Status	Jan 2019 Voting Packet
Date of Last Update	2018/11/12
Person Assigned	David Clunie mailto:dclunie@dclunie.com
Submitter Name	Markus Daniel Herrmann mailto:mdherrmann@partners.org
Submission Date	2018/06/26

Correction Number CP-1843

Log Summary: JPEG-LS Planar Configuration constraints conflict with WSI, US, Enhanced Color MR and compressed RGB images

Name of Standard

PS3.3, PS3.5

Rationale for Correction:

Planar Configuration is required to be 0 for WSI, US, VL and Multi-frame SC True Color RGB images.

CP 1653 specified that Planar Configuration is required to be 1 for JPEG-LS (unlike most other JPEG family compression schemes, which specify 0).

YBR_FULL may be used as the Photometric Interpretation for some compressed images, or RGB when no decorrelating reversible color space transformation is used, e.g., JPEG-LS.

JPEG-LS specifies line and sample as well as component interleave methods in its bitstream, not that there is any need for the DICOM attributes to recapitulate that since it is handled by the codec.

RLE described the Planar Configuration as being required to be 1 and this is reflected in US image constraints.

Historically, the intent of constraining Planar Configuration at all was to avoid gratuitously different uncompressed encoding, and such constraints are meaningless for compressed images.

Clarify that Planar Configuration is required to be 0 instead of 1 (constrained in CP 1653) for JPEG-LS to be consistent with the other JPEG family schemes, and relax the constraint on RLE to be and to allow 0, except for US where it has traditionally been sent as 1.

The US requirements in Table C.8-23 for RGB and YBR_FULL are historically different and there is a large installed base, so these are not changed to be consistent with each other, since this would be too disruptive.

Ideally, the requirement for Planar Configuration to be present at all for compressed Transfer Syntaxes would be removed since it is meaningless, as is already stated in C.7.6.3.1.3, but this would be too disruptive.

[Ed.Note. A separate CP will address the matter of JPEG-LS decorrelating color space transformations, whether or not a SPIFF header may be present, the use (or not) of the 14495-2 reversible color transform from Annex F and Annex G.1.2.8, as well as other transforms such as YCoCg24, and corresponding new Photometric Interpretations.

Correction Wording:

*Amend DICOM PS3.3 as follows (changes to existing text are bold and **underlined** for additions and **~~struckthrough~~** for removals):*

A.8.5.4 Multi-frame True Color SC Image IOD Content Constraints

In the Image Pixel Module, the following constraints apply:

- Samples per Pixel (0028,0002) shall be 3
- Photometric Interpretation (0028,0004) shall be RGB for uncompressed or lossless compressed Transfer Syntaxes that do not have defined color space transformations, YBR_ICT for irreversible JPEG 2000 Transfer Syntaxes, YBR_RCT for reversible JPEG 2000 Transfer Syntaxes, YBR_PARTIAL_420 for MPEG2, MPEG-4 AVC/H.264, HEVC/H.265 Transfer Syntaxes and YBR_FULL_422 for JPEG lossy compressed Transfer Syntaxes and YBR_FULL or RGB for RLE Transfer Syntaxes

  Note

  Future lossless and lossy Transfer Syntaxes may lead to the need for new definitions and choices for Photometric Interpretation.

- ...
- Planar Configuration (0028,0006) shall be 0 (color-by-pixel) if Photometric Interpretation (0028,0004) is RGB

C.7.6.3.1.2 Photometric Interpretation

...

See ???? for additional restrictions imposed by compressed Transfer Syntaxes.

The following values are defined. Other values are permitted if supported by the Transfer Syntax but the meaning is not defined by this Standard.

Defined Terms:

RGB Pixel data represent a color image described by red, green, and blue image planes. The minimum sample value for each color plane represents minimum intensity of the color. This value may be used only when Samples per Pixel (0028,0002) has a value of 3. Planar Configuration (0028,0006) may be 0 or 1. May be used for pixel data in a Native (uncompressed) or Encapsulated (compressed) format; see ????.

YBR_FULL Pixel data represent a color image described by one luminance (Y) and two chrominance planes (CB and CR). This photometric interpretation may be used only when Samples per Pixel (0028,0002) has a value of 3. May be used for pixel data in a Native (uncompressed) or Encapsulated (compressed) format; see Section 8.2. Planar Configuration (0028,0006) may be 0 or 1.

This Photometric Interpretation is primarily used with RLE compressed bit streams, for which the Planar Configuration (0028,0006) **~~is required to~~may** be **0 or** 1; see PS3.5 Section 8.2.2 **and PS3.5 Section G.2**. When used in the US Image Module, the Planar Configuration (0028,0006) is required to be 1; see Section C.8.5.6.1.16 “Planar Configuration”.

Black is represented by Y equal to zero. The absence of color is represented by both CB and CR values equal to half full scale.

Note

In the case where Bits Allocated (0028,0100) has value of 8 half full scale is 128.

In the case where Bits Allocated (0028,0100) has a value of 8 then the following equations convert between RGB and YCBCR Photometric Interpretation.

Y = + .2990R + .5870G + .1140B

CB= - .1687R - .3313G + .5000B + 128

CR= + .5000R - .4187G - .0813B + 128

Note

The above is based on CCIR Recommendation 601-2 dated 1990.

