MIIT 2009

DICOM Support for Compression: More than JPEG

Dr. David Clunie, MBBS, FRACR

CTO, RadPharm, Inc.

RADPHARM
Princeton Radiology Pharmaceutical Research

Policy versus Technology

- DICOM provides technology
  - only provides schemes to use
  - does NOT “approve” their use
  - does NOT recommend use for any purpose
  - goal is interoperability, not quality
- Regulators largely avoid the issue
  - in US, MQSA forbids lossy for mammo primary reads
- Discretion of the user
- Guidance from (some) professional societies

Regulation of compression

- Impact on device classification
- Guidance draft 8/93
- Classification final rule 4/98
  - Reversible (lossless) -> Class I (exempt)
  - Irreversible (lossy) -> Class II (510k)
- Revision 1/00 (FDMA) – irreversible exempted
- FDA – on-screen labeling
  - irreversible compression has been applied
  - approximate compression ratio
- ACR – amount and method of data compression

Standards Prior to DICOM

- ACR-NEMA 1985 – no compression
- PS2 Data Compression Standard
- “Toolkit” of various “components”
  - image conversion techniques
    - DPCM, DCT, S-transform, pyramid transform
  - coding schemes
    - Huffman, RLE, LZ, perimeter
- Rarely used
  - Siemens private CT compression

DICOM uses standards

- “Standard”
  - ISO/IEC 10918-1 / ITU T.81 – JPEG
  - ISO/IEC 14495-1 / ITU T.87 – JPEG-LS
  - ISO/IEC 15444-1 / ITU T.800 – JPEG 2000
  - ISO/IEC 13818-2 / ITU H.262 – MPEG-2
- Not so “standard”
  - RLE (for ultrasound) – TIFF PackBits

“Original” JPEG

- DICOM in 1993 included all JPEG processes
  - 29 processes in 18 different Transfer Syntaxes
- Only one consumer process common
  - 8 bit per channel color DCT Huffman
- Only one DICOM process common
  - Lossless 16 bit Selection Value 1 Huffman
- Lossy compression of CT/MR
  - 12 bit per channel monochrome DCT Huffman
- Unused processes “retired” 2002/01
  - 4 remaining

JPEG

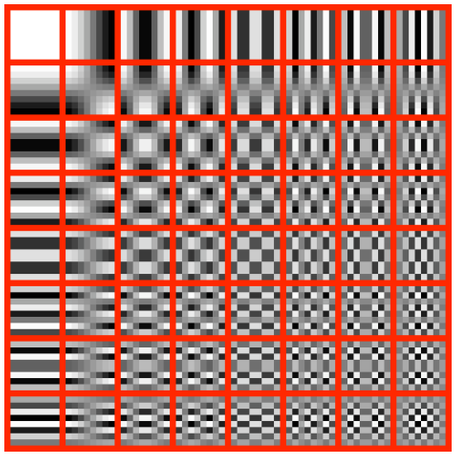

- Lossy process
  - Color space transform RGB -> YCbCr
  - 8x8 blocks
  - Discrete Cosine Transform (DCT)
  - Quantization of DCT coefficients
  - Huffman entropy coding of block

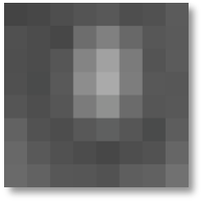

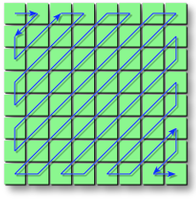

- Lossless process
  - Entire image (no blocks)
  - Difference coding (SV1 – previous pixel)
  - Huffman entropy coding

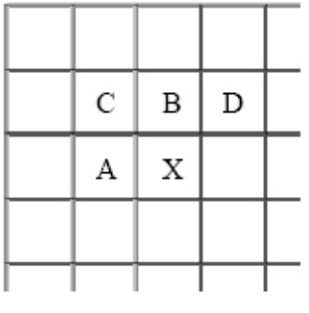

http://www.wikipedia.org/

JPEG-LS

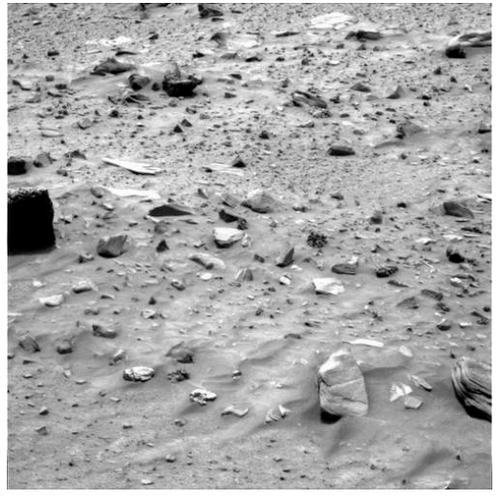

- Distinct from “lossless JPEG”
- ISO 14495-1 / ITU T.87
- Based on HP LOCO-I (“LOw COmplexity LOssless COmpression for Images”) used in Mars Spirit Rover
- Fast, simple, better compression than lossless JPEG
- Performs as well as more complex, slower lossless JPEG 2000
- Prediction, residual modeling, context-based coding, Rice-Golomb entropy coder
- Run-length mode handles large uniform areas
- Added to DICOM 2000/09
- Has *not* seen widespread adoption
- Benefit over JPEG lossless not enough to justify ?

