



HAL
open science

An OpenEXR Layout for Spectral Images

Alban Fichet, Romain Pacanowski, Alexander Wilkie

► **To cite this version:**

Alban Fichet, Romain Pacanowski, Alexander Wilkie. An OpenEXR Layout for Spectral Images. Journal of Computer Graphics Techniques, 2021, 10 (3). hal-03252797v2

HAL Id: hal-03252797

<https://inria.hal.science/hal-03252797v2>

Submitted on 30 Sep 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NoDerivatives 4.0 International License

An OpenEXR Layout for Spectral Images

Alban Fichet
IOGS, CNRS
Charles University

Romain Pacanowski
CNRS, INRIA

Alexander Wilkie
Charles University

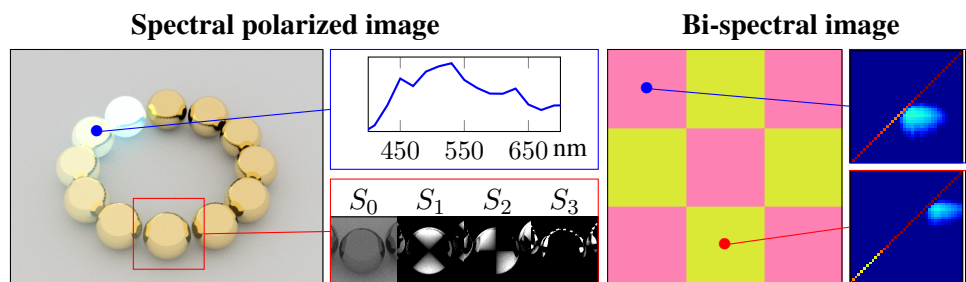


Figure 1. Examples of using the proposed OpenEXR layout to store spectral and bi-spectral information along with the polarization state. Left: For any given pixel of a radiance image, we can extract the full spectral data (blue box). We also store the corresponding Stokes components per wavelength, shown here averaged for all stored wavelengths in a specific area of the image (red box). Right: A bi-spectral image can store a full re-radiation matrix for each pixel. This is obviously only meaningful as input data for a reflective texture. The sample points are two locations with different bi-spectral reflectance, fluorescent pink (blue box) and fluorescent yellow (red box). The corresponding RGB image uses a D65 illuminant. These two OpenEXR spectral images are provided in the supplemental materials.

Abstract

We propose a standardized layout to organize spectral data stored in OpenEXR images. We motivate why we chose the OpenEXR format as the basis for our work, and we explain our choices with regard to data selection and organization: our goal is to define a standard for the exchange of measured or simulated spectral and bi-spectral data. We also provide sample code to store spectral images in OpenEXR format.

1. Introduction

In recent years, spectral rendering has become increasingly popular, e.g., for high end visual effects work [Fascione et al. 2018]. The main reason for its increasing adoption is that it allows for a far more physically plausible handling of various important vi-

sual effects such as skin appearance, at least when compared to tristimulus systems. It also allows proper integration of still rarely found features such as dispersion, polarization, and fluorescence. Spectral renderer output is not tied to any particular color space, whereas tristimulus images are intrinsically only valid for a given combination of white point and RGB primaries. However, the output of spectral rendering systems such as Manuka [Fascione et al. 2018] is usually not stored in spectral form: the main reasons why this is rarely done are spectral aliasing issues, storage considerations, and the fact that, for most workflows, wide-gamut RGB data in floating point format is sufficient from the point of rendering onward.

However, there are applications in both research and production environments where storing the actual spectral output of a renderer is useful, so we still need a standardized way of doing this. And there is one additional application area where the ability to use spectral images is arguably far more important anyway: as *input* for the shading subsystem of a spectral renderer. If accurate matches between measured surface appearance and renderings are desired, the ability to handle (bi-)spectral reflectance textures is essential.

Many RGB image formats exist, but—somewhat surprisingly—the same is not true for spectral images. Each pipeline uses either its own proprietary flavor of image format (e.g., ENVI hyperspectral [Harris Geospatial Solutions 2020]) or a collection of monochromatic images stored in a set of RGB image formats (e.g., META-COW [Fairchild and Johnson 2004]). This makes the exchange, support in software, and long-term storage of spectral images difficult.

We propose a spectral image format based on the OpenEXR framework: the result is suitable not only for rendering but also for various applications such as acquisition, and as a storage format for (bi-)spectral reflectance. In addition, we retain the user friendliness of an RGB image format allowing a quick preview of the spectral image in commonly used graphics software.

2. Related Work

To store spectral images, a simple technique is to provide a collection of standard RGB image files, where each image corresponds to a specific spectral band [Yasuma et al. 2010]. The key advantage of this method is that it allows uncomplicated viewing of all spectral channels with standard software. On the other hand, handling multiple separate files is not convenient, and reconstruction of color images from the spectral data is tedious.