YBR_FULL_422 The same as YBR_FULL except that the CB and CR values are sampled horizontally at half the Y rate and as a result there are half as many CB and CR values as Y values.

Planar Configuration (0028,0006) shall be 0. May be used for pixel data in a Native (uncompressed) or Encapsulated (compressed) format; see ????.

Note

1. This Photometric Interpretation is primarily used with JPEG compressed bit streams, but is also occasionally used for pixel data in a Native (uncompressed) format.

2. Though the chrominance channels are downsampled, there are still nominally three channels, hence Samples per Pixel (0028,0002) has a value of 3, not 2. I.e., for pixel data in a Native (uncompressed) format, the Value Length of Pixel Data (7FE0,0010) is not:

   Rows (0028,0010) * Columns (0028,0011) * Number of Frames (0028,0008) * Samples per Pixel (0028,0002) * (⌊(Bits Allocated (0028,0100)-1)/8⌋+1)

   padded to an even length, as it would otherwise be, but rather is:

   Rows (0028,0010) * Columns (0028,0011) * Number of Frames (0028,0008) * 2 * (⌊(Bits Allocated (0028,0100)-1)/8⌋+1)

   padded to an even length.

3. When used to describe JPEG compressed bit streams, the chrominance sub-sampling in the JPEG bit stream may differ from this description. E.g., though many JPEG codecs produce only horizontally sub-sampled chrominance components (4:2:2), some sub-sample vertically as well (4:2:0). Though inaccurate, the use of YBR_FULL_422 to describe both has proven harmless. For a discussion of the sub-sampling notation, see ???.

Two Y values shall be stored followed by one CB and one CR value. The CB and CR values shall be sampled at the location of the first of the two Y values. For each Row of Pixels, the first CB and CR samples shall be at the location of the first Y sample. The next CB and CR samples shall be at the location of the third Y sample etc.

Note

This subsampling sited on the even luminance pixels is often referred to as cosited sampling. The cositing applies when describing pixel data in a Native (uncompressed) form. When used to describe compressed bit streams, the siting depends on the compression scheme. E.g., for JPEG according to JFIF ???, the siting is midway between luminance samples, whereas for MPEG2 ???, the sampling is cosited with the even luminance pixels. See also ???.

YBR_PARTIAL_422 Retired. See PS3.3 2017b.

YBR_PARTIAL_420 Pixel data represent a color image described by one luminance (Y) and two chrominance planes (CB and CR).

This photometric interpretation may be used only when Samples per Pixel (0028,0002) has a value of 3. The CB and CR values are sampled horizontally and vertically at half the Y rate and as a result there are four times less CB and CR values than Y values.

Planar Configuration (0028,0006) shall be 0. Shall only be used for pixel data in an Encapsulated (compressed) format; see ????.

Note

This Photometric Interpretation is primarily used with MPEG compressed bit streams. For a discussion of the sub-sampling notation and siting, see ???.

Luminance and chrominance values are represented as follows:

1. black corresponds to Y = 16;
2. Y is restricted to 220 levels (i.e., the maximum value is 235);
3. CB and CR each has a minimum value of 16;
4. CB and CR are restricted to 225 levels (i.e., the maximum value is 240);
5. lack of color is represented by CB and CR equal to 128.

In the case where Bits Allocated (0028,0100) has value of 8 then the following equations convert between RGB and YBR_PARTIAL_420 Photometric Interpretation

Y = + .2568R + .5041G + .0979B + 16

CB= - .1482R - .2910G + .4392B + 128

CR= + .4392R - .3678G - .0714B + 128

Note

The above is based on CCIR Recommendation 601-2 dated 1990.

The CB and CR values shall be sampled at the location of the first of the two Y values. For the first Row of Pixels (etc.), the first CB and CR samples shall be at the location of the first Y sample. The next CB and CR samples shall be at the location of the third Y sample etc. The next Rows of Pixels containing CB and CR samples (at the same locations than for the first Row) will be the third etc.

YBR_ICT Irreversible Color Transformation:

Pixel data represent a color image described by one luminance (Y) and two chrominance planes (CB and CR).

This photometric interpretation may be used only when Samples per Pixel (0028,0002) has a value of 3. Planar Configuration (0028,0006) shall be 0. Shall only be used for pixel data in an Encapsulated (compressed) format; see ????.

Note

This Photometric Interpretation is primarily used with JPEG 2000 compressed bit streams.

Black is represented by Y equal to zero. The absence of color is represented by both CB and CR values equal to zero.

Regardless of the value of Bits Allocated (0028,0100), the following equations convert between RGB and YCBCR Photometric Interpretation.

Y = + .29900R + .58700G + .11400B

CB= - .16875R - .33126G + .50000B

CR= + .50000R - .41869G - .08131B

Note

1. The above is based on ??? (JPEG 2000).

2. In a JPEG 2000 bit stream, DC level shifting (used if the untransformed components are unsigned) is applied before forward color transformation, and the transformed components may be signed (unlike in JPEG ISO/IEC 10918-1).

3. In JPEG 2000, spatial down-sampling of the chrominance components, if performed, is signaled in the JPEG 2000 bit stream.

YBR_RCT Reversible Color Transformation:

Pixel data represent a color image described by one luminance (Y) and two chrominance planes (CB and CR).