http://marsrover.nasa.gov/

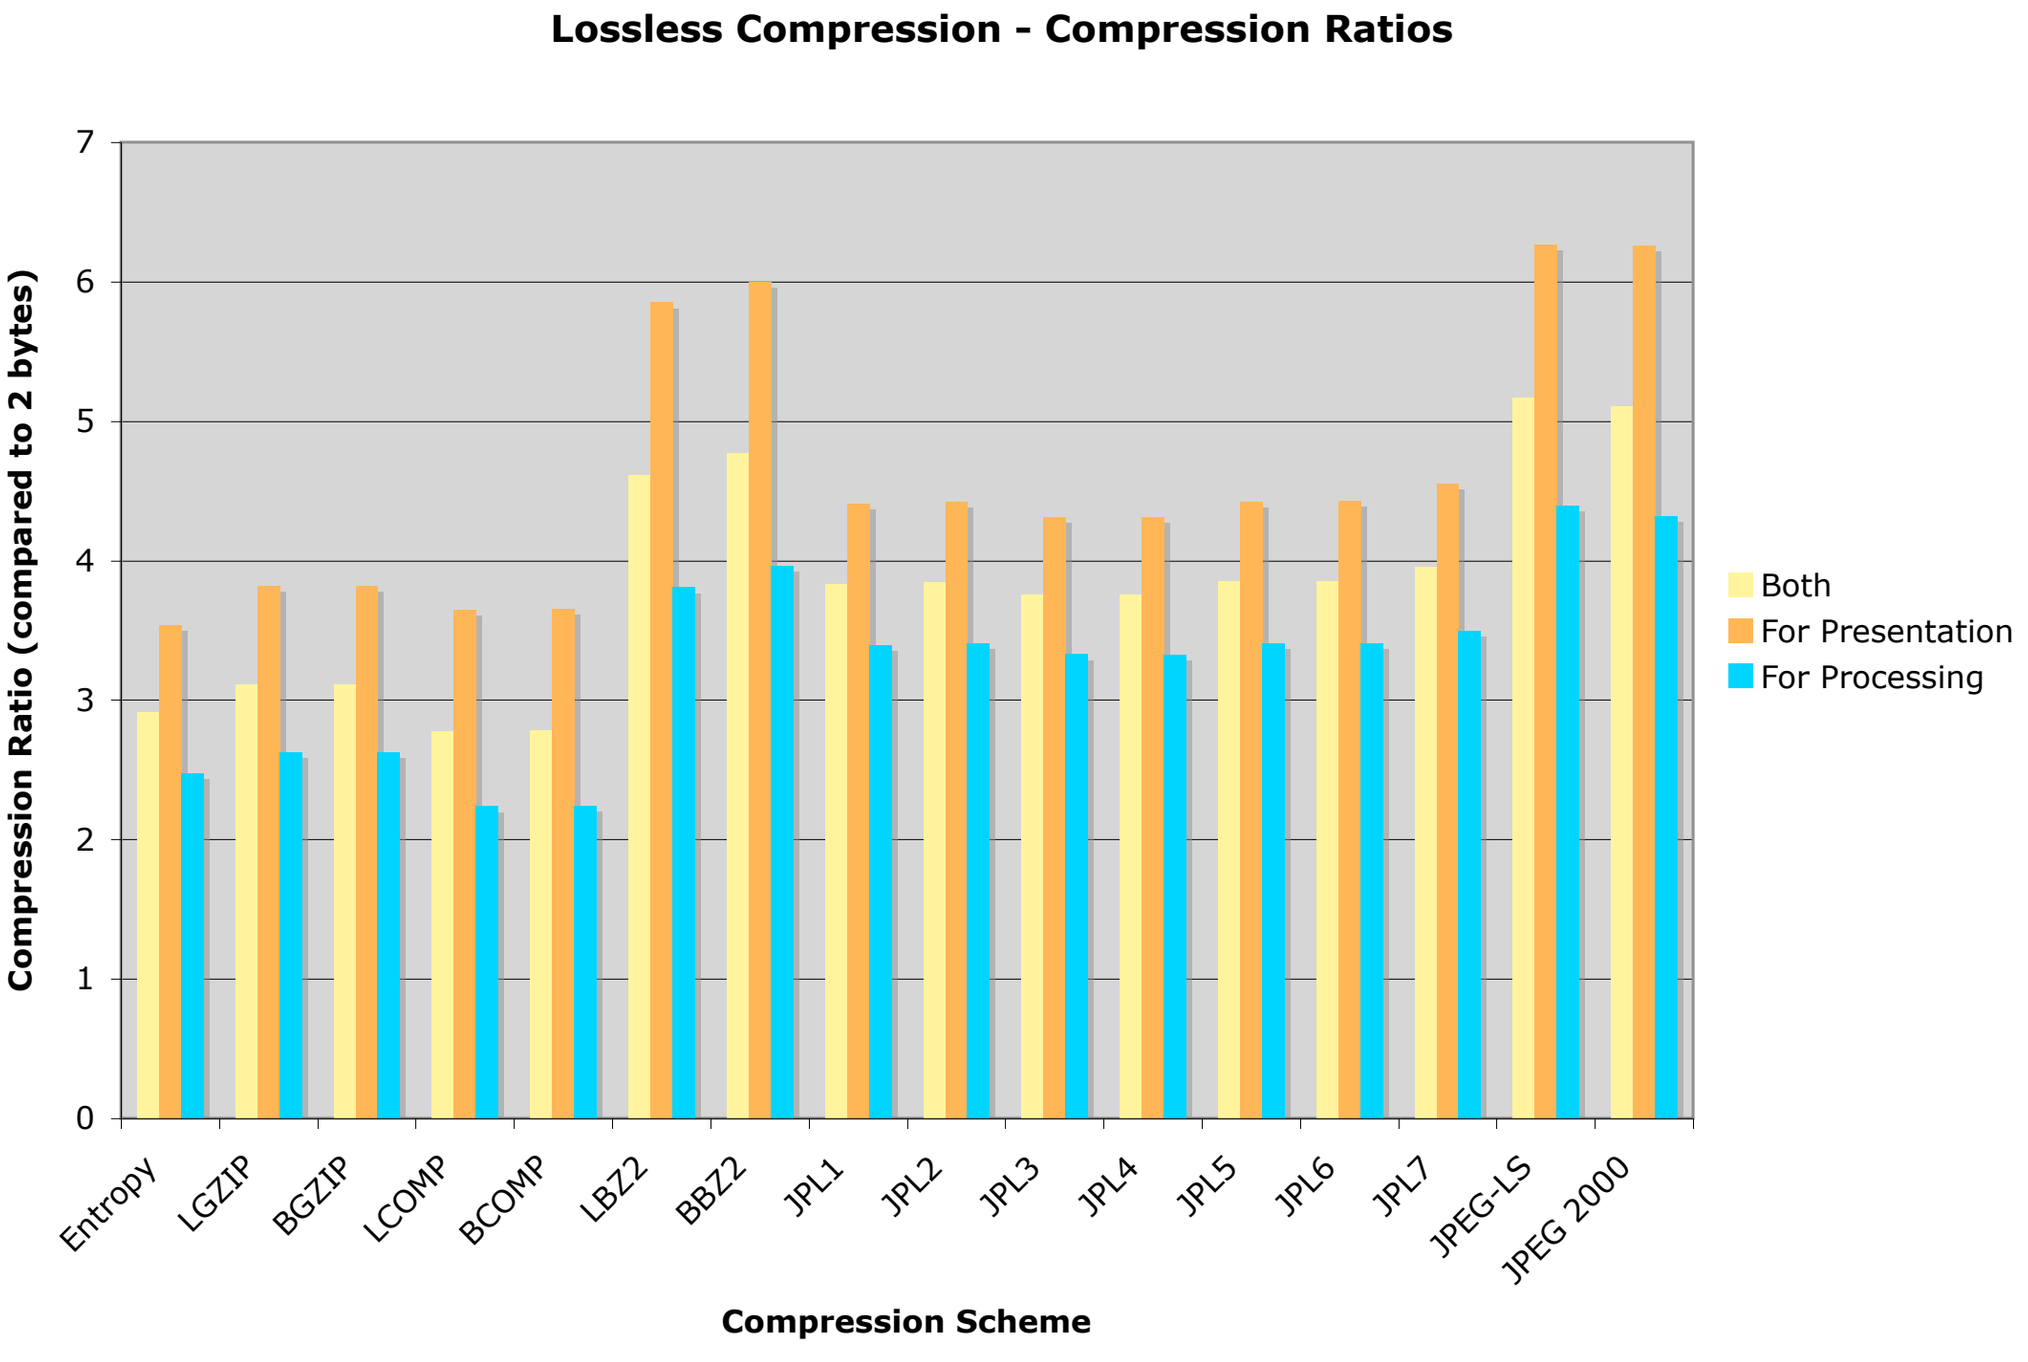

20 uncropped pairs of MG images from 3 vendors (40 images)

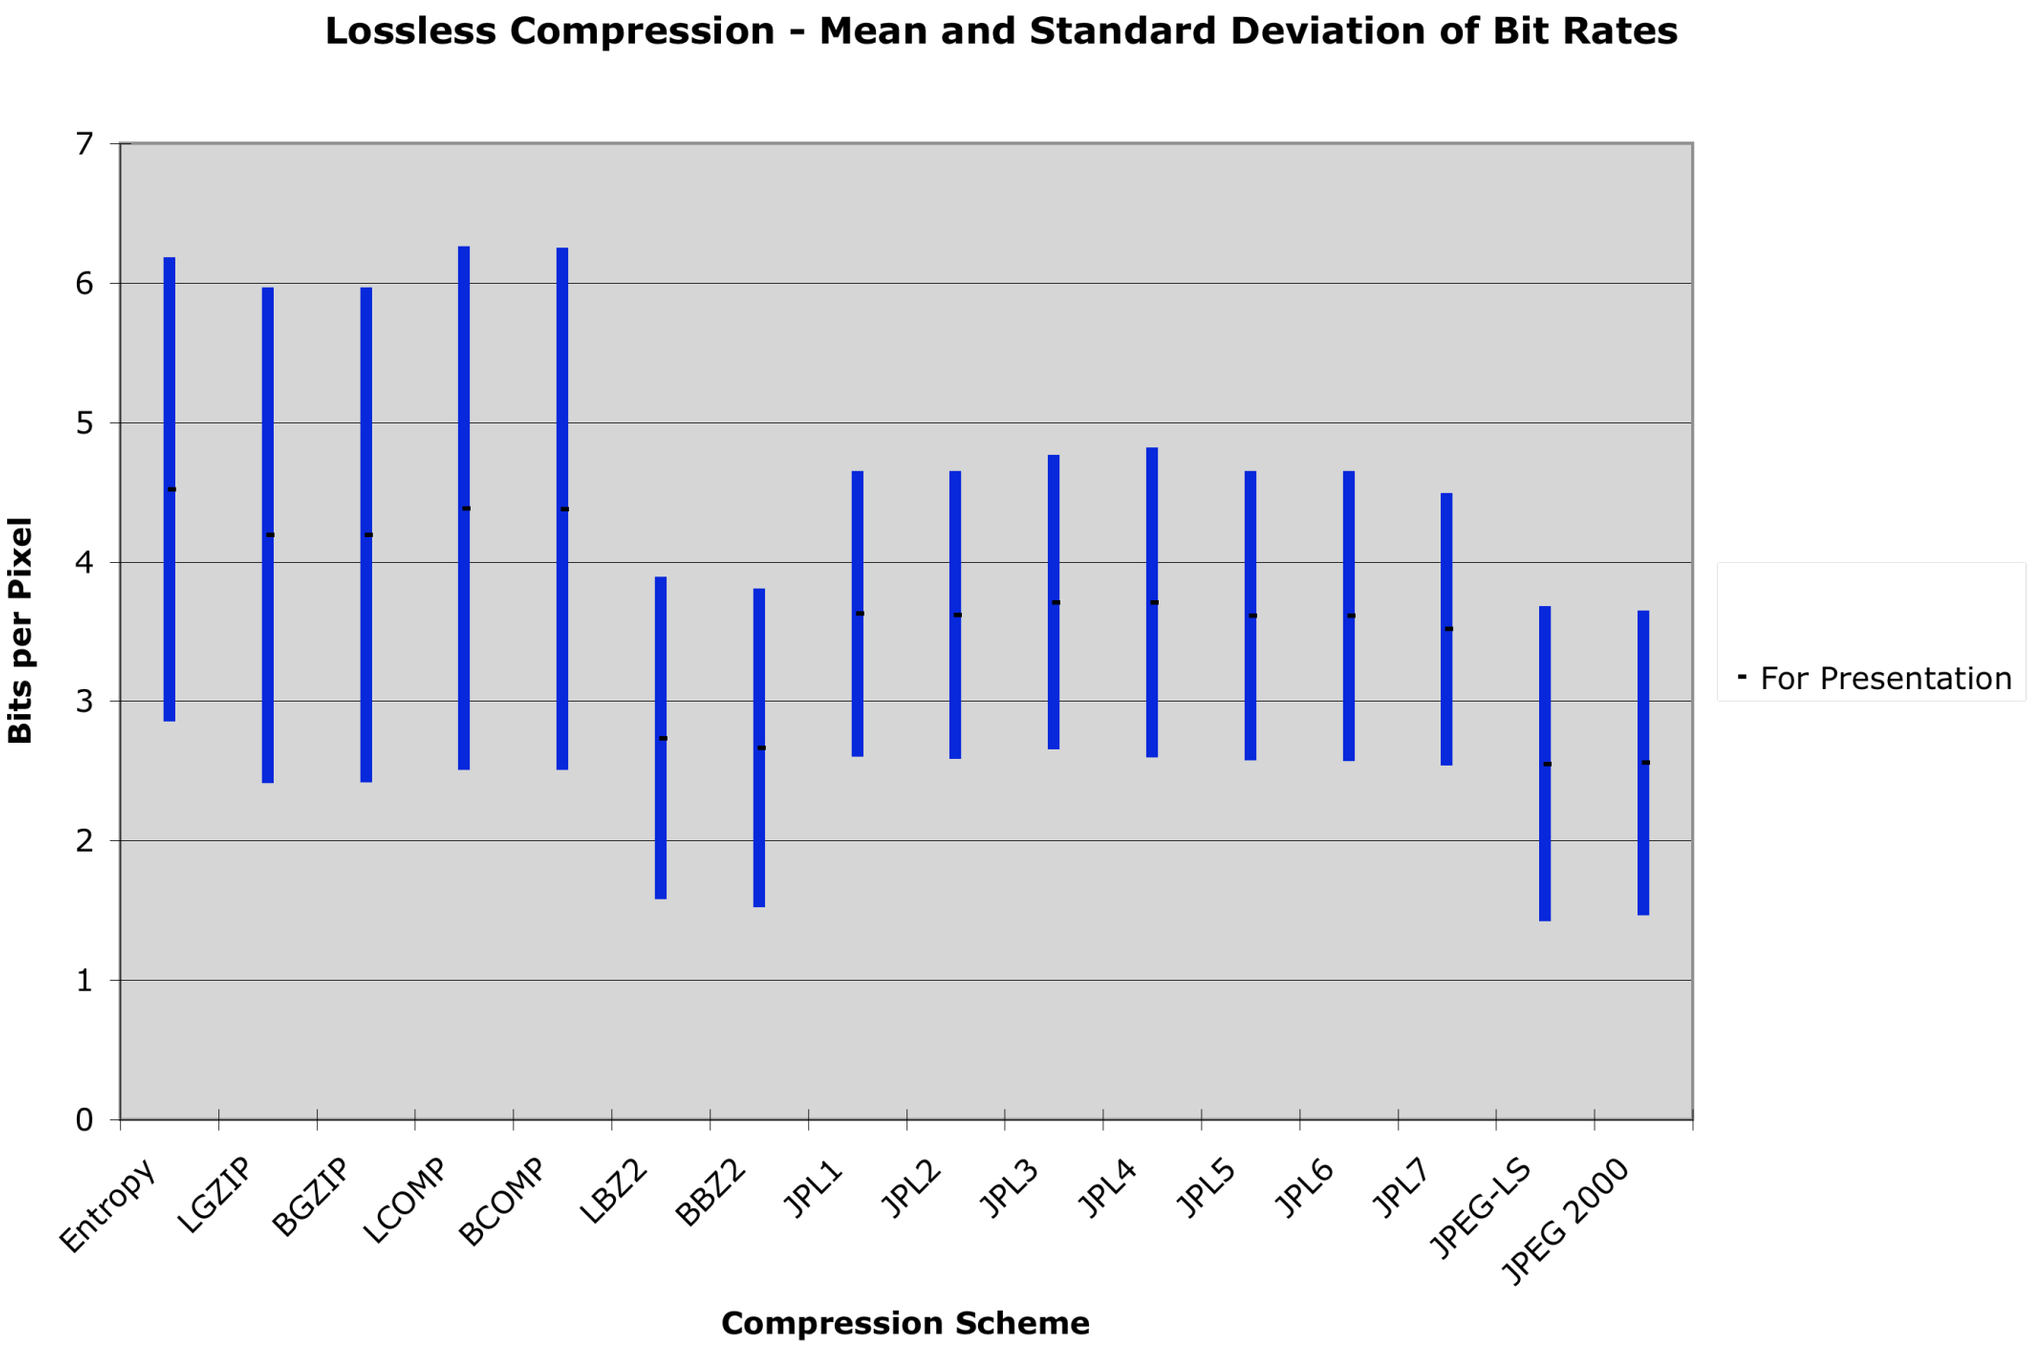
Lossless Compression - Mean and Standard Deviation of Bit Rates
Bits per Pixel
7
6
5
4
3
2
1
0
Entropy
LGZIP
BGZIP
LCOMP
BCOMP
LBZ2
BBZ2
JPL1
JPL2
JPL3
JPL4
JPL5
JPL6
JPL7
JPEG-LS
JPEG 2000
Compression Scheme
- For Presentation