Specialized formats exist for storing spectral data in specialized fields. For example, the ENVI format [Harris Geospatial Solutions 2020] targets satellite imagery, and separate data blocks are stored for each spectral channel. A header file, in plain text format, describes the information stored in the images, and a binary file is used

as a dump of the image data. This reduces the amount of separate file types, but still requires two separate files in order to read the image. Moreover, there is no standardized way to store polarization and bi-spectral information, as especially the latter never occurs in satellite imaging contexts.

In the rendering context, the ARTRAW format was introduced along with the Advanced Rendering Toolkit (ART) renderer [Advanced Rendering Toolkit 2018]. It supports polarized emissive images, and the spectral sampling options are those of the renderer itself (8, 11, 18 and 46 spectral samples, respectively). The header and data are stored in a single file, in a mixture of plain text (header) and binary (data) formats. This format is mostly used only by this specific renderer and cannot be displayed without specialized software: it also intentionally eschews all attempts at reducing file size, as it was only ever intended as a lossless reference data storage for raw rendering output. As spectral images tend to be quite large, any reduction in file size would be very welcome; so, adopting ARTRAW (which is not very flexible in terms of its internal structure anyway) as a standard would be decidedly suboptimal.

Finally, there is OpenEXR [Kainz et al. 2004], the de facto standard for storing HDR color space images. The official OpenEXR specifications do not cover the use and description of spectral data, but allow very flexible storage of multiple layers of binary data. Some software packages already use layers of the OpenEXR format to store spectral data [Eclat Digital 2019], but, to the best of our knowledge, the specifications used there are limited. These formats generally lack support of polarization, bi-spectral, and capture metadata information with a unified interface.

3. Scope of Our Spectral Image Representation

In this section, we describe the range of information we intend to store with our spectral representation.

3.1. Reflective and Emissive Spectra

First, we distinguish between *emissive* and *reflective* spectra. Both are used in distinct places within a rendering system, and do not share the same physical units.

An emissive spectrum $L_e(\lambda)$ describes radiant energy within a scene, such as the radiance emitted by light sources. The image computed by a path tracer is a collection of emissive spectral values: it represents the spectral radiance reaching the sensor. The pixels in an environment map are a specialized case of light source and also contain emissive spectral data. $L_e(\lambda)$ is expressed in $\text{W} \cdot \text{m}^{-2} \cdot \text{sr}^{-1} \cdot \text{nm}^{-1}$ in radiometric units.¹

On the other hand, a reflective spectrum describes the attenuation of energy af-

¹The corresponding photometric unit is nits, which is used when working with such quantities in a tristimulus pipeline.

ter reflection, scattering, or transmission/volume absorption events. For instance, an albedo texture image describes how a surface attenuates incoming radiant energy upon reflection. The albedo is unitless and is normally used in the context of a BRDF model (in sr^{-1}) that further describes the directional dependency of surface reflectance:

$$\rho_\lambda(\omega_o) = \int_{\theta_i} \int_{\phi_i} \text{brdf}_\lambda(\omega_o, \theta_i, \phi_i) \cos \theta_i \sin \theta_i d\theta_i d\phi_i \quad [\text{sr}^{-1}]. \quad (1)$$

3.2. Polarization

Although it is not directly visible to human eyes,² the polarization state of electromagnetic radiation can play an important role in light transport for some scenes. Specular inter-reflections are mis-predicted to varying degrees by non-polarizing rendering computations, as well as the reflections of blue sky radiance off specular surfaces. For various purposes related to rendering, such as being able to store a polarized environment map of sky dome radiation and to inspect the polarization state of reflective highlights in a scene, we want a spectral image format to optionally also store polarization information.

The mathematical formalism of Stokes vectors and Mueller matrices allows the representation of fully, partially, or non-polarized light in a compact and intuitive manner that blends in well with non-polarized spectral rendering. Stokes vectors describe light polarization,

$$S = \begin{bmatrix} S_0 \\ S_1 \\ S_2 \\ S_3 \end{bmatrix}, \quad (2)$$

and Mueller matrices describe the effect of a scattering event on light polarization: it encodes the transform of a Stokes vector [Goldstein 2010; Wilkie and Weidlich 2011].

In order to achieve the desired functionality for emissive images, we store the following information, per pixel and per wavelength:

- all Stokes components S_0, S_1, S_2, S_3 if the image contains polarization information, or
- only the first Stokes component S_0 if the image does not contain polarization information.

To use a consistent notation, we always refer to emissive spectral radiance in images as Stokes component S_0 , even in non-polarized images.