This photometric interpretation may be used only when Samples per Pixel (0028,0002) has a value of 3. Planar Configuration (0028,0006) shall be 0. Shall only be used for pixel data in an Encapsulated (compressed) format; see ????.

Note

This Photometric Interpretation is primarily used with JPEG 2000 compressed bit streams.

Black is represented by Y equal to zero. The absence of color is represented by both CB and CR values equal to zero.

Regardless of the value of Bits Allocated (0028,0100), the following equations convert between RGB and YBR_RCT Photometric Interpretation.

$Y = \lfloor (R + 2G + B) / 4 \rfloor$ (Note: $\lfloor \ldots \rfloor$ mean floor)

$CB = B - G$

$CR = R - G$

The following equations convert between YBR_RCT and RGB Photometric Interpretation.

$G = Y - \lfloor (CR + CB) / 4 \rfloor$

$R = CR + G$

$B = CB + G$

Note

1. The above is based on ??? (JPEG 2000).

2. In a JPEG 2000 bit stream, DC level shifting (used if the untransformed components are unsigned) is applied before forward color transformation, and the transformed components may be signed (unlike in JPEG ISO/IEC 10918-1).

3. This photometric interpretation is a reversible approximation to the YUV transformation used in PAL and SECAM.

C.7.6.3.1.3 Planar Configuration

Planar Configuration (0028,0006) indicates whether the color pixel data are encoded color-by-plane or color-by-pixel. This Attribute shall be present if Samples per Pixel (0028,0002) has a value greater than 1. It shall not be present otherwise.

Enumerated Values:

0 The sample values for the first pixel are followed by the sample values for the second pixel, etc. For RGB images, this means the order of the pixel values encoded shall be R1, G1, B1, R2, G2, B2, …, etc.

1 Each color plane shall be encoded contiguously. For RGB images, this means the order of the pixel values encoded is R1, R2, R3, …, G1, G2, G3, …, B1, B2, B3, etc.

Note

Planar Configuration (0028,0006) is not meaningful when a compression Transfer Syntax is used that involves reorganization of sample components in the compressed bit stream. In such cases, since the Attribute is required to be present, then an appropriate value to use may be specified in the description of the Transfer Syntax in ????, though in all likelihood the value of the Attribute will be ignored by the receiving implementation.

C.8.5.6.1.16 Planar Configuration

For US Images, Planar Configuration (0028,0006) is specified to use the following values for specific Photometric Interpretations:

Table C.8-23. US Planar Configuration

Photometric Interpretation	Planar Configuration Value
RGB	0 - color-by-pixel, or 1 - color-by-plane
YBR_FULL	0 or 1 if uncompressed 0 if lossless JPEG, lossless JPEG-LS or reversible JPEG 2000 1 if RLE
YBR_FULL_422	0
YBR_RCT	0
YBR_ICT	0
YBR_PARTIAL_420	0

C.8.12.1.1 VL Image Module Attribute Descriptions

...

C.8.12.1.1.5 Planar Configuration

This value shall be present if Samples per Pixel (0028,0002) has a value greater than 1.

Enumerated Values:

0

Note

~~**The prohibition of a value of 1 for Planar Configuration (0028,0006) prevents the use of the RLE Transfer Syntax.**~~

C.8.12.4 Whole Slide Microscopy Image Module

Table C.8.12.4-1. Whole Slide Microscopy Image Module Attributes

Attribute Name	Tag	Type	Attribute Description
...			
Planar Configuration	(0028,0006)	1C	Indicates whether the pixel data are encoded color-by-plane or color-by-pixel. Required if Samples per Pixel (0028,0002) has a value greater than 1. **Enumerated Values:** **0** color-by-pixel

Attribute Name	Tag	Type	Attribute Description
..			

C.8.12.4.1.5 Photometric Interpretation and Samples Per Pixel

See Section C.7.6.3.1.2.

Enumerated Values for Photometric Interpretation (0028,0004):

MONOCHROME2
RGB
YBR_FULL_422
YBR_ICT
YBR_RCT

The value shall be appropriate to the compression Transfer Syntax used, if any, and shall be MONOCHROME2 or RGB for uncompressed or lossless compressed Transfer Syntaxes that do not have defined color space transformations, YBR_ICT for irreversible JPEG 2000 Transfer Syntaxes, YBR_RCT for reversible JPEG 2000 Transfer Syntaxes, and YBR_FULL_422 for JPEG lossy compressed Transfer Syntaxes.

Note

1. Future lossless and lossy Transfer Syntaxes may lead to the need for new definitions and choices for Photometric Interpretation. The Enumerated Values may therefore be extended with additional Photometric Interpretation values directly associated with new Transfer Syntaxes that are negotiated, and hence do not render existing implementations non-conformant.

2. Motion compression Transfer Syntaxes are not expected to be used for Whole Slide Imaging, so the use of YBR_PARTIAL_420 for MPEG2, MPEG-4 AVC/H.264 and HEVC/H.265 Transfer Syntaxes is not permitted.

3. ~~**The prohibition of a value of 1 for Planar Configuration (0028,0006) prevents the use of the RLE Transfer Syntax.**~~

...