Variation in compressibility

JPEG-LS Lossless

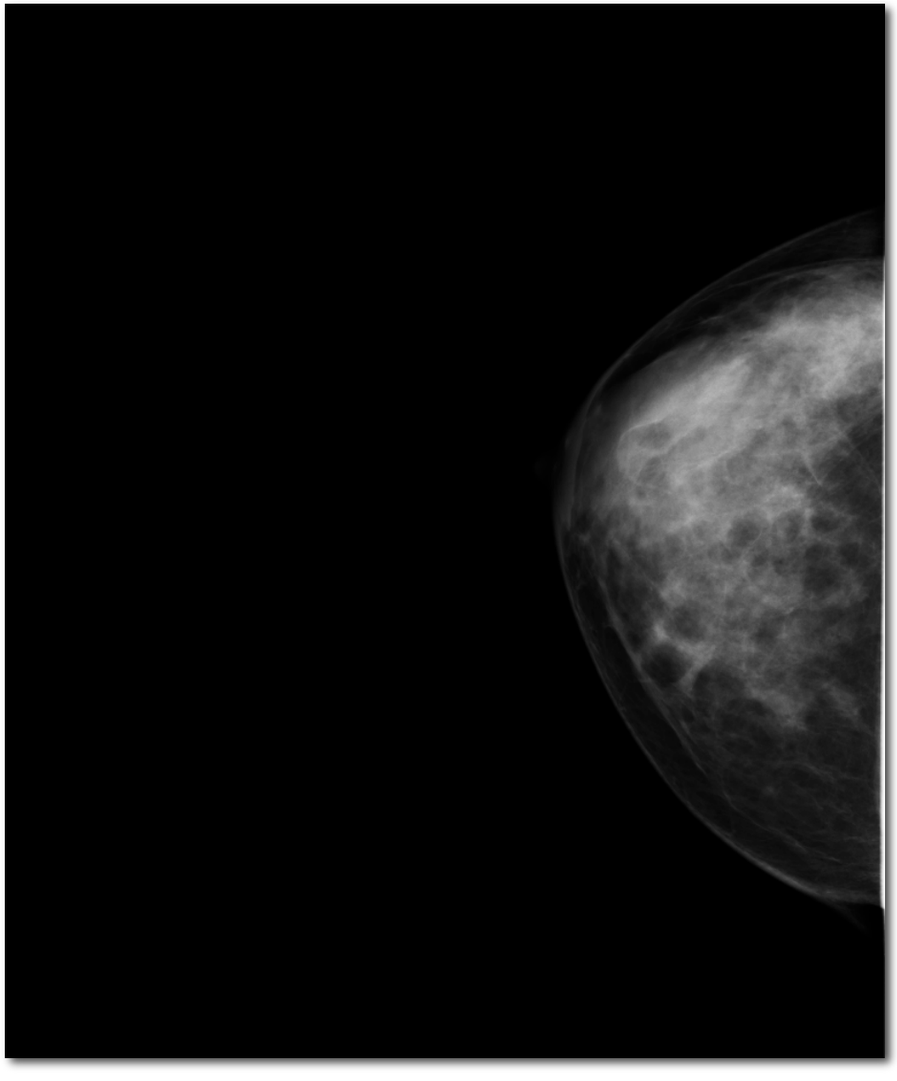

Best - CR 12.9

Worst - CR 3.19

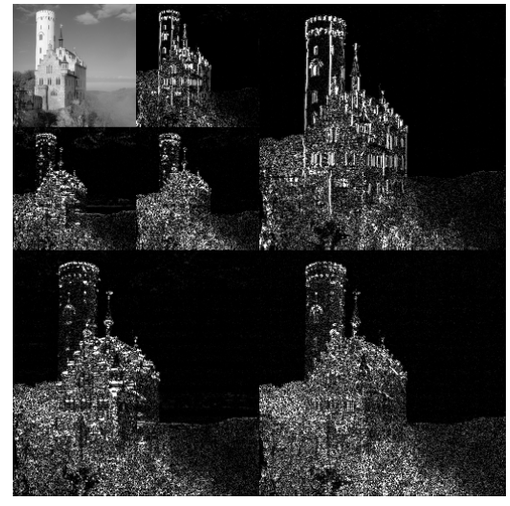

JPEG 2000 (J2K)