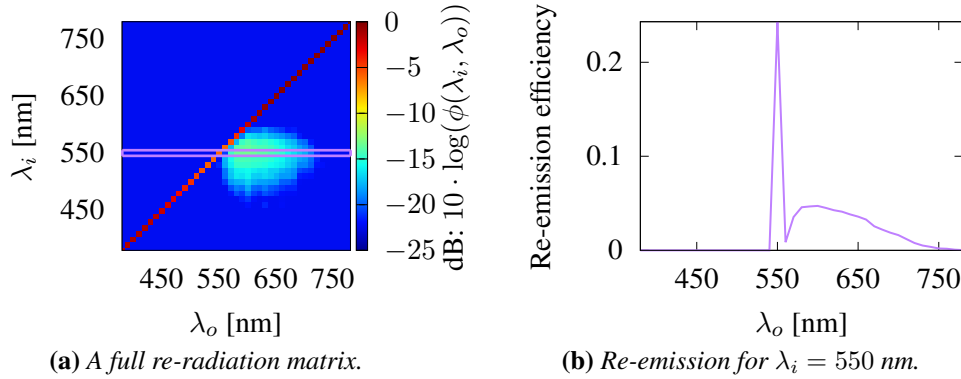


Figure 2. (a) A full re-radiation matrix like this example completely describes bi-spectral reflectance for the given input (λ_i) and output (λ_o) wavelength ranges. For a given illumination wavelength λ_i , the light re-emitted at various wavelengths is shown as a row in the 2D diagram. (b) A plot of the row highlighted in purple in the full matrix at $\lambda_i = 550$ nm, including the diagonal part corresponding to the non-bi-spectral reflectance.

3.3. Bi-spectral Reflectance

Even though this behavior is not yet commonly simulated in current rendering software, a significant number of manmade and natural materials exhibit wavelength shifting behavior when reflecting or scattering light. Apart from obviously fluorescent artificial materials such as Day-Glo warning paint, examples that show at least some degree of such behavior include wood (lignin is fluorescent), green plants (chlorophyll), and paper and textiles (mostly due to washing powders usually containing optical brighteners that make white cloth look “whiter” by adding a blue fluorescent tinge to it).

Similar to storing “normal” spectral albedo textures that are then used in the context of a BRDF model, complex appearance modeling requires the ability to store extended albedo textures that also describe wavelength shifting. As it can be seen in Figure 2, one needs to store a two-dimensional spectrum—the re-radiation matrix—for each pixel in this case: for a given incoming wavelength λ_i , re-emission can occur at various intensities for a wide range of possible output wavelengths λ_o .

3.4. Capture Information

For measured spectral images, we also need to store the capture device characteristics as metadata. The Cornell Box website [Cornell University Program of Computer Graphics 1998] serves as an example of the sort of data we need to retain: as reference data, they provide a set of monochromatic images of the real Cornell Box from

²At least mostly so: strongly polarized light is weakly visible to some observers, but the effect is not thought to significantly impact how we perceive scenes in general.

400 nm to 700 nm with a sample taken every 50 nm.

In the Cornell dataset, each image was taken with a monochrome camera with a specific narrow-band filter placed in front of the lens: detailed information about the measurement setup is provided along with the raw data on the website. If one wants to be able to interpret the contents of such a measured image in a stand-alone form, directly from the image, this information should be included in the image header. There are two optical elements that alter the energy received by the sensor: First, the camera lens has a characteristic transmission spectrum. Second, each measurement filter has a transmission spectrum. In addition to this, the camera sensor response is usually not linear and also needs to be stored. If all this information is retained in the image header, a captured spectral image can be used for long-term archival storage without losing the ability to properly interpret its contents.

4. Using OpenEXR for Spectral Data

OpenEXR supports the storage of multiple layers within a single file, but there are some constraints on using this functionality. In this section, we explain our choices made in order to meet our requirements for storing spectral data without violating the OpenEXR format specification.

4.1. Layers and Channels Organization

The official OpenEXR specification [Kainz 2013] states that an OpenEXR image can consist of different layers and that each layer can contain an arbitrary number of channels. In other words, a layer is a group of channels, and from an implementation point of view, layers and channels are indistinguishable. They are merely separated by a naming convention that is enforced by convenient functions in the official OpenEXR library. A layer name is followed by a dot character. Layers can be nested and the tailing string names a channel.

The following outlines the layer naming conventions that we propose.

4.1.1. Layers

Emissive images. Layers that store individual Stokes components of a polarized image are named S_0 , S_1 , S_2 , and S_3 . If there is no polarization information, only the S_0 layer is present.

Reflective images. Layers that store reflective or attenuation values are named T .

