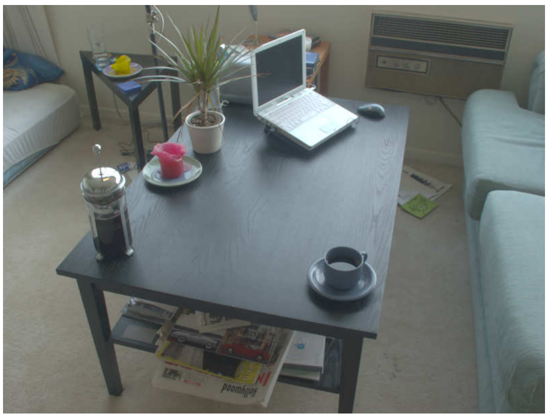

b)

c)

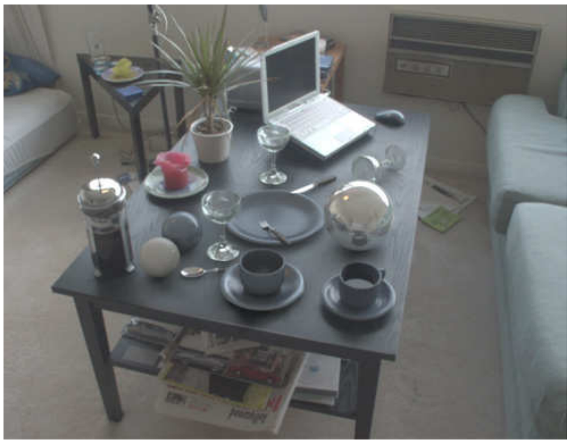

d)

e)

Figure 22: Image based lighting in a mixed reality scene: a) HDR-lightprobe, b) photo of real scene, c) virtual objects in the scene, d) final rendering e) close up of the rendering [18].

4.1.3 Digital cinema/video

Like digital photography, digital cinema is also an fast-moving application of HDR imaging. It is currently used there for editing, production and in digital movie projectors (although conventional movie projectors already have a higher range).

New standards like DVB-T and HDTV are already finding their way into our living rooms, and sooner or later the consumer market will gain a need for HDR technology, including storage-efficient HDR-formats.

4.1.4 Image Editing

The use of multiple image operations requires high accuracy to avoid accumulation errors. Also, extreme contrast or color changes are nearly impossible to accomplish with standard image formats. There is already a set of image editing software on the market, which is capable of handling 32 bit pixel data, such as Adobe Photoshop CS 2, Cinepaint, Photogenics [1].

a)

b)

c)

d)

Figure 23: Image operations applied on HDR and LDR images: a) shows the original photograph, b) the motion-blurred LDR image, c) the motion-blurred HDR image, d) the photograph with real motion blur (made with the camera) [25].

4.1.5 Remote sensing

There are applications that deal with a much larger range of light than any human can perceive. Satellites often record a wide area of electromagnetic wavelengths, for example from infrared light to x-rays. The needed accuracy may vary with the type of application, and it is often necessary to annotate using metadata.

4.1.6 Games

HDRI technology is recently used in computer games to add more realistic lighting effects. The most popular one, in which very bright parts of a scene are overlapping darker parts in a blurry way, is called *blooming*. This is meant to emulate a camera's overexposure.

HDR skyboxes are used to allow real-time exposure adjustment of the scene and therefore simulating human adaptation to light changes, as demonstrated in figure 24. Also, HDR-cube maps are used to achieve realistic impressions of reflective surfaces.

a)

b)

c)

d)

Figure 24: Tunnel lighting exposure in Need for Speed Most Wanted. Image a) shows HDR tunnel lights reflect onto the car, b)-d) demonstrate the auto-adjusting exposure [26].

5. Conclusion

Generally, HDR formats are designed to cover the abilities of human vision, making them independent from the process of display and therefore appropriate for future use.

Although many tasks for HDR images are existing, and there is already a variety of formats that are suitable for most applications, the right choice depends on the requirements. From highest precision formats for minimizing computation errors to feature rich formats used for compositing in movie-production, there are many fields of work that take advantage of the developments described in this document.

For a future mainstream use of HDRI, either in digital cameras, computer games or video, lossy formats seem to be a promising field of research.