- Wavelet-based
- Full frame not blocks (may be tiled)
- Full 16 bits (not 12 like JPEG lossy)
- Reversible and irreversible forms of *same process*
  - DICOM Transfer Syntaxes for each
- Sophisticated
  - relatively complex, slow, but many features (ROI, progressive)
- Different artifact than DCT JPEG
  - blurring and “rice grains” versus blocks
- Effective
  - lossy – large matrix continuous tone images
  - lossless – state of the art (== JPEG-LS)

http://www.wikipedia.org/

Single frame lossless

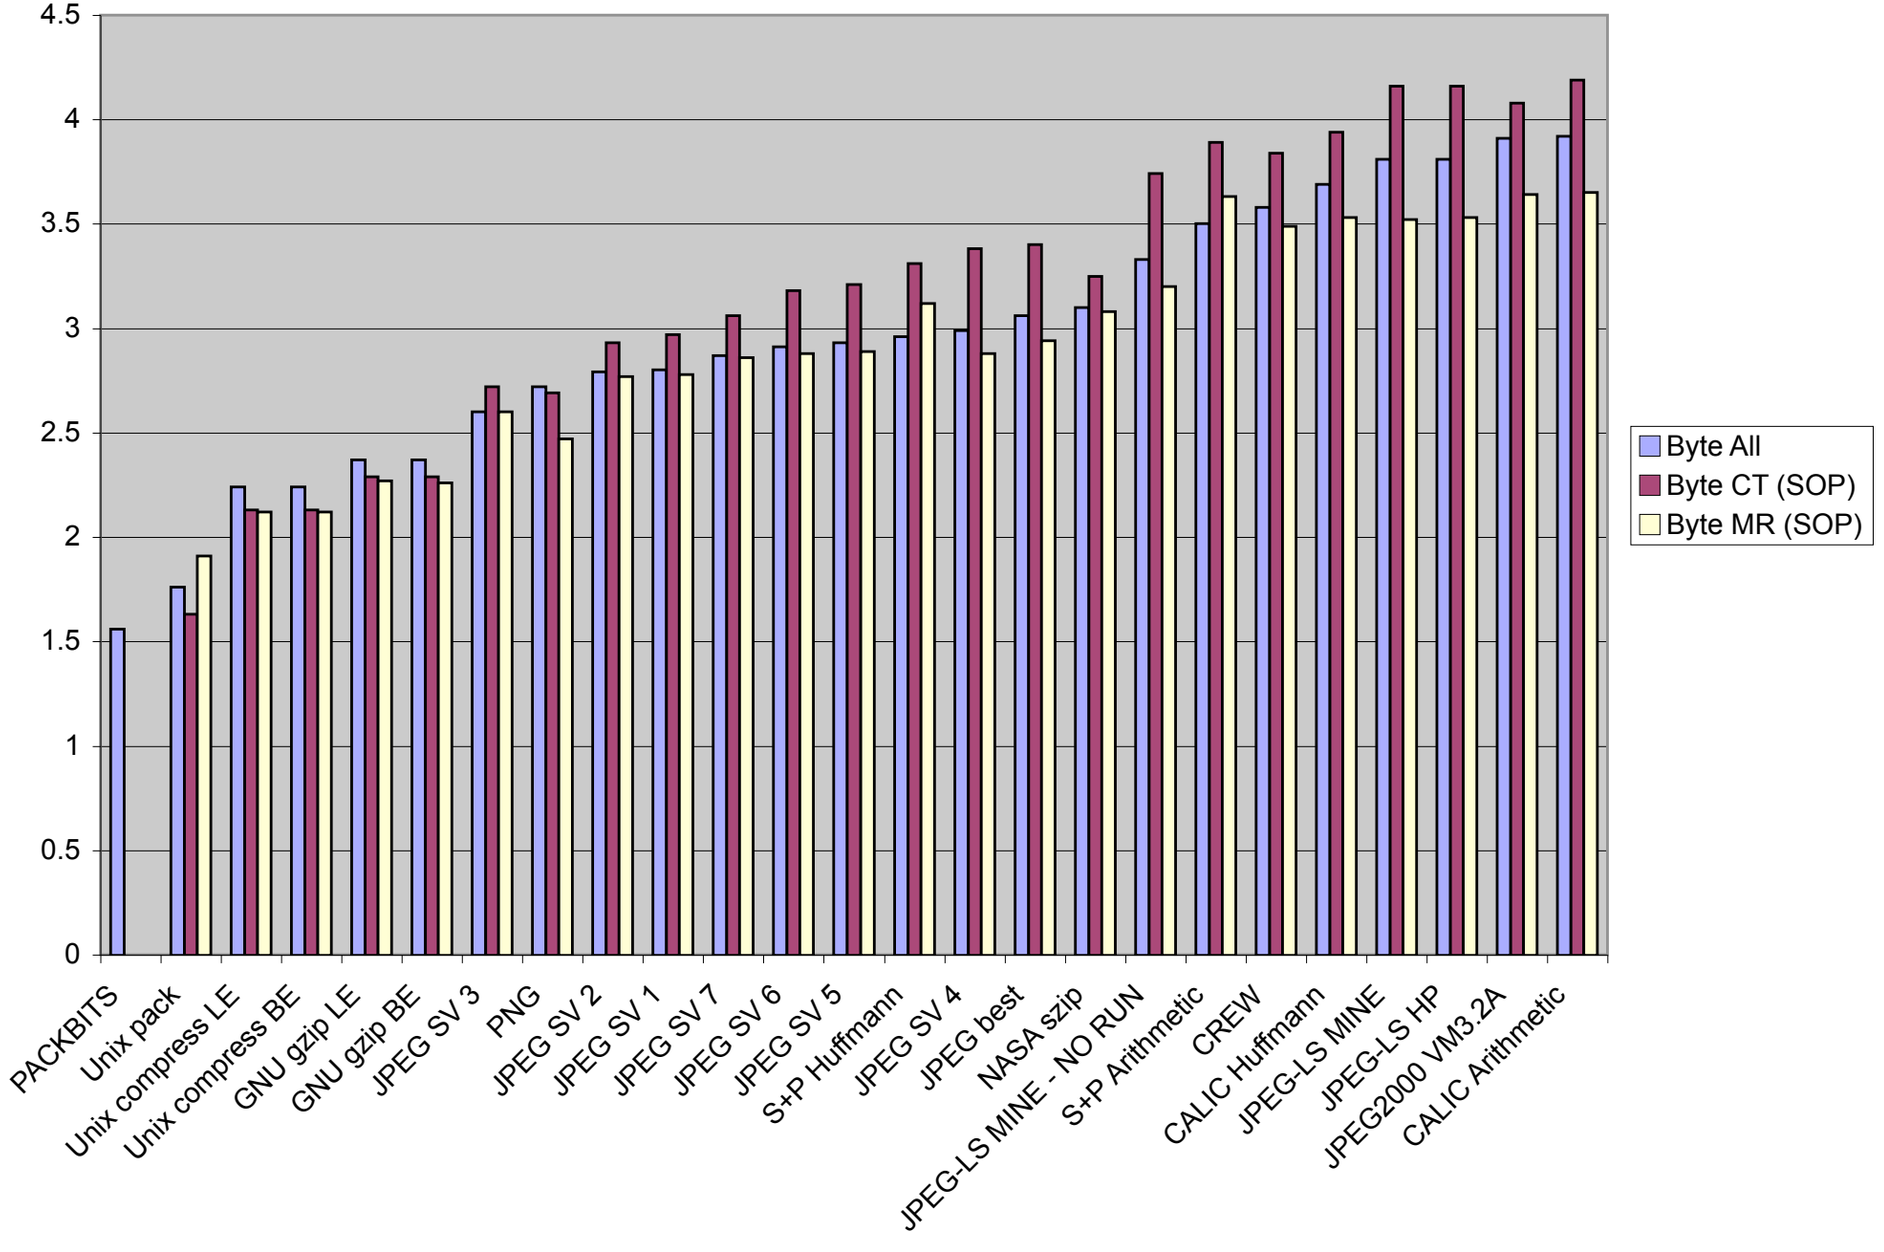

Lossless JPEG 2000 Compression (Alexis Tzannes, Aware, 2003)