4.1.2. Channels

Each channel is named according to the wavelength or frequency bracket it represents. Because the OpenEXR format reserves dots for separating different layers, we need to use a decimal comma notation for wavelength or frequency values. Wave-

	Information	Description	Example	Mandatory
1	Layer name	Stokes component or absorption	S0.	Yes
2	Value	Comma separated	350,5	Yes
3	Exponent	Power of ten	E-3	No
4	Multiplier	Unit multiplier	k	No
5	Units	Hz or m	Hz	Yes

Table 1. Description of channel naming convention.

length/frequency values are always positive values, and an optional integer power of ten exponent also can be provided with E or e followed by an optional sign and an integer value.

This value is followed by an optional unit multiplier and by the used unit. These are either Hertz (Hz) or meters (m). Table 1 enumerates the different fields that describe a channel name. The channel naming convention is also described in Appendix A.2.

For instance, the channel “S0.560,5nm” corresponds to an S₀ component at 560.5 nm.

A regex matching this naming convention can be written as follows:

```
layer = "((S[0-3])|T) "
value = "(\\d*,?\\d*([eE][+-]?\\d+)?) "
units = "(Y|Z|E|P|T|G|M|k|h|da|d|c|m|u|n|p|f|a|z|y)?(m|Hz) "
regex = "^" + layer + '\\.' + value + units + "$"
```

4.1.3. Bi-spectral Data

To store a re-radiation matrix for each pixel of a reflective image, we can also use the OpenEXR layering system, although the resulting image is of considerable size. The main diagonal, which corresponds to the case where there is no re-radiation, follows the same convention as described in Section 4.1.2. Nested layers store the re-radiation part named with the re-radiation wavelength or frequency using the same comma separated format followed by the units. For instance, re-radiation from 560 nm to 600 nm would be stored in a channel named “T.560nm.600nm.”

4.1.4. Nested Hierarchy

We also define a layer convention for child layers, so that it is possible to use a spectral layout within an existing layer hierarchy. For instance, a stereo image with 550 nm layers will have its “default” (left) spectral layer stored in “S0.550nm” and its right view in “right.S0.550nm.”

4.2. Metadata

We allow the storage of additional metadata to retain capture information using OpenEXR custom attributes in the image header. OpenEXR documentation recommends using types provided from `ImfStandardAttributes.h`. A summary of all metadata we support is provided in Appendix A.1.

4.2.1. Layout Version

To allow future evolution of the proposed spectral layout, we add an attribute that specifies the format version. If the layout is exactly as described in this article, the value shall be set to “1.0” (see Table 2).

Description	Attribute Name	Value
Version	<code>spectralLayoutVersion</code>	“1.0”

Table 2. Mandatory version information of the layout in the OpenEXR header.

4.2.2. Radiometric Units

When the image contains emissive spectral layers, the header has to specify the used radiometric units (see Table 3). For instance, the result of a rendering or a hyperspectral image shall be in $W \cdot m^{-2} \cdot sr^{-1}$.

Description	Attribute Name	Value
Emissive layers radiometric units	<code>emissiveUnits</code>	“W,” “W.m ⁻² ,” “W.sr ⁻¹ ,” or “W.m ⁻² .sr ⁻¹ ”

Table 3. Emissive unit information in the OpenEXR header: this is only mandatory for emissive images.

4.2.3. Polarization

When using emissive polarized images, we need to specify the orientation convention used (see Table 4). Figure 3 illustrates how the same signal can be described as left circular polarized or right circular polarized depending on the convention used.

Description	Attribute Name	Value
Polarization frame	<code>polarisationHandedness</code>	“left” or “right”

Table 4. Mandatory metadata in the OpenEXR header for emissive polarized spectral images.




Figure 3. To avoid confusion when interpreting Stokes parameters, a polarization orientation convention must be specified. Notice that (a) and (b) have opposite circular polarization orientation for the same signal depending on the convention used. All extant polarization renderers we know of (Mitsuba 2 and ART) use the right-handed convention.

4.2.4. Capture Information

We use optional custom attributes to store properties of the acquisition device (see Table 5):

- the lens system transmission spectrum,
- the camera sensitivity curve,
- a response curve for each S0 channel.


In the case that the lens transmission spectrum or the camera sensitivity spectrum metadata field is not present, we assume that this optical element is perfect. When no filter is specified in the header for a spectral bin, the filter is assumed to be a gate function. The filter boundaries $[B_n(\lambda^-), B_n(\lambda^+)]$ for a given bin wavelength B_n are

$$B_n(\lambda^-) = B_n - \frac{B_n - B_{n-1}}{2}, \quad (3)$$


$$B_n(\lambda^+) = B_n + \frac{B_{n+1} - B_n}{2}. \quad (4)$$

Description	Attribute Name	Type
Lens system transmission spectrum	lensTransmission	Spectrum
Camera sensitivity spectrum	cameraResponse	Spectrum
Filter/channel bandwidth	Same as S0 channel names	Spectrum

Table 5. Optional metadata in the OpenEXR header for capture setup information.