6. References

[1] Erik Reinhard, Greg Ward, Sumanta Pattanaik, Paul Debevec: *High dynamic Range Imaging: Acquisition, display and image-based lighting*, Morgan Kaufmann, 2006

[2] Patrick Ledda, Alan Chalmers, Tom Troscianko, Helge Seetzen: Evaluation of Tone Mapping operators using a High Dynamic Range Display, 2005

[3] Paul Debevec, Dan Lemmon: *Image-based lighting.* Siggraph 01, Course #14 slides, 2001

[4] Patrick Ledda, Luis Paulo Santos, Alan Chalmers: A Local Model of Eye Adaptation for High Dynamic Range Images, 2004

[5] Sumanta Pattanaik, James Ferwerda, Mark Fairchild, Donald Greenberg: *A Multiscale Model of Adaptation and Spatial Vision for Realistic Image Display*, Program of Computer Graphics, Cornell University, 1998

[6] James A. Ferwerda, Sumanta N. Pattanaik, Peter Shirley, Donald P. Greenberg: *A Model of Visual Adaptation for Realistic Image Synthesis*, Program of Computer Graphics, Cornell University, 1996

[7] Charles Poynton: *A Technical Introduction to Digital Video*, John Wiley & Sons, 1996

[8] Greg Ward: *High Dynamic Range Image Encodings*, http://www.anyhere.com/gward/hdrenc/hdr_encodings.html, January 2006

[9] Wikipedia: *Tone mapping*, http://en.wikipedia.org/wiki/tone_mapping, December 2006

[10] Frédo Durand, Julie Dorsey: *Fast bilateral filtering for the display of high-dynamic-range images*, Laboratory for Computer Science, MIT, 2002

[11] Micheal Stokes, Matthew Anderson, Srinivasan Chandrasekar, Ricardo Motta: *A Standard Default Color Space for the Internet*, www.w3.org/Graphics/Color/sRGB, November 1996

[12] Greg Ward: *JPEG-HDR Whitepaper*, BrightSide Technologies, March 2006

[13] Gregory Ward Larson: *Overcoming Gamut and Dynamic Range Limitations in Digital Images*, Silicon Graphics, Inc., 1998

[14] *IEEE Standard for Binary Floating-Point Arithmetic*, The Institute of Electrical and Electronics Engineers, Inc. 1985

[15] Kimmo Roimela, Tomi Aarnio, Joonas Itäranta: *High Dynamic Range Texture Compression*, Nokia Research Center, 2006

[16] *Technical Introduction to OpenEXR*, Industrial Light & Magic, September 2006

[17] Bruce Fraser: *Understanding Digital Raw Capture*, Adobe Whitepaper, 2004

[18] Bloch, Christian: *Praktischer Einsatz von High Dynamic Range Imaging in der Postproduktion*, FH Leipzig, 2003

[19] Nathan Willis: *High Dynamic Range images under Linux*, http://www.linux.com/article.pl?sid=05/12/06/2115258, December 2006

[20] *Digital Negative (DNG) Specification*, Adobe Systems Incorporated, 2005

[21] *Defnition of the Flexible Image Transport System*, NASA/Science Office of Standards and Technology, 1999

[22] *DDS File Reference*, http://msdn.microsoft.com/archive/default.asp?url=/archive/en-us/directx9_c/directx/graphics/reference/ddsfilereference/ddsfileformat.asp, Microsoft, 2007

[23] Kevin Spaulding, Geoffrey Woolfe, Rajan Joshi: *Using a Residual Image to Extend the Color Gamut and Dynamic Range of an sRGB Image*, ERI-JPEG Whitepaper, Eastman Kodak Company, 2003

[24] *DPX File Format Summary*, http://www.fileformat.info/format/dpx/, December 2006

[25] Paul Debevec, http://www.debevec.org/, December 2006

[26] Erik Reinhard, Paul Debevec, Greg Ward, Karol Myszkowski, Helge Seetzen, Drew Hess, Gary McTaggart, Habib Zargarpour: *High Dynamic Range Imaging: Theory and Practice*. Siggraph 06, Course #5 slides, 2006

[27] Debevec, Reinhard, Ward, Pattanaik: *High Dynamic Range Imaging*, Siggraph 04 Course #13 slides, 2004

[28] *IEEE Standard for Binary Floating-Point Arithmetic*, The Institute of Electrical and Electronics Engineers, Inc., 1985

[29] Greg Ward, Maryann Simmons: *Subband Encoding of High Dynamic Range Imagery*, In Proceedings of APGV 04, 2004

[30] Ruifeng Xu, Sumanta Pattanaik, Charles E. Hughes: *High-dynamic range still-image encoding in JPEG 2000*. IEEE Computer Graphics and Applications 25, 6, 57–64, 2005