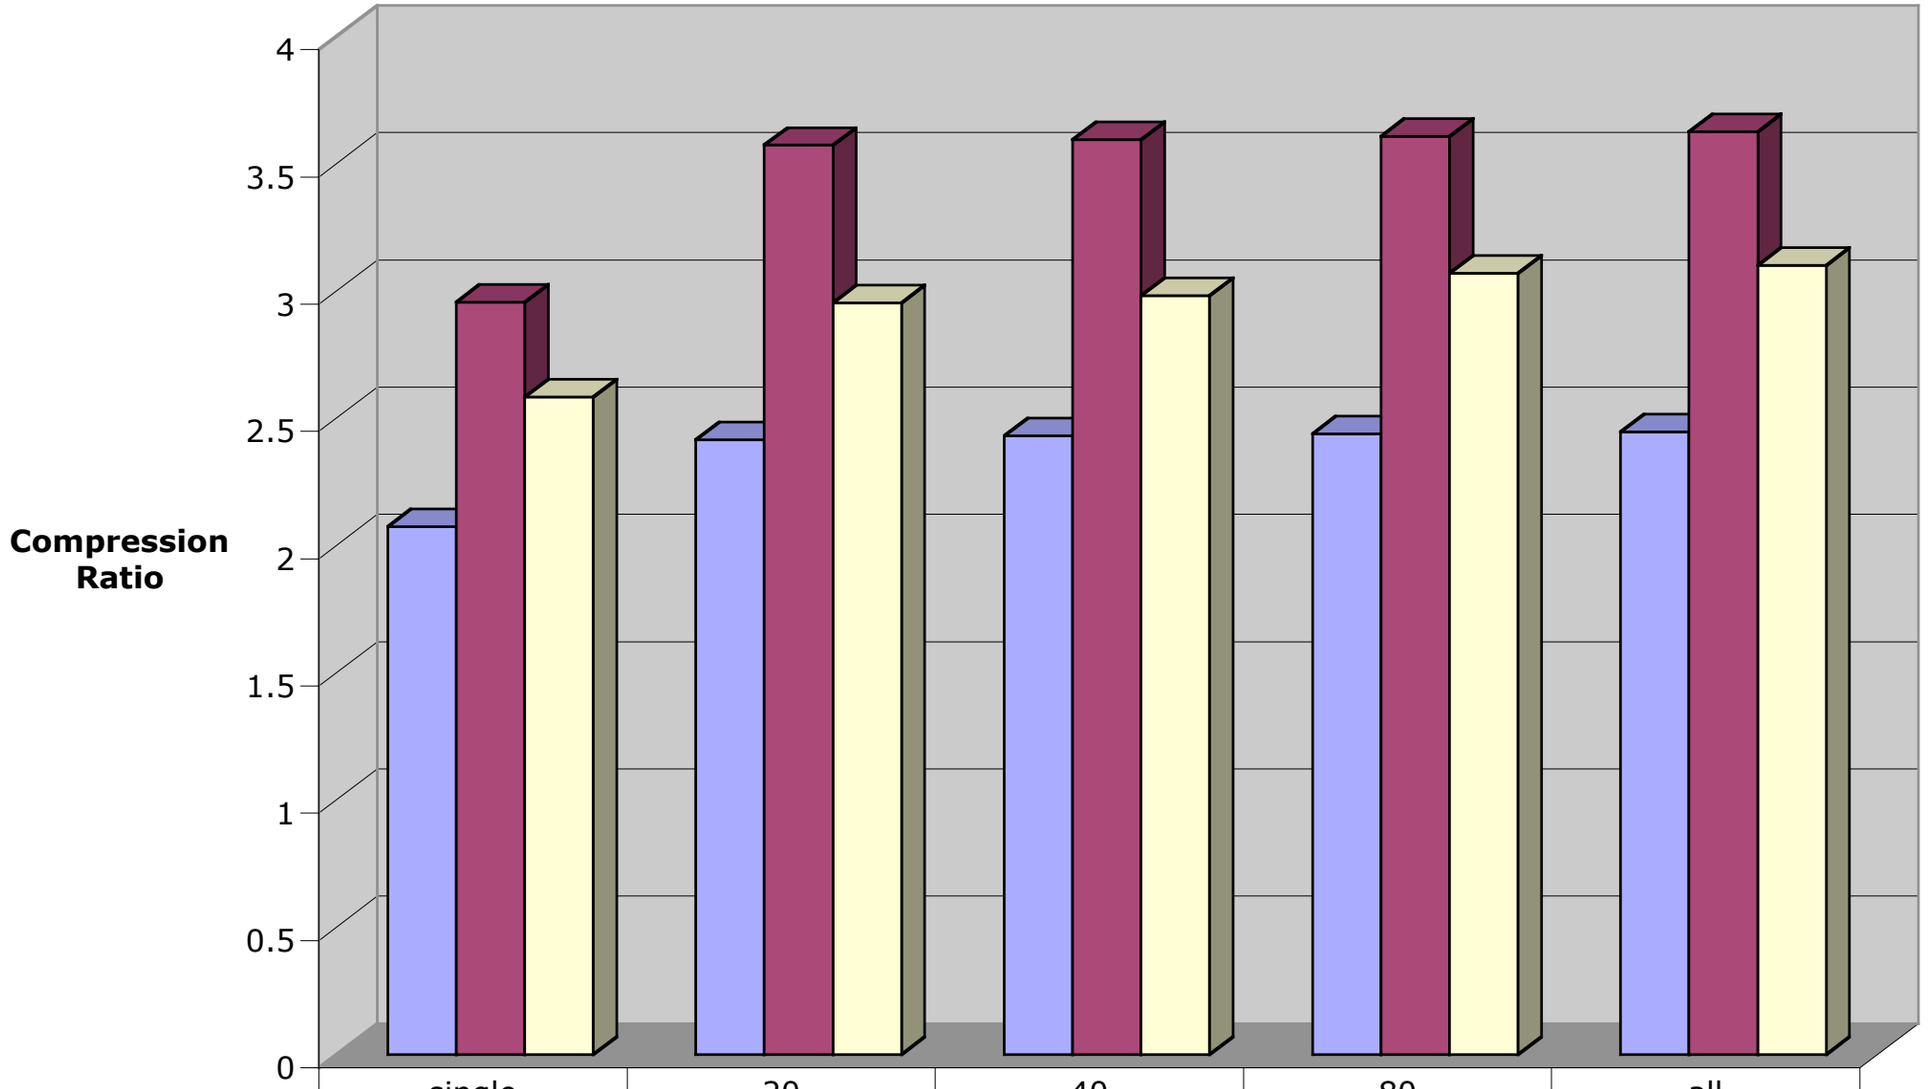

	single	20	40	80	all
127x256x8 7.9MB	2.073490814	2.415902141	2.430769231	2.438271605	2.445820433
449x512x16 224MB	2.955145119	3.572567783	3.595505618	3.607085346	3.624595469
620x512x16 310MB	2.583333333	2.952380952	2.980769231	3.069306931	3.1

Slices in 3rd dimension

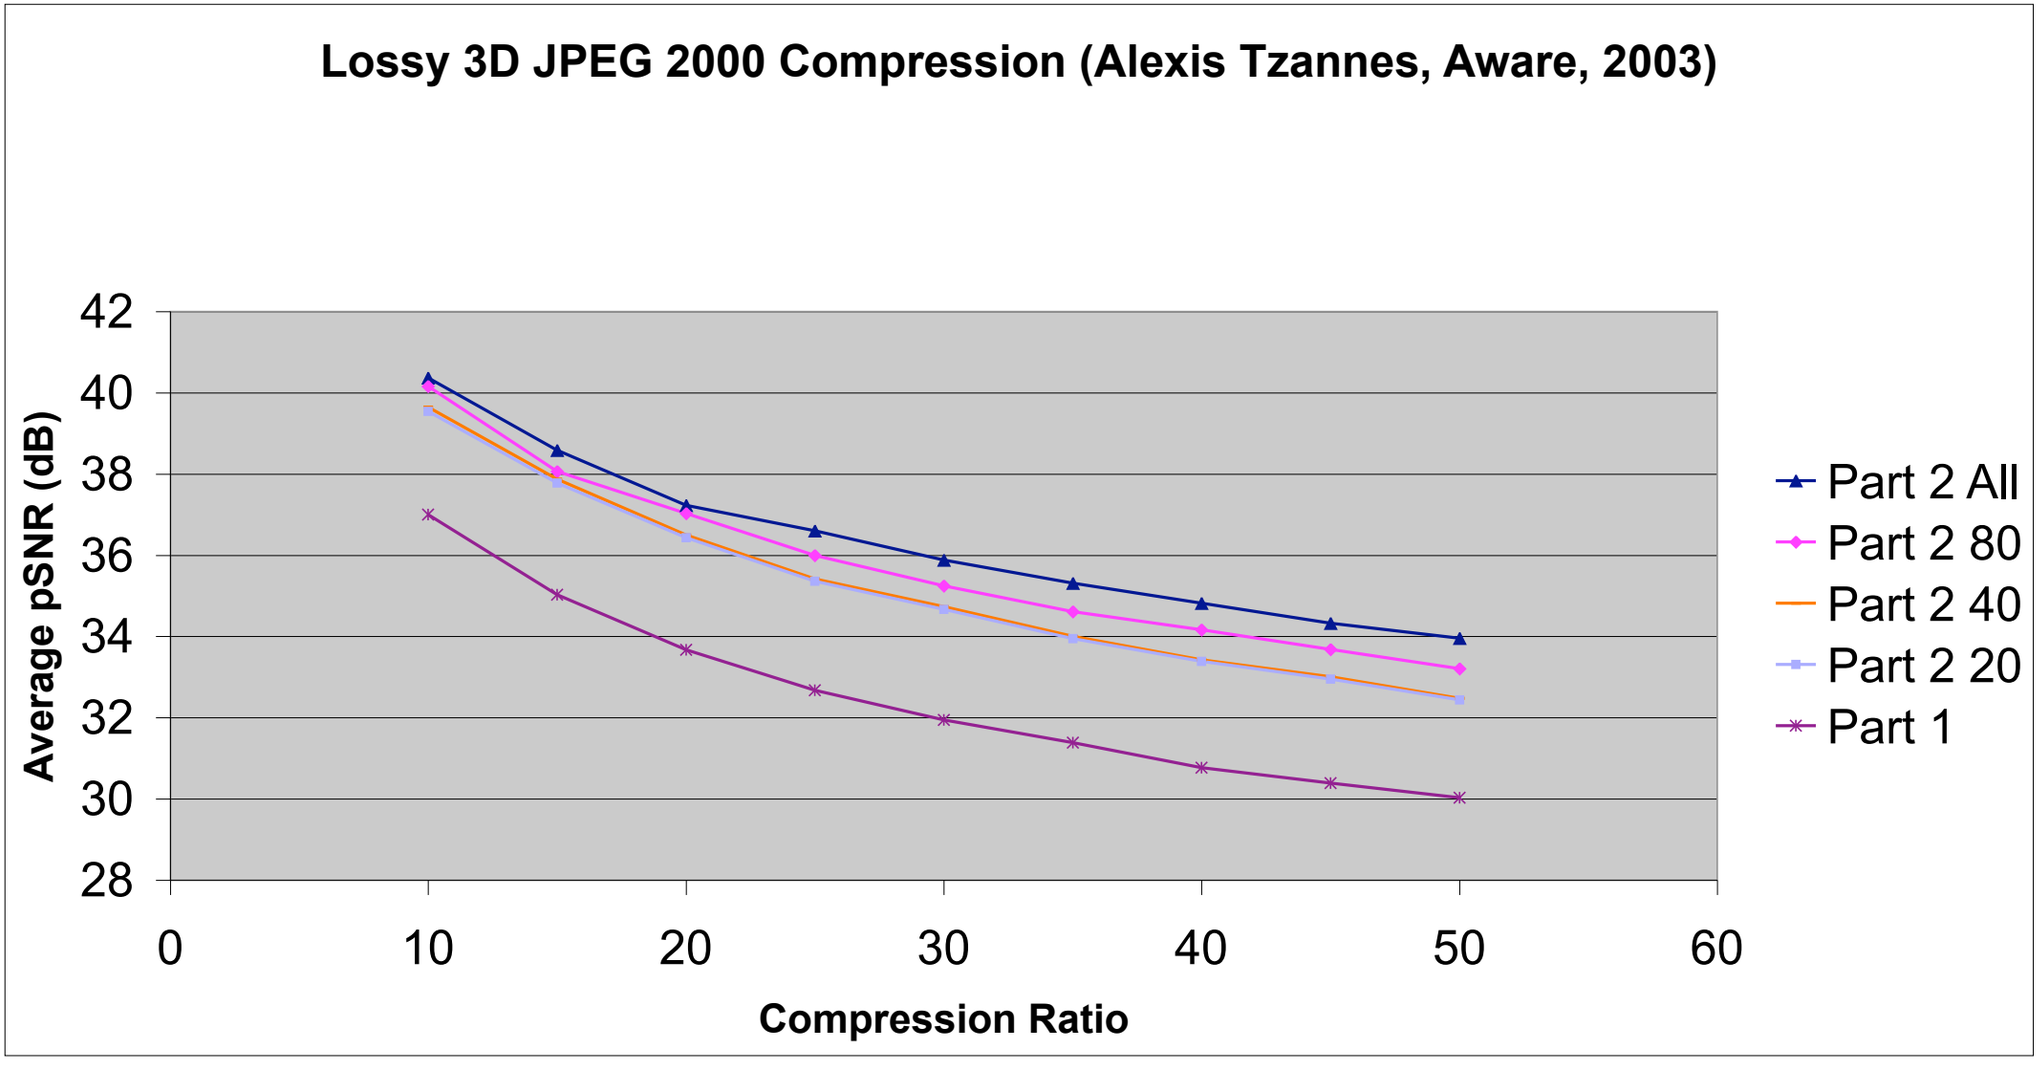
Lossy 3D JPEG 2000 Compression (Alexis Tzannes, Aware, 2003)
Average pSNR (dB)
42
40
38
36
34
32
30
28
0
10
20
30
40
50
60
Compression Ratio
Part 2 All
Part 2 80
Part 2 40
Part 2 20
Part 1