Enumerated Values for Samples per Pixel (0028,0002) when Photometric Interpretation (0028,0004) is not MONOCHROME2:

3

C.8.13.1.1 Enhanced MR Image Module Attribute Description

C.8.13.1.1.2 Photometric Interpretation, Pixel Representation, Samples Per Pixel, Planar Configuration, Bits Allocated and Bits Stored

Table C.8-82 specifies the Enumerated Values and allowed combinations of Samples per Pixel (0028,0002), Planar Configuration (0028,0006), Pixel Representation (0028,0103), Bits Allocated (0028,0100) and Bits Stored (0028,0101) for each allowable Photometric Interpretation allowed by the IOD that invokes this Module.

Table C.8-82. Allowed Combinations of Attribute Values for Photometric Interpretation, Samples Per Pixel, Planar Configuration, Pixel Representation, Bits Allocated and Bits Stored

Photometric Interpretation	Samples per Pixel	Planar Configuration	Pixel Representation	Bits Allocated	Bits Stored
MONOCHROME2	1	-	0 or 1	8	8
MONOCHROME2	1	-	0 or 1	16	12, 16
RGB	3	0	0	8	8

Photometric Interpretation	Samples per Pixel	Planar Configuration	Pixel Representation	Bits Allocated	Bits Stored
YBR_ICT	3	0	0	8	8
YBR_RCT	3	0	0	8	8
YBR_PARTIAL_420	3	0	0	8	8
YBR_FULL_422	3	0	0	8	8
YBR_FULL	3	~~1~~**0**	0	8	8

C.8.17.2.1 Ophthalmic Photography Image Module Attribute Descriptions

C.8.17.2.1.3 Photometric Interpretation

Specifies the intended interpretation of the pixel data.

Enumerated Values:

MONOCHROME2
RGB
YBR_FULL_422
YBR_PARTIAL_420
YBR_ICT
YBR_RCT

When Samples per Pixel (0028,0002) is greater than 1, Photometric Interpretation (0028,0004) shall be RGB for uncompressed or lossless compressed Transfer Syntaxes that do not have defined color space transformations, YBR_ICT for irreversible JPEG 2000 Transfer Syntaxes, YBR_RCT for reversible JPEG 2000 Transfer Syntaxes, YBR_PARTIAL_420 for MPEG2, MPEG-4 AVC/H.264 and HEVC/H.265 Transfer Syntaxes and YBR_FULL_422 for JPEG lossy compressed Transfer Syntaxes.

Note

~~**The prohibition of a value of 1 for Planar Configuration (0028,0006) prevents the use of the RLE Transfer Syntax.**~~

*Amend DICOM PS3.5 as follows (changes to existing text are bold and **underlined** for additions and ~~**struckthrough**~~ for removals):*

8.2 Native or Encapsulated Format Encoding

Pixel data conveyed in the Pixel Data (7FE0,0010) may be sent either in a Native (uncompressed) Format or in an Encapsulated Format (e.g., compressed) defined outside the DICOM standard.

If Pixel Data (7FE0,0010) is sent in a Native Format, then the Photometric Interpretation (0028,0004) shall be other than:

- YBR_RCT
- YBR_ICT
- YBR_PARTIAL_420

Note

These values are not permitted because they are not encodable in an uncompressed form.

...

8.2.1 JPEG Image Compression

DICOM provides a mechanism for supporting the use of JPEG Image Compression through the Encapsulated Format (see PS3.3). Annex A defines a number of Transfer Syntaxes that reference the JPEG Standard and provide a number of lossless (bit preserving) and lossy compression schemes.

Note

The context where the usage of lossy compression of medical images is clinically acceptable is beyond the scope of the DICOM Standard. The policies associated with the selection of appropriate compression parameters (e.g., compression ratio) for JPEG lossy compression is also beyond the scope of this standard.

In order to facilitate interoperability of implementations conforming to the DICOM Standard that elect to use one or more of the Transfer Syntaxes for JPEG Image Compression, the following policy is specified:

- Any implementation that conforms to the DICOM Standard and has elected to support any one of the Transfer Syntaxes for lossless JPEG Image Compression, shall support the following lossless compression: The subset (first-order horizontal prediction [Selection Value 1) of JPEG Process 14 (DPCM, non-hierarchical with Huffman coding) (see Annex F).
- Any implementation that conforms to the DICOM Standard and has elected to support any one of the Transfer Syntaxes for 8-bit lossy JPEG Image Compression, shall support the JPEG Baseline Compression (coding Process 1).
- Any implementation that conforms to the DICOM Standard and has elected to support any one of the Transfer Syntaxes for 12-bit lossy JPEG Image Compression, shall support the JPEG Compression Process 4.

Note

The DICOM conformance statement shall differentiate whether or not the implementation is capable of simply receiving or receiving and processing JPEG encoded images (see ????).

The use of the DICOM Encapsulated Format to support JPEG Compressed Pixel Data requires that the Data Elements that are related to the Pixel Data encoding (e.g., Photometric Interpretation, Samples per Pixel, Planar Configuration, Bits Allocated, Bits Stored, High Bit, Pixel Representation, Rows, Columns, etc.) shall contain values that are consistent with the characteristics of the compressed data stream.

The requirements when using a Standard Photometric Interpretation (i.e., a Defined Term from ????) are specified in Table 8.2.1-1 and Table 8.2.1-2. No other Standard Photometric Interpretation values shall be used.