(a) Regularly sampled spectrum.



(b) Varying sampling rate over the spectrum.

Figure 4. Default spectral bin filtering behavior when no filter is specified in the header.

And the special cases for the first and the last bins are

$$B_1(\lambda^-) = B_1 - \frac{B_2 - B_1}{2}, \quad (5)$$

$$B_N(\lambda^+) = B_N + \frac{B_N - B_{N-1}}{2}. \quad (6)$$

Figure 4 illustrates this default filtering behavior with regular and varying sampling rate.

4.2.5. Spectrum Type Metadata


To store spectrum types, we use an `Imf::StringAttribute` packing spectrum in a string. Each pair of wavelength or frequency and value are separated by a colon and followed by a semicolon. We enforce specification of the units of the wavelength or frequency and allow a unit multiplier. Contrary to channel naming, we do not have a constraint on the character set, so we use dot decimals.

A regex matching this naming convention can be written as follows:

```
value = "(\\d*\\.?.?\\d*([eE][+-]?\\d+)?) "  
units = "(Y|Z|E|P|T|G|M|k|h|da|d|c|m|u|n|p|f|a|z|y)?(m|Hz) "  
regex = value + units + ":" + value + ";"
```

For instance, a triangular filter response from 620.5 nm to 679.5 nm with its maximum at 650 nm is described as follows:

```
620.5nm:0;650nm:1;679.5nm:0;
```



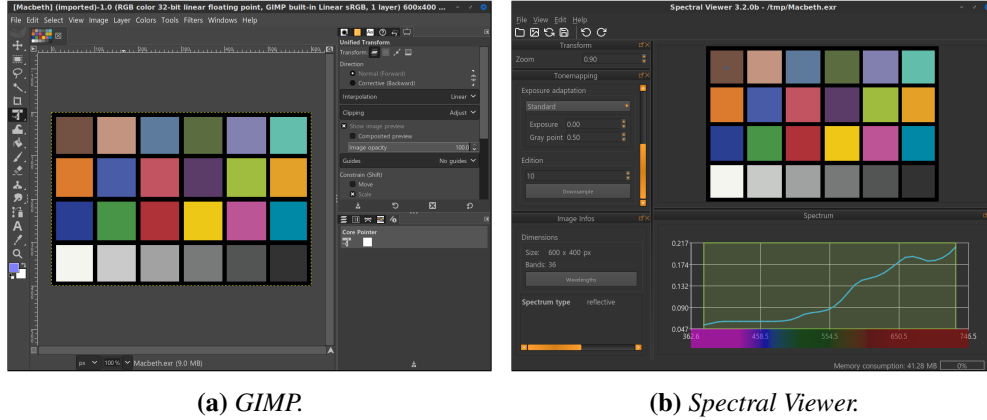


Figure 5. The same spectral OpenEXR image opened in two different programs: first, a standard graphic software, GIMP [GIMP 1996], which is only capable of handling RGB images, and second, our specialized software for spectral images, Spectral Viewer [Dufay et al. 2018]. (a) Gimp cannot access the spectral information but still can render the RGB version thanks to the additional R, G, and B channels provided in the spectral OpenEXR image. (b) Our specialized software can read the spectral information stored in each pixel.

4.3. Data Compression

The OpenEXR format offers a wide variety of lossless and lossy compression algorithms, which is interesting for spectral and bi-spectral data that have a significant memory footprint. An overview of the various compression formats available are discussed in the technical introduction to OpenEXR [Kainz et al. 2013]: choosing any particular one of them depends on the application one needs spectral image storage for (that is, whether lossy compression is acceptable, or not). When in doubt, lossless compression should be chosen.

5. Enforcing Backward Compatibility

OpenEXR is widely supported in graphic packages (GIMP, Photoshop, etc.) in its RGB variant. Although these software packages only support normal color spaces (RGB, CMYK, etc.), we can provide additional layers in spectral OpenEXR files that allow opening and previewing an RGB version of the spectral image (see Figure 5).

This is simply done by adding R, G, and B channels to the layer set. This behavior, which is not necessary to have a valid spectral OpenEXR file, is strongly encouraged in the OpenEXR official documentation page [Hillman and Welford 2007].

Multiple spectral-to-RGB conversions exist. In this section we explain how we keep track of the conversion pipeline.