2D JPEG 2000 0.625mm slices

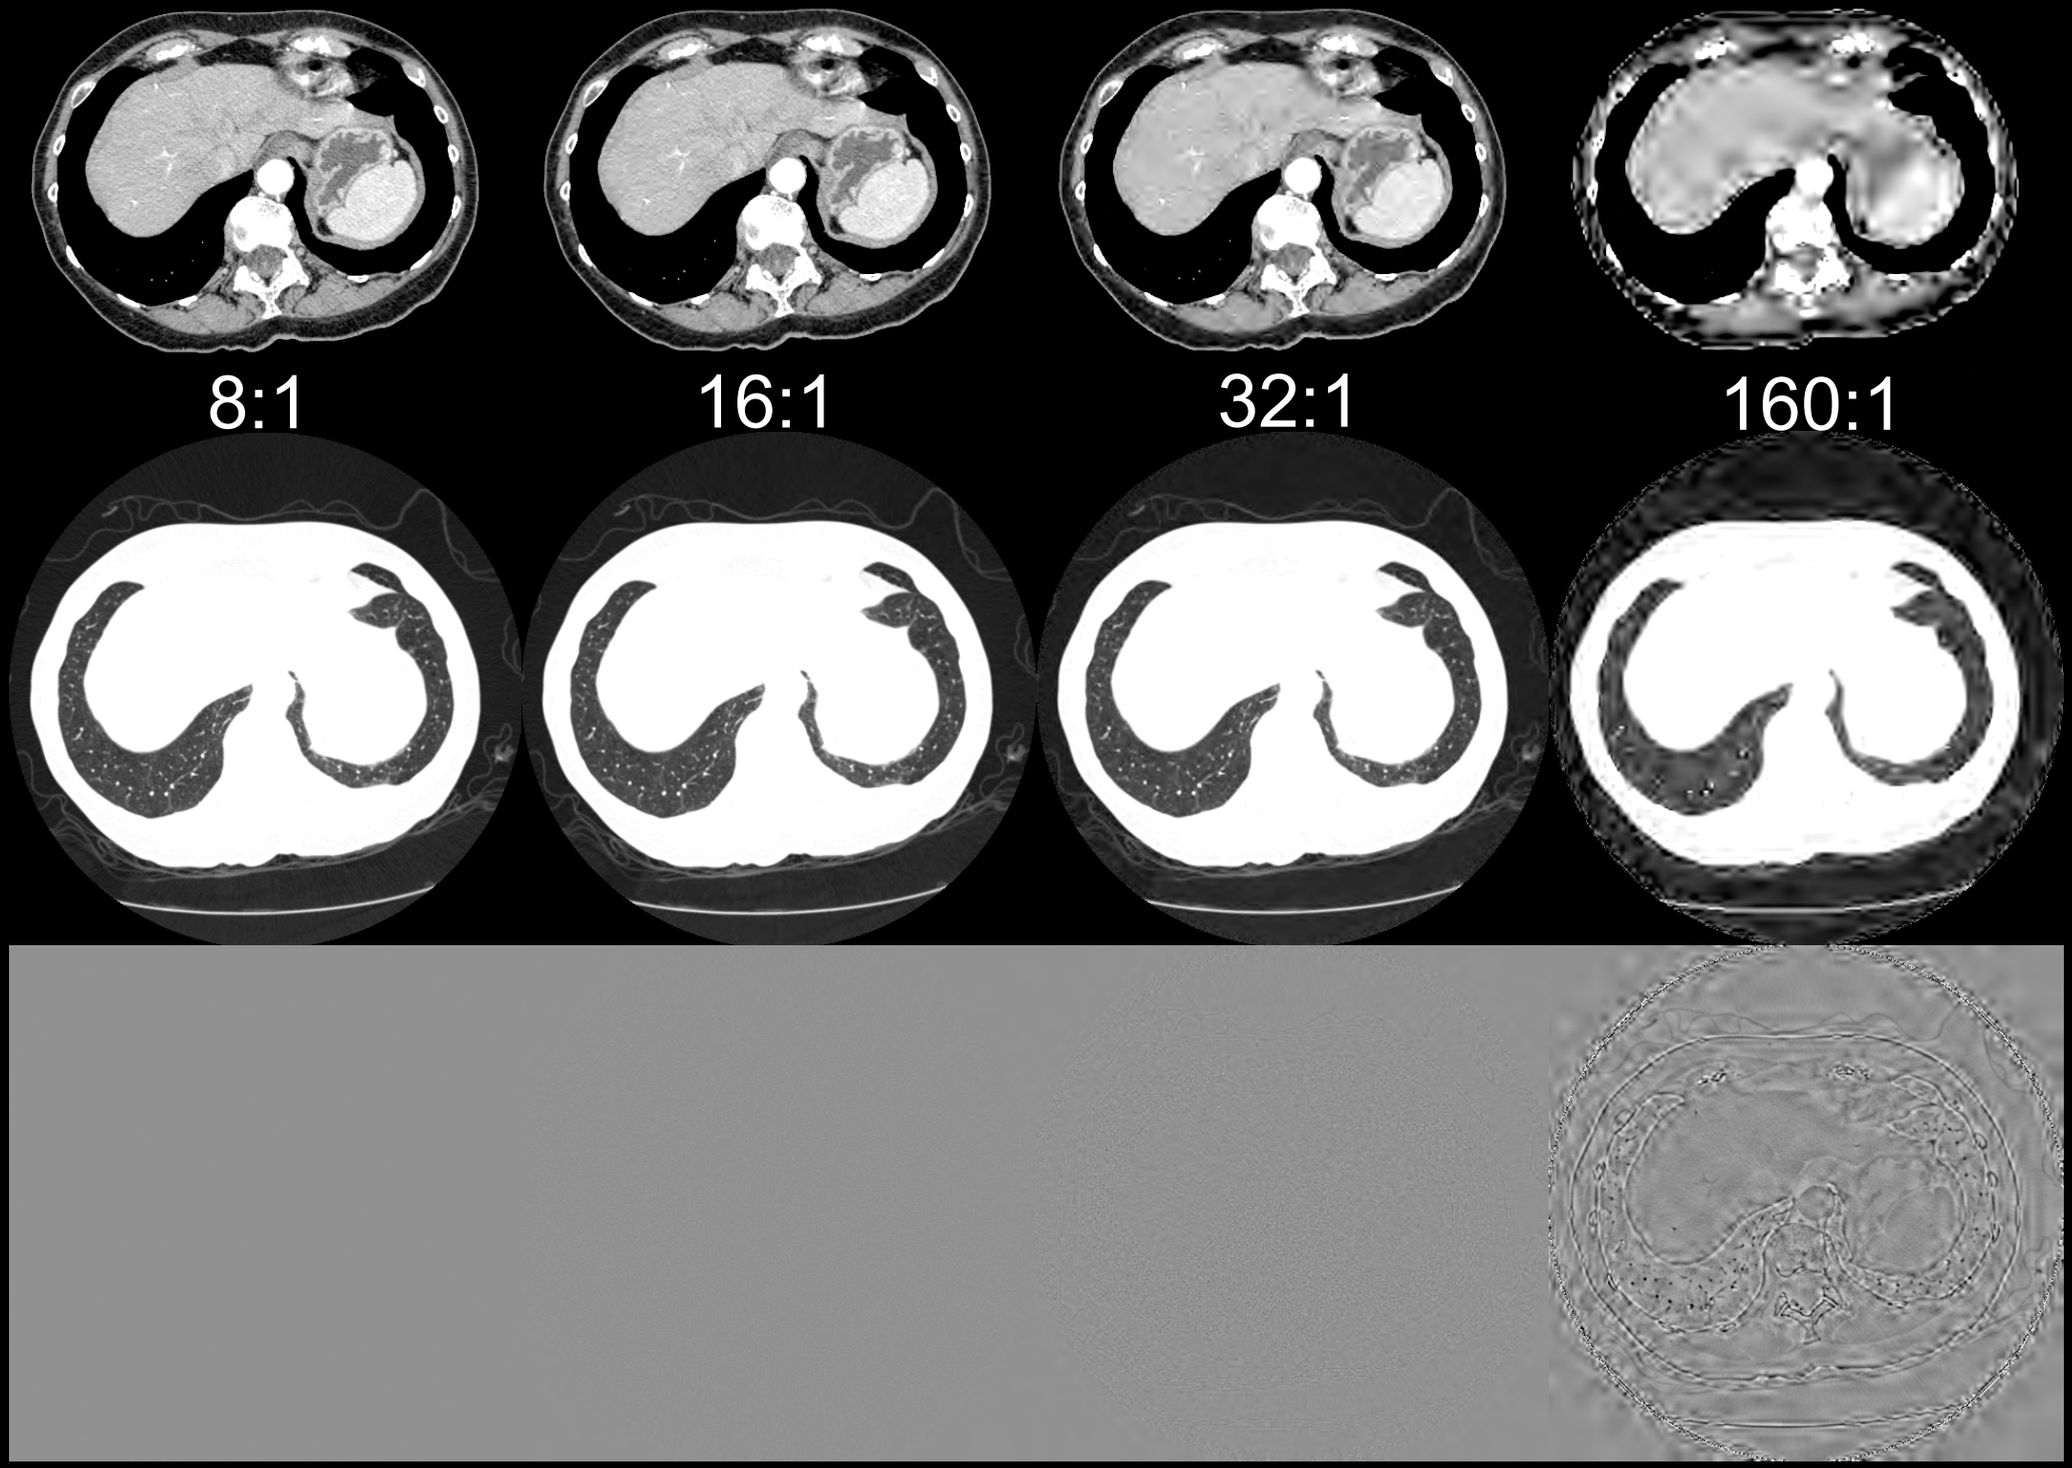


16 bpp
1:1


2 bpp
8:1
J2K


1 bpp
16:1
J2K


1 bpp
16:1
3D


16 bpp
1:1


1 bpp
16:1
J2K


1 bpp
16:1
JPEG


0.5 bpp
32:1
J2K

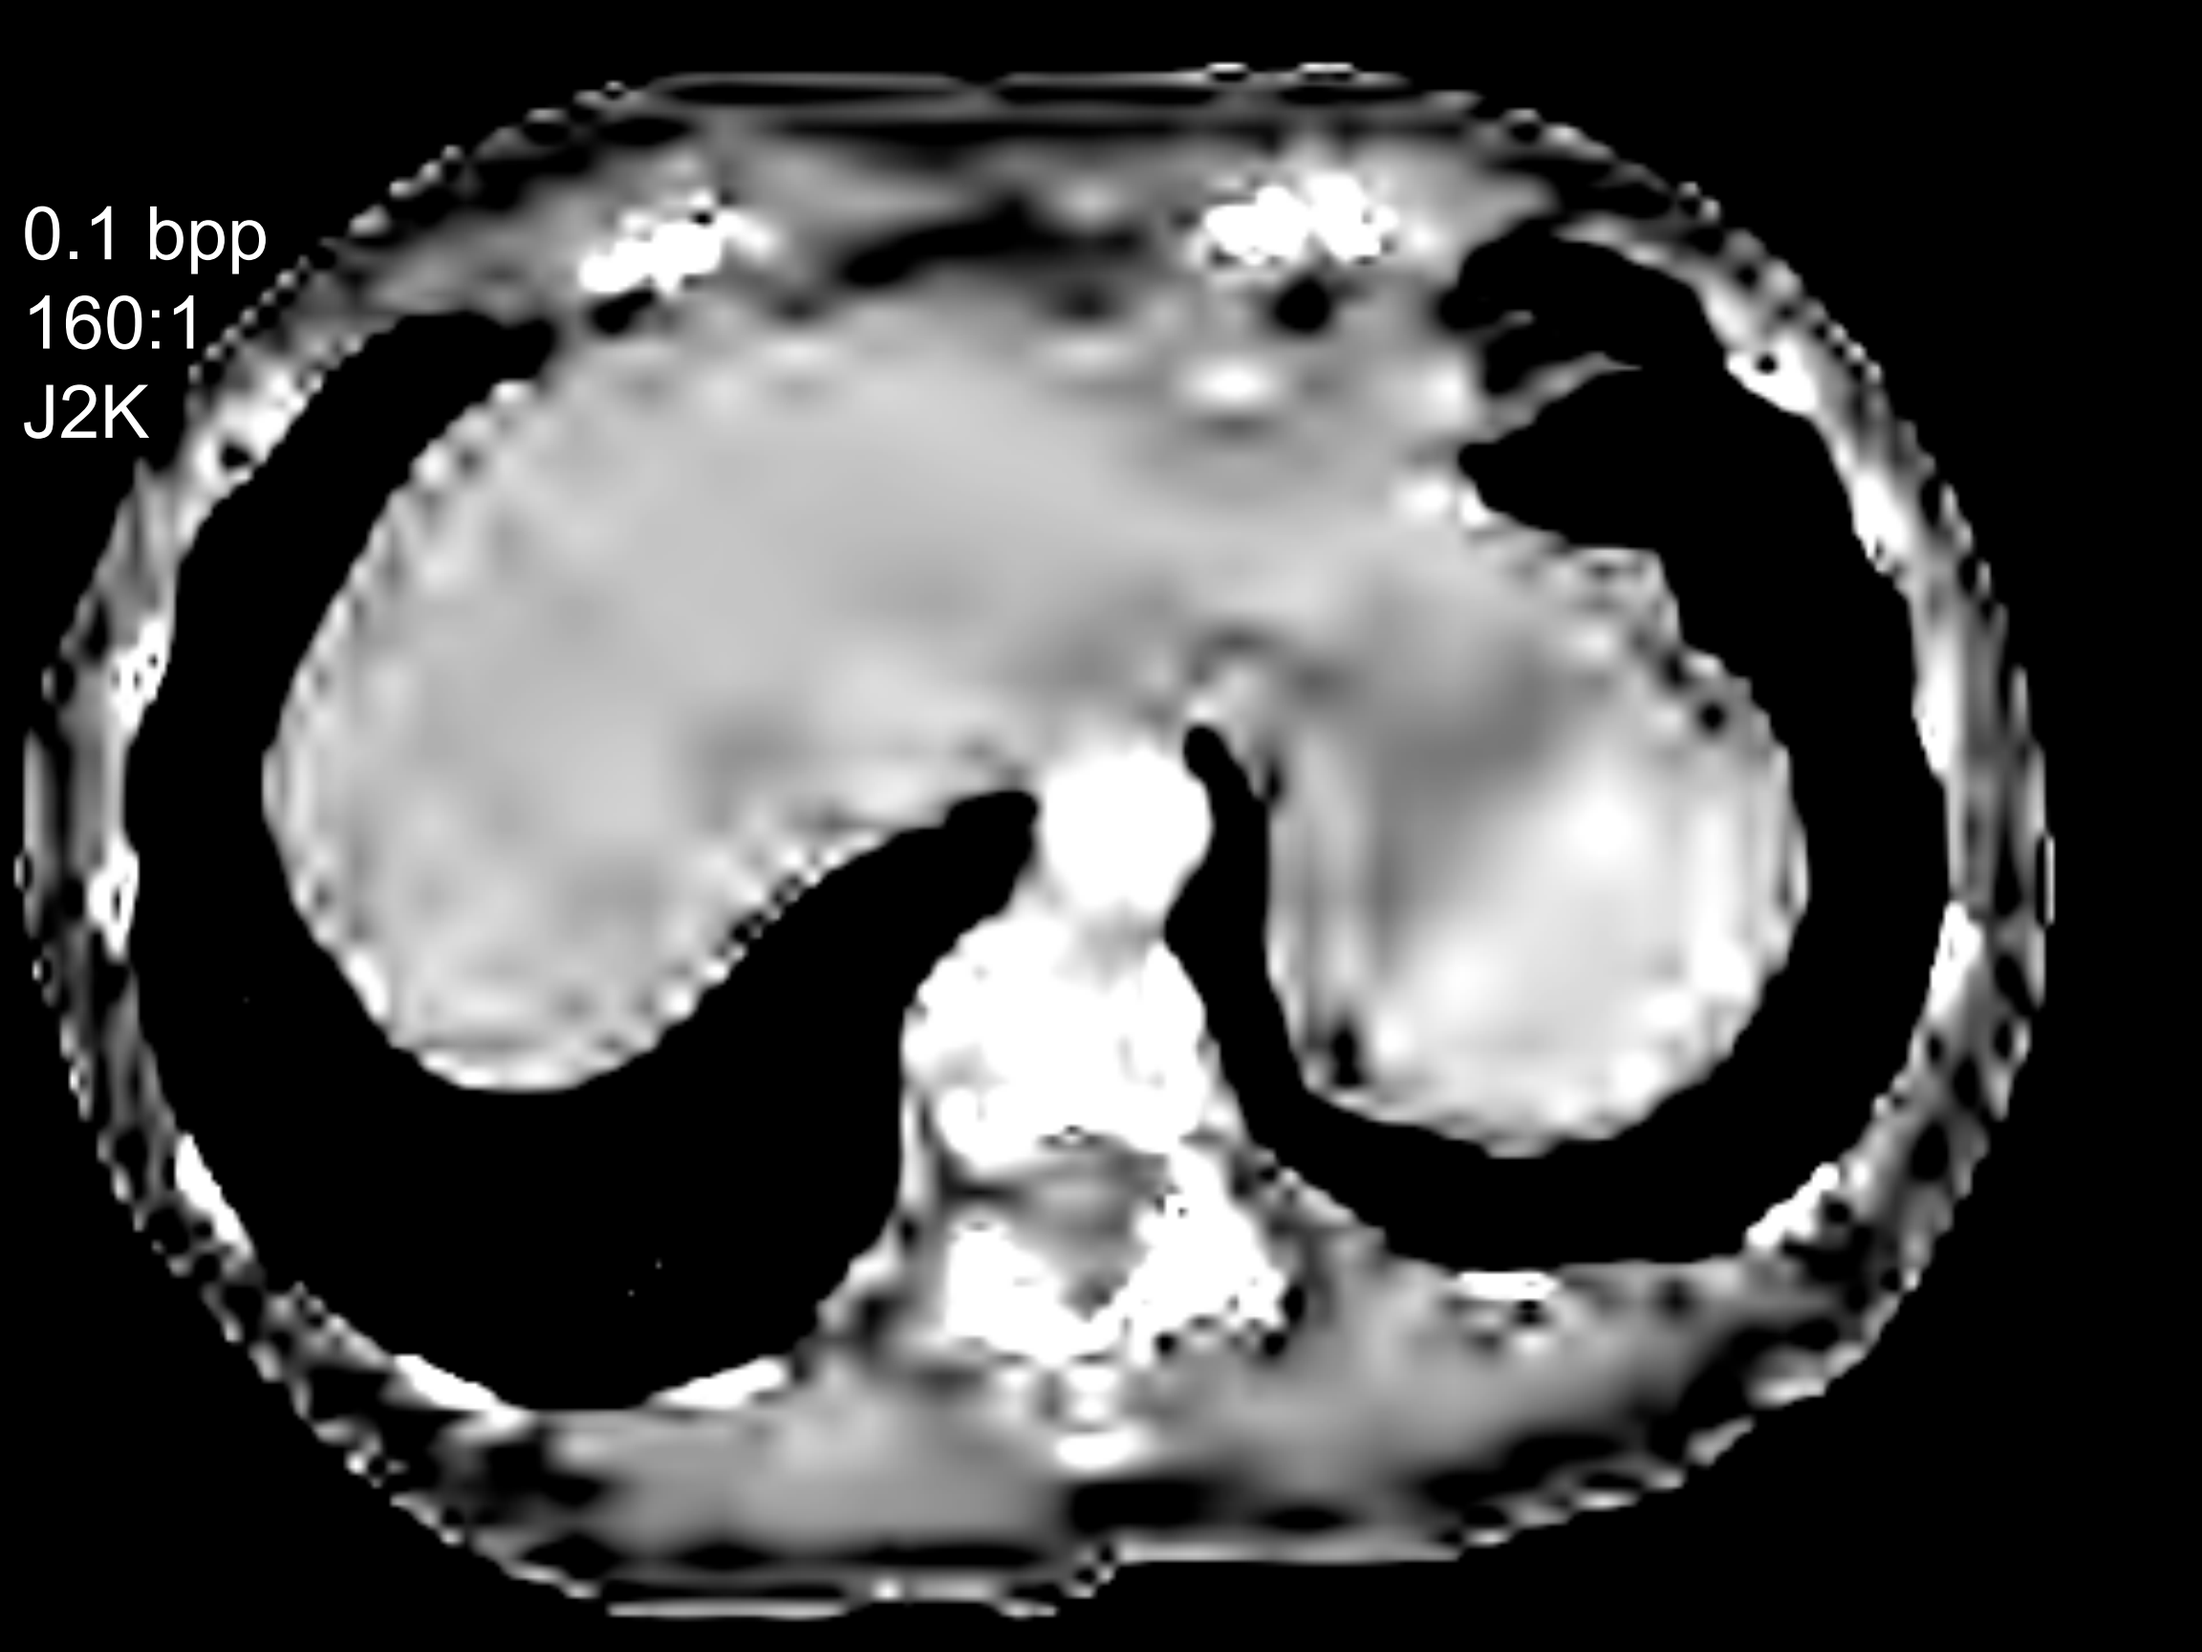
0.1 bpp
160:1
J2K

1 bpp
(16:1)

J2K 3D

JPEG

J2K

Multi-frame compression performance reality check

- Lossless compression in 3D
  - slight gain - 15 to 20% smaller than 2D
- Lossy compression in 3D
  - modest gain - possibly 50% smaller than 2D
  - but - only relatively modest loss before noticeable
- Recent studies of JPEG 2000 on CT, 2D and 3D
  - perceived image quality & detectable difference
  - observer performance studies
  - Ringl et al – 12.5:1
- DICOM has transfer syntax support for 3D J2K
  - no media profile though

Defining volumes

- What to compress in 3D ?
  - entire “volume” ?
  - sub-sets of adjacent contiguous slices ?
- How do you find a “volume” ?
  - in a bunch of separate single frame images ?
- What is a “volume” anyway ?
  - one traversal through space
- What about other dimensions ?
  - time (e.g. contrast phase), cardiac cycle, diffusion B value, etc. ?
- Not so easy to define a compressible volume !

DICOM volumes

- Existing DICOM CT and MR objects in common use are single frame
  - *CANNOT be used to transmit a 3D compressed volume !*
- New “enhanced” objects are multi-frame
  - Can be used to transmit or store a 3D compressed volume
  - Presupposing frames are ordered “appropriately” (e.g., sorted by spatial location)

3D versus “multi-component”

- JPEG 2000 multi-component transform
  - is not really “3D” per se
  - is simply “another” dimension in which a wavelet transform can be applied
- ITU-T Rec.T.800 | ISO/IEC 15444-1 Annex J