Table 8.2.1-1. Valid Values of Pixel Data Related Attributes for JPEG Lossy Transfer Syntaxes using Standard Photometric Interpretations

Photometric Interpretation	Transfer Syntax	Transfer Syntax UID	Samples per Pixel	Planar Configuration	Pixel Representation	Bits Allocated	Bits Stored	High Bit
MONOCHROME1 MONOCHROME2	JPEG Baseline	1.2.840.10008.1.2.4.50	1	absent	0	8	8	7
MONOCHROME1 MONOCHROME2	JPEG Extended	1.2.840.10008.1.2.4.51	1	absent	0	8	8	7
MONOCHROME1 MONOCHROME2	JPEG Extended	1.2.840.10008.1.2.4.51	1	absent	0	16	12	11
YBR_FULL_422	JPEG Baseline	1.2.840.10008.1.2.4.50	3	0	0	8	8	7

Table 8.2.1-2. Valid Values of Pixel Data Related Attributes for JPEG Lossless Transfer Syntaxes using Standard Photometric Interpretations

Photometric Interpretation	Transfer Syntax	Transfer Syntax UID	Samples per Pixel	Planar Configuration	Pixel Representation	Bits Allocated	Bits Stored	High Bit
MONOCHROME1 MONOCHROME2	JPEG Lossless, Non-Hierarchical JPEG Lossless, Non-Hierarchical, SV1	1.2.840.10008.1.2.4.57 1.2.840.10008.1.2.4.70	1	absent	0 or 1	8 or 16	1-16	0-15
PALETTE COLOR	JPEG Lossless, Non-Hierarchical JPEG Lossless, Non-Hierarchical, SV1	1.2.840.10008.1.2.4.57 1.2.840.10008.1.2.4.70	1	absent	0	8 or 16	1-16	0-15
YBR_FULL RGB	JPEG Lossless, Non-Hierarchical JPEG Lossless, Non-Hierarchical, SV1	1.2.840.10008.1.2.4.57 1.2.840.10008.1.2.4.70	3	0	0	8 or 16	1-16	0-15

The Pixel Data characteristics included in the JPEG Interchange Format shall be used to decode the compressed data stream.

Note

1. These requirements were formerly specified in terms of the "uncompressed pixel data from which the compressed data stream was derived". However, since the form of the "original" uncompressed data stream could vary between different implementations, this requirement is now specified in terms of consistency with what is encapsulated.

   When decompressing, should the characteristics explicitly specified in the compressed data stream (e.g., spatial sub-sampling or number of components or planar configuration) be inconsistent with those specified in the DICOM Data Elements, those explicitly specified in the compressed data stream should be used to control the decompression. The DICOM data elements, if inconsistent, can be regarded as suggestions as to the form in which an uncompressed Data Set might be encoded, subject to the general and IOD-specific rules for uncompressed Photometric Interpretation and Planar Configuration, which may require that decompressed data be converted to one of the permitted forms.

2. Those characteristics not explicitly specified in the compressed data stream (e.g., the color space of the compressed components, which is not specified in the JPEG Interchange Format), or implied by the definition of the compression scheme (e.g., always unsigned in JPEG), can therefore be determined from the DICOM Data Element in the enclosing Data Set. For example a Photometric Interpretation of "YBR_FULL_422" would describe the color space that is commonly used to lossy compress images using JPEG. It is unusual to use an RGB color space for lossy compression, since no advantage is taken of correlation between the red, green and blue components (e.g., of luminance), and poor compression is achieved.

3. The JPEG Interchange Format is distinct from the JPEG File Interchange Format (JFIF). The JPEG Interchange Format is defined in ??? section 4.9.1, and refers to the inclusion of decoding tables, as distinct from the "abbreviated format" in which these tables are not sent (and the decoder is assumed to already have them). The JPEG Interchange Format does NOT specify the color space. The JPEG File Interchange Format, not part of the original JPEG standard, but defined in ??? and ???, is often used to store JPEG bit streams in consumer format files, and does include the ability to specify the color space of the components. The JFIF APP0 marker segment is NOT required to be present in DICOM encapsulated JPEG bit streams, and should not be relied upon to recognize the color space. Its presence is not forbidden (unlike the JP2 information for JPEG 2000 Transfer Syntaxes), but it is recommended that it be absent.

4. Should the compression process be incapable of encoding a particular form of pixel data representation (e.g., JPEG cannot encode signed integers, only unsigned integers), then ideally only the appropriate form should be "fed" into the compression process. However, for certain characteristics described in DICOM Data Elements but not explicitly described in the compressed data stream (such as Pixel Representation), then the DICOM Data Element should be considered

to describe what has been compressed (e.g., the pixel data really is to be interpreted as signed if Pixel Representation so specifies).

5. DICOM Data Elements should not describe characteristics that are beyond the capability of the compression scheme used. For example, JPEG lossy processes are limited to 12 bits, hence the value of Bits Stored should be 12 or less. Bits Allocated is irrelevant, and is likely to be constrained by the Information Object Definition in ???? to values of 8 or 16. Also, JPEG compressed data streams are always color-by-pixel and should be specified as such (a decoder can essentially ignore this element however as the value for JPEG compressed data is already known).

6. If JPEG Compressed Pixel Data is decompressed and re-encoded in Native (uncompressed) form, then the Data Elements that are related to the Pixel Data encoding are updated accordingly. If color components are converted from YBR_FULL_422 to RGB during decompression and Native re-encoding, the Photometric Interpretation will be changed to RGB in the Data Set with the Native encoding.