5.1. Proposed Conversion Pipeline

5.1.1. Emissive Images

For emissive images, we use the S0 layer and convert the spectrum to RGB using the following pipeline:

1. Convert emissive spectrum $L_e(\lambda)$ to CIE-XYZ 2° 1931.
2. Convert XYZ to linear sRGB D65.

Steps 1 and 2 correspond to the following equation:

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \mathbf{M}_{\text{XYZ} \rightarrow \text{RGB}} \cdot \begin{bmatrix} \int_{\Lambda} L_e(\lambda) \cdot \bar{x}(\lambda) d\lambda \\ \int_{\Lambda} L_e(\lambda) \cdot \bar{y}(\lambda) d\lambda \\ \int_{\Lambda} L_e(\lambda) \cdot \bar{z}(\lambda) d\lambda \end{bmatrix}, \quad (7)$$

where

$$\bullet \mathbf{M}_{\text{XYZ} \rightarrow \text{RGB}} = \begin{bmatrix} 3.2404542 & -1.5371385 & -0.4985314 \\ -0.9692660 & 1.8760108 & 0.0415560 \\ 0.0556434 & -0.2040259 & 1.0572252 \end{bmatrix},$$

- $L_e(\lambda)$ is the emissive spectrum,
- $\bar{x}(\lambda), \bar{y}(\lambda), \bar{z}(\lambda)$ are the sensitivity curves for the CIE-XYZ 2° 1931 color space.

5.1.2. Reflective Images

For a reflective image, we use the T layer and convert the spectrum to RGB using the following pipeline:

1. Multiply the spectrum $S(\lambda)$ with CIE D65 illuminant $L_{\text{D65}}(\lambda)$.
2. Project the spectrum to CIE-XYZ 2° 1931 color space.
3. Normalize the components with CIE D65 Y component (CIE 1931 Y).
4. Convert XYZ to sRGB D65 linear.

Spectral case. The conversion is defined as follows:

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \frac{1}{Y_{\text{D65}}} \cdot \mathbf{M}_{\text{XYZ} \rightarrow \text{RGB}} \cdot \begin{bmatrix} \int_{\Lambda} L_{\text{D65}}(\lambda) \cdot S(\lambda) \cdot \bar{x}(\lambda) d\lambda \\ \int_{\Lambda} L_{\text{D65}}(\lambda) \cdot S(\lambda) \cdot \bar{y}(\lambda) d\lambda \\ \int_{\Lambda} L_{\text{D65}}(\lambda) \cdot S(\lambda) \cdot \bar{z}(\lambda) d\lambda \end{bmatrix}, \quad (8)$$

where

- $Y_{\text{D65}} = \int_{\Lambda} L_{\text{D65}}(\lambda) \cdot \bar{y}(\lambda) d\lambda$,
- $S(\lambda)$ is the reflective spectrum.

Bi-spectral case. The integration becomes two dimensional:

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \frac{1}{Y_{D65}} \cdot \mathbf{M}_{XYZ \rightarrow RGB} \cdot \begin{bmatrix} \int_{\Lambda_o} \int_{\Lambda_i} L_{D65}(\lambda_i) \cdot S(\lambda_i, \lambda_o) \cdot \bar{x}(\lambda_o) d\lambda_i d\lambda_o \\ \int_{\Lambda_o} \int_{\Lambda_i} L_{D65}(\lambda_i) \cdot S(\lambda_i, \lambda_o) \cdot \bar{y}(\lambda_o) d\lambda_i d\lambda_o \\ \int_{\Lambda_o} \int_{\Lambda_i} L_{D65}(\lambda_i) \cdot S(\lambda_i, \lambda_o) \cdot \bar{z}(\lambda_o) d\lambda_i d\lambda_o \end{bmatrix}. \quad (9)$$

5.2. Using a Different Conversion Pipeline

One can use one’s own conversion pipeline for a specific application. In this case, additional information can be provided as OpenEXR header attributes: the XYZ sensitivity curves, the illuminant spectrum for reflective images, and an additional exposure compensation value (see Table 6). This additional information is especially important for bi-spectral images where the illuminant greatly influences the RGB conversion even for two illuminants with the same white point.

The exposition compensation is defined as follows:

$$\mathbf{RGB}' = \mathbf{RGB} \cdot 2^{EV}, \quad (10)$$

where EV is the exposure compensation value (unitless).

Description	Attribute Name	Type	Default
Illuminant	illuminant	Spectrum	CIE D65
X color matching function	X	Spectrum	CIE X 1931 2°
Y color matching function	Y	Spectrum	CIE Y 1931 2°
Z color matching function	Z	Spectrum	CIE Z 1931 2°
Exposure compensation	EV	Float	0

Table 6. Optional metadata for the spectral-to-tristimulus conversion pipeline. The spectrum format follows the same convention as the one described in Section 4.2.5.

Specification of the full conversion pipeline, from spectral to linear RGB values, is necessary to validate one’s specific conversion pipeline. With this information on hand, a producer and a consumer of a spectral image with RGB channels can validate the consistency of their spectral conversion pipelines.