8.2.2 Run Length Encoding Image Compression

DICOM provides a mechanism for supporting the use of Run Length Encoding (RLE) Image Compression, which is a byte oriented lossless compression scheme through the encapsulated Format (see ???? of this Standard). Annex G defines RLE Image Compression and its Transfer Syntax.

Note

The RLE Image Compression algorithm described in Annex G is the compression used in the TIFF 6.0 specification known as the "PackBits" scheme.

The use of the DICOM Encapsulated Format to support RLE Compressed Pixel Data requires that the Data Elements that are related to the Pixel Data encoding (e.g., Photometric Interpretation, Samples per Pixel, Planar Configuration, Bits Allocated, Bits Stored, High Bit, Pixel Representation, Rows, Columns, etc.) shall contain values that are consistent with the compressed data.

The requirements when using a Standard Photometric Interpretation (i.e., a Defined Term from PS.3. C.7.6.3.1.2) are specified in Table 8.2.2-1. No other Standard Photometric Interpretation values shall be used.

Table 8.2.2-1. Valid Values of Pixel Data Related Attributes for RLE Compression using Standard Photometric Interpretations

Photometric Interpretation	Samples per Pixel	Planar Configuration	Pixel Representation	Bits Allocated	Bits Stored	High Bit
MONOCHROME1 MONOCHROME2	1	absent	0 or 1	8 or 16	1-16	0-15
PALETTE COLOR	1	absent	0	8 or 16	1-16	0-15
YBR_FULL	3	**0 or** 1	0	8	1-8	0-7
RGB	3	**0 or** 1	0	8 or 16	1-16	0-15

Note

1. These requirements were formerly specified in terms of the "uncompressed pixel data from which the compressed data was derived". However, since the form of the "original" uncompressed data stream could vary between different implementations, this requirement is now specified in terms of consistency with what is encapsulated.

2. Those characteristics not implied by the definition of the compression scheme (e.g., always color-by-plane in RLE), can therefore be determined from the DICOM Data Element in the enclosing Data Set. For example a Photometric Interpretation of "YBR_FULL" would describe the color space that is commonly used to losslessly compress images using RLE. It is unusual to use an RGB color space for RLE compression, since no advantage is taken of correlation between the red, green and blue components (e.g., of luminance), and poor compression is achieved (note however that the conversion from RGB to YBR_FULL is itself lossy. A new photometric interpretation may be proposed in the future that allows lossless conversion from RGB and also results in better RLE compression ratios).

3. DICOM Data Elements should not describe characteristics that are beyond the capability of the compression scheme used. For example, RLE compressed data streams (using the algorithm mandated in the DICOM Standard) are always color-by-plane.

4. If RLE Compressed Pixel Data is decompressed and re-encoded in Native (uncompressed) form, then the Data Elements that are related to the Pixel Data encoding are updated accordingly. If color components are converted from YBR_FULL to RGB during decompression and Native re-encoding, the Photometric Interpretation will be changed to RGB in the Data Set with the Native encoding. It is permitted, however, to leave the YBR_FULL color components unconverted but decompressed in the Native format, in which case the Photometric Interpretation in the Data Set with the Native encoding would be YBR_FULL.

8.2.3 JPEG-LS Image Compression

DICOM provides a mechanism for supporting the use of JPEG-LS Image Compression through the Encapsulated Format (see PS3.3). Annex A defines a number of Transfer Syntaxes that reference the JPEG-LS Standard and provide a number of lossless (bit preserving) and lossy (near-lossless) compression schemes.

Note

The context where the usage of lossy (near-lossless) compression of medical images is clinically acceptable is beyond the scope of the DICOM Standard. The policies associated with the selection of appropriate compression parameters (e.g., compression ratio) for JPEG-LS lossy (near-lossless) compression is also beyond the scope of this standard.

The use of the DICOM Encapsulated Format to support JPEG-LS Compressed Pixel Data requires that the Data Elements that are related to the Pixel Data encoding (e.g., Photometric Interpretation, Samples per Pixel, Planar Configuration, Bits Allocated, Bits Stored, High Bit, Pixel Representation, Rows, Columns, etc.) shall contain values that are consistent with the characteristics of the compressed data stream. The Pixel Data characteristics included in the JPEG-LS Interchange Format shall be used to decode the compressed data stream.

The requirements when using a Standard Photometric Interpretation (i.e., a Defined Term from PS.3. C.7.6.3.1.2) are specified in Table 8.2.3-1. No other Standard Photometric Interpretation values shall be used.

Table 8.2.3-1. Valid Values of Pixel Data Related Attributes for JPEG-LS Compression using Standard Photometric Interpretations

Photometric Interpretation	Transfer Syntax	Transfer Syntax UID	Samples per Pixel	Planar Configuration	Pixel Representation	Bits Allocated	Bits Stored	High Bit
MONOCHROME1 MONOCHROME2	JPEG-LS Lossless JPEG-LS Lossy (Near-Lossless)	1.2.840.10008.1.2.4.80 1.2.840.10008.1.2.4.81	1	absent	0 or 1	8 or 16	2-16	1-15
PALETTE COLOR	JPEG-LS Lossless	1.2.840.10008.1.2.4.80	1	absent	0	8 or 16	2-16	1-15
YBR_FULL	JPEG-LS Lossless JPEG-LS Lossy (Near-Lossless)	1.2.840.10008.1.2.4.80 1.2.840.10008.1.2.4.81	3	~~1~~**0**	0	8	2-8	1-7
RGB	JPEG-LS Lossless JPEG-LS Lossy (Near-Lossless)	1.2.840.10008.1.2.4.80 1.2.840.10008.1.2.4.81	3	~~1~~**0**	0	8 or 16	2-16	1-15

Note

1. See also the notes in Section 8.2.1.

2. No color transformation Photometric Interpretation specific for JPEG-LS is currently defined in DICOM. Annex F of ISO 14495-2 describes a "Sample transformation for inverse colour transform" and a marker segment to encode its parameters, but this is not known to have been implemented. Common practice is to compress the RGB components unconverted, which sacrifices compression performance, and send the Photometric Interpretation as RGB. Though the YBR_RCT Photometric Interpretation and component conversion could theoretically be used, in the absence of DC shifting it results in signed values to be encoded, which are not supported by JPEG-LS.

3. If JPEG-LS Compressed Pixel Data is decompressed and re-encoded in Native (uncompressed) form, then the Data Elements that are related to the Pixel Data encoding are updated accordingly. If color components are converted from any other Photometric Interpretation to RGB during decompression and Native re-encoding, the Photometric Interpretation will be changed to RGB in the Data Set with the Native encoding.

4. The lower limit of 2 on Bits Stored (0028,0101) reflects the minimum JPEG-LS sample precision of 2.

The value of Planar Configuration (0028,0006) is irrelevant since the manner of encoding components is specified in the JPEG-LS bit stream as component, line or sample interleaved, hence it shall be set to 0.

8.2.4 JPEG 2000 Image Compression

DICOM provides a mechanism for supporting the use of JPEG 2000 Image Compression through the Encapsulated Format (see PS.3). Annex A defines a number of Transfer Syntaxes that reference the JPEG 2000 Standard and provide lossless (bit preserving) and lossy compression schemes.

Note

The context where the usage of lossy compression of medical images is clinically acceptable is beyond the scope of the DICOM Standard. The policies associated with the selection of appropriate compression parameters (e.g., compression ratio) for JPEG 2000 lossy compression are also beyond the scope of this standard.

The use of the DICOM Encapsulated Format to support JPEG 2000 Compressed Pixel Data requires that the Data Elements that are related to the Pixel Data encoding (e.g., Photometric Interpretation, Samples per Pixel, Planar Configuration, Bits Allocated, Bits Stored, High Bit, Pixel Representation, Rows, Columns, etc.) shall contain values that are consistent with the characteristics of the compressed data stream. The Pixel Data characteristics included in the JPEG 2000 bit stream shall be used to decode the compressed data stream.

The requirements when using a Standard Photometric Interpretation (i.e., a Defined Term from PS.3. C.7.6.3.1.2) are specified in Table 8.2.4-1. No other Standard Photometric Interpretation values shall be used.

Table 8.2.4-1. Valid Values of Pixel Data Related Attributes for JPEG 2000 Transfer Syntaxes using Standard Photometric Interpretations

Photometric Interpretation	Transfer Syntax	Transfer Syntax UID	Samples per Pixel	Planar Configuration	Pixel Representation	Bits Allocated	Bits Stored	High Bit
MONOCHROME1 MONOCHROME2	JPEG 2000 (Lossless Only) JPEG 2000	1.2.840.10008.1.2.4.90 1.2.840.10008.1.2.4.91	1	absent	0 or 1	8, 16, 24, 32 or 40	1-38	0-37
PALETTE COLOR	JPEG 2000 (Lossless Only)	1.2.840.10008.1.2.4.90	1	absent	0	8 or 16	1-16	0-15
YBR_RCT	JPEG 2000 (Lossless Only) JPEG 2000	1.2.840.10008.1.2.4.90 1.2.840.10008.1.2.4.91	3	0	0	8, 16, 24, 32 or 40	1-38	0-37
YBR_ICT	JPEG 2000	1.2.840.10008.1.2.4.91	3	0	0	8, 16, 24, 32 or 40	1-38	0-37
RGB	JPEG 2000 (Lossless Only) JPEG 2000	1.2.840.10008.1.2.4.90 1.2.840.10008.1.2.4.91	3	0	0	8, 16, 24, 32 or 40	1-38	0-37

Photometric Interpretation	Transfer Syntax	Transfer Syntax UID	Samples per Pixel	Planar Configuration	Pixel Representation	Bits Allocated	Bits Stored	High Bit
YBR_FULL	JPEG 2000 (Lossless Only) JPEG 2000	1.2.840.10008.1.2.4.90 1.2.840.10008.1.2.4.91	3	0	0	8, 16, 24, 32 or 40	1-38	0-37

Note

These requirements are specified in terms of consistency with what is encapsulated, rather than in terms of the uncompressed pixel data from which the compressed data stream may have been derived.