6. Conclusion

We introduce a versatile, unambiguous, and flexible layout for storing spectral data within the existing framework of the OpenEXR image format. This layout supports emissive and reflective spectral images, and also handles polarization and bi-spectral reflective spectra. In addition to pixel data, we describe how to retain metadata from capture devices: the camera spectral response, lens transmission, and filter properties. It also allows the storage of all spectral and capture information in a single file, and

common graphic software can display the spectral image in its tristimulus form if the RGB channels are also provided. The format also intrinsically supports various compression schemes on the individual channels, by virtue of being part of the OpenEXR framework. Finally, the format is modular and allows additional metadata and layers for specific applications.

Acknowledgements


We want to thank the reviewers for their meaningful advice, among these, contacting the OpenEXR developers. We also want to thank the OpenEXR developers and mailing list members for their valuable feedback. This work was funded in part by French National Research Agency grant ANR-17-CE23-0017 “VIDA.” We also acknowledge funding by Czech Science Foundation grant 19-07626S, EU H2020-MSCA-ITN-2020 grant number 956585 (PRIME), and Charles University grant number SVV-260588.

A. Appendix: Summary of the Spectral OpenEXR Layout

A.1. Header Metadata

Attribute Name	Description	Type	Mandatory	Default	Table
spectralLayoutVersion	Version of the format	String	Yes ³	-	2
polarisationHandedness	Handedness of polarization observation frame	String (“left” or “right”)	Yes ⁴	-	4
emissiveUnits	Radiometric units for emissive layers	String	Yes ⁵	-	3
lensTransmission	Lens transmission spectrum	Spectrum	No	Flat (perfect optical system)	5
cameraResponse	Camera sensitivity spectrum	Spectrum	No	Flat (perfect optical system)	5
Same as S0 channel names	Filter / Channel bandwidth	Spectrum	No	Flat (perfect optical system)	5
illuminant	Illuminant	Spectrum in a string	No	CIE D65	6
X	X color matching function	Spectrum in a string	No	CIE X 1931 2°	6
Y	Y color matching function	Spectrum in a string	No	CIE Y 1931 2°	6
Z	Z color matching function	Spectrum in a string	No	CIE Z 1931 2°	6
EV	Exposure compensation (see Equation (10))	Float	No	0	6

A.2. Layer Naming



³Value must be set to 1.0 for the layout described in this article.

⁴Only for polarized images.

⁵Only for emissive images. Value must be “W,” “W.m⁻²,” “W.sr⁻¹,” or “W.m⁻².sr⁻¹.”

⁶Comma separated integer and decimal parts.

Index of Supplemental Materials

Sample Code

We provide C++ sample code for reading and writing spectral and bi-spectral OpenEXR files under BSD License 2.0 terms.

- **Spectral:** The spectral classes (`[EXR]?SpectralImage.*`) are a simpler version for one interested only in the spectral part of the OpenEXR image. This reads as well bi-spectral files but ignores the re-radiation information.
- **Bi-spectral:** The bi-spectral classes (`*BiSpectralImage.*`) support both spectral and bi-spectral OpenEXR images.

Sample Images

We provide several sample images.

Emissive Images

- **D65.exr:** The spectral power distribution of the CIE D65 illuminant. A 1×1 pixel image giving an example of the application for storing spectra.
- **CornellBox.exr:** A single-file spectral EXR version compiled from the official Cornell box data: <http://www.graphics.cornell.edu/online/box/data.html>. It contains the aggregation of the monochromatic OpenEXR images and the filter responses for each band.
- **PolarisedRendering.exr:** A spectral EXR file with polarization information from the ART renderer [Advanced Rendering Toolkit 2018].
- **MatProbe.exr:** A spectral, stereo, non-polarized EXR file from the Malia Rendering Framework (MRF) renderer [Dufay et al. 2019].

Reflective Images

- **Macbeth.exr:** A reflective spectral image. It represents a Macbeth color chart. The spectral data comes from measurements of several Macbeth color chart and was taken from <http://www.babelcolor.com/colorchecker-2.htm>.
- **BiSpectral.exr:** A bi-spectral texture made of two different fluorescent reflectances.