When decompressing, should the characteristics explicitly specified in the compressed data stream be inconsistent with those specified in the DICOM Data Elements, those explicitly specified in the compressed data stream should be used to control the decompression. The DICOM data elements, if inconsistent, can be regarded as suggestions as to the form in which an uncompressed Data Set might be encoded, subject to the general and IOD-specific rules for uncompressed Photometric Interpretation and Planar Configuration, which may require that decompressed data be converted to one of the permitted forms.

The JPEG 2000 bit stream specifies whether or not a reversible or irreversible multi-component (color) transformation [ISO 15444-1 Annex G], if any, has been applied. If no multi-component transformation has been applied, then the components shall correspond to those specified by the DICOM Attribute Photometric Interpretation (0028,0004). If the JPEG 2000 Part 1 reversible multi-component transformation has been applied then the DICOM Attribute Photometric Interpretation (0028,0004) shall be YBR_RCT. If the JPEG 2000 Part 1 irreversible multi-component transformation has been applied then the DICOM Attribute Photometric Interpretation (0028,0004) shall be YBR_ICT.

Note

1. For example, single component may be present, and the Photometric Interpretation (0028,0004) may be MONOCHROME2.

2. The application of a JPEG 2000 Part 1 reversible multi-component transformation is signaled in the JPEG 2000 bit stream by a value of 1 rather than 0 in the SGcod Multiple component transformation type of the COD marker segment [ISO 15444-1 Table A.17]. No other value of Photometric Interpretation than YBR_RCT or YBR_ICT is permitted when SGcod Multiple component transformation type is 1.

3. Though it would be unusual, would not take advantage of correlation between the red, green and blue components, and would not achieve effective compression, a Photometric Interpretation of RGB could be specified as long as no multi-component transformation [ISO 15444-1 Annex G] was specified by the JPEG 2000 bit stream. Alternative methods of decorrelation of the color components than those specified in [ISO 15444-1 Annex G] are permitted as defined in PS3.3, such as a Photometric Interpretation of YBR_FULL; this may be useful when converting existing YBR_FULL Pixel Data (e.g., in a different Transfer Syntax) without further loss.

   In either case (Photometric Interpretation of RGB or YBR_FULL), the value of SGcod Multiple component transformation type would be 0.

   ???? may constrain the values of Photometric Interpretation for specific IODs.

4. Despite the application of a multi-component color transformation and its reflection in the Photometric Interpretation attribute, the "color space" remains undefined. There is currently no means of conveying "standard color spaces" either by fixed values (such as sRGB) or by ICC profiles. Note in particular that the JP2 file header is not sent in the JPEG 2000 bit stream that is encapsulated in DICOM.

5. If JPEG 2000 Compressed Pixel Data is decompressed and re-encoded in Native (uncompressed) form, then the Data Elements that are related to the Pixel Data encoding are updated accordingly. If color components are converted from YBR_ICT or YBR_RCT to RGB during decompression and Native re-encoding, the Photometric Interpretation will be changed to RGB in the Data Set with the Native encoding.

6. The upper limit of 40 on Bits Allocated (0028,0100) and 38 on Bits Stored (0028,0101) reflects the maximum JPEG 2000 sample precision of 38 and the DICOM requirement to describe Bits Allocated (0028,0100) as multiples of bytes (octets).

The JPEG 2000 bit stream is capable of encoding both signed and unsigned pixel values, hence the value of Pixel Representation (0028,0103) may be either 0 or 1 for monochrome Photometric Interpretations depending on what has been encoded (as specified in the SIZ marker segment in the precision and sign of component parameter).

The value of Planar Configuration (0028,0006) is irrelevant since the manner of encoding components is specified in the JPEG 2000 standard, hence it shall be set to 0.

G Encapsulated RLE Compressed Images (Normative)

G.1 Summary

This annex describes how to apply RLE Image Compression to an image or an individual frame of a multi-frame image. This method can be used for any image, independent of the values of the data elements that describe the image (i.e., Photometric Interpretation (0028,0004) and Bits Stored (0028,0101)).

RLE Image Compression consists of the following steps:

1. The image is converted to a sequence of Composite Pixel Codes (see ????).
2. The Composite Pixel Codes are used to generate a set of Byte Segments (see Section G.2).
3. Each Byte Segment is RLE compressed to produce a RLE Segment (see Section G.4).
4. The RLE Header is appended in front of the concatenated RLE Segments (see Section G.5).

G.2 Byte Segments

A Byte Segment is a series of bytes generated by decomposing the Composite Pixel Code (see ????).

If the Composite Pixel Code is not an integral number of bytes in size, sufficient Most Significant zero bits are added to make it an integral byte size. This is known as the Padded Composite Pixel Code.

The first Segment is generated by stripping off the most significant byte of each Padded Composite Pixel Code and ordering these bytes sequentially. The second Segment is generated by repeating this process on the stripped Padded Composite Pixel Code continuing until the last Pixel Segment is generated by ordering the least significant byte of each Padded Component Pixel Code sequentially.

Note

1. If Photometric Interpretation (0028, 0004) equals RGB and Bits Allocated equals 8, then three Segments are generated. The first one holds all the Red values, the second all the Green values, and the third all the Blue values.
2. The use of separate segments implies that the Planar Configuration (0028,0006) **~~will always~~**<u>**could theoretically**</u> be 1 for RLE compressed images<u>**, but for consistency with other Encapsulated (compressed) Transfer Syntaxes and restrictions on Planar Configuration in many IODs, it may be 0**</u>.