References

- ADVANCED RENDERING TOOLKIT, 2018. The ART homepage. URL: <https://cg.mff.cuni.cz/ART/>. 3, 16
- CORNELL UNIVERSITY PROGRAM OF COMPUTER GRAPHICS, 1998. The Cornell Box. URL: <http://www.graphics.cornell.edu/online/box/>. 5
- DUFAY, A., PACANOWSKI, R., BARLA, P., AND FICHET, A., 2018. Spectral Viewer. URL: <https://adufay.gitlabpages.inria.fr/SpectralViewer/>. 11
- DUFAY, A., MURRAY, D., PACANOWSKI, R., ET AL., 2019. The Malia Rendering Framework. URL: <https://pacanows.gitlabpages.inria.fr/MRF/>. 16
- ECLAT DIGITAL, 2019. Ocean 2019 documentation: Spectral EXR format. URL: <http://www.eclat-digital.com/ocean2019-docs/reference/fileformat/spectralexr.html>. 3
- FAIRCHILD, M. D., AND JOHNSON, G. M. 2004. METACOW: A public-domain, high-resolution, fully-digital, noise-free, metamerically extended-dynamic-range, spectral test target for imaging system analysis and simulation. *Color and Imaging Conference 2004*, 1, 239–245. URL: <https://www.ingentaconnect.com/content/ist/cic/2004/00002004/00000001/art00043>. 2
- FASCIONE, L., HANIKA, J., LEONE, M., DROSKE, M., SCHWARZHAUPT, J., DAVIDOVIC, T., WEIDLICH, A., AND MENG, J. 2018. Manuka: A batch-shading architecture for spectral path tracing in movie production. *ACM Trans. Graph* 37, 3, 31:1–31:18. URL: <https://dl.acm.org/doi/abs/10.1145/3182161>. 1, 2
- GIMP, 1996. GNU Image Manipulation Program. URL: <https://www.gimp.org>. 11
- GOLDSTEIN, D. H. 2010. *Polarized Light*, third ed. CRC Press. URL: <https://doi.org/10.1201/b10436>. 4
- HARRIS GEOSPATIAL SOLUTIONS, 2020. ENVI header files. URL: <http://www.harrisgeospatial.com/docs/ENVIHeaderFiles.html>. 2
- HILLMAN, P., AND WELFORD, M., 2007. Storing multi-view images in OpenEXR files. URL: <https://www.openexr.com/documentation/MultiViewOpenEXR.pdf>. 11
- KAINZ, F., BOGART, R., AND HESS, D. 2004. The OpenEXR image file format. In *GPU Gems: Programming Techniques, Tips, and Tricks for Real-Time Graphics*, R. Fernando, Ed. Addison-Wesley, ch. 26. URL: <https://developer.nvidia.com/gpugems/gpugems/part-iv-image-processing/chapter-26-openexr-image-file-format>. 3
- KAINZ, F., BOGART, R., STANCZYK, P., AND HILLMAN, P., 2013. Technical Introduction to OpenEXR. URL: <https://www.openexr.com/documentation/TechnicalIntroduction.pdf>. 11
- KAINZ, F., 2013. Reading and writing OpenEXR image files with the IlmImf Library. URL: <https://www.openexr.com/documentation/ReadingAndWritingImageFiles.pdf>. 6

WILKIE, A., AND WEIDLICH, A. 2011. How to write a polarisation ray tracer. In *SIGGRAPH Asia 2011 Courses*, Association for Computing Machinery, 8:1–8:36. URL: <https://doi.org/10.1145/2077434.2077442>. 4

YASUMA, F., MITSUNAGA, T., ISO, D., AND NAYAR, S. K. 2010. Generalized assorted pixel camera: Postcapture control of resolution, dynamic range, and spectrum. *IEEE Transactions on Image Processing* 19, 9, 2241–2253. URL: <https://ieeexplore.ieee.org/abstract/document/5439932>. 2

Author Contact Information

Alban Fichet
Institut d’Optique Graduate School, CNRS
Charles University
alban.fichet@gmx.fr
<https://afichet.github.io>

Romain Pacanowski
CNRS, INRIA
romain.pacanowski@institutoptique.fr
<http://manao.inria.fr/perso/~pac>

Alexander Wilkie
Charles University
wilkie@cgg.mff.cuni.cz
<https://cgg.mff.cuni.cz/~wilkie>

A. Fichet, R. Pacanowski, and A. Wilkie, An OpenEXR Layout for Spectral Images, *Journal of Computer Graphics Techniques (JCGT)*, vol. 10, no. 3, 1–18, 2021
<http://jcgt.org/published/0010/03/01/>

Received: 2021-01-07
Recommended: 2021-04-20 Corresponding Editor: Marc Stamminger
Published: 2021-09-29 Editor-in-Chief: Marc Olano

© 2021 A. Fichet, R. Pacanowski, and A. Wilkie (the Authors).
The Authors provide this document (the Work) under the Creative Commons CC BY-ND 3.0 license available online at <http://creativecommons.org/licenses/by-nd/3.0/>. The Authors further grant permission for reuse of images and text from the first page of the Work, provided that the reuse is for the purpose of promoting and/or summarizing the Work in scholarly venues and that any reuse is accompanied by a scientific citation to the Work.