  - “The most common multiple component transformation application is the compression of colour images … are transformed into a colour space that is more conducive to spatial compression … technique can be extended for images that have more components; for example, LANDSAT images have seven components, six of which are highly correlated … can be used for the compression of CMYK images, multiple component medical images, and any other multiple component data.”

Multi-component types

- Anything correlated between frames
- Spatial dimension
  - a single 3D volume
- Time dimension
  - contrast perfusion study
  - cardiac gated (prospectively or retrospectively)
- Other dimensions
  - diffusion B value
  - functional MR paradigm

True “3D” J2K

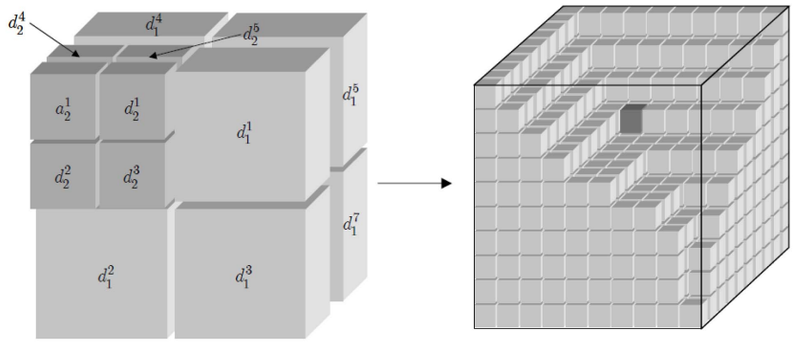

- Part 2 Annex J MCT is not the final word
- JPEG 2000 Part 10
  - “extensions for three-dimensional data”
  - for “logically rectangular 3-dimensional data sets with no time component”
  - extends MCT to support 3D “context models”
  - goal is “moderate” improvement
- Status
  - ISO standard as of 2008/12
  - Not yet in DICOM as Transfer Syntax
- Informal results - may be additional 5% improvement

Bruylants et, SPIE News

Multi-frame versus video

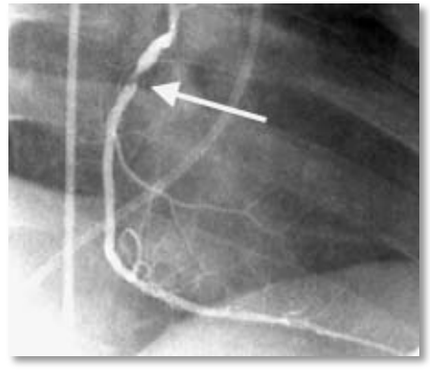

- First use of JPEG in DICOM was cardiac XA CD
  - *lossless* 8 bit multi-frame
- Other use is ultrasound
  - especially echocardiography
  - *lossy* 8 bit per channel RGB
  - precursor was VHS video (limited quality requirement)
- No inter-frame prediction - distinguish
  - multi-frame from cine
  - cine from “video”
- Visible light video – e.g., endoscopy
  - MPEG-2 MP@ML, MP@HL (HD) in DICOM
  - Work in progress – H.264 (MPEG-10)

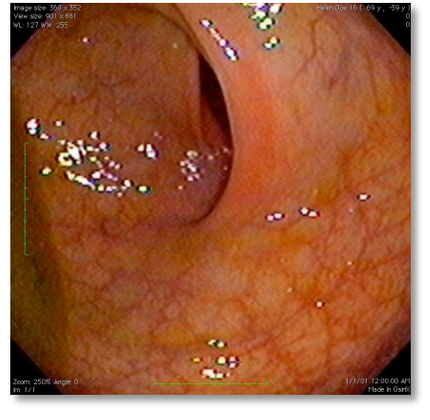

Reality Check

- What is currently mostly used ?
  - Lossless JPEG CT and MR MODs
  - Lossless JPEG XA CDs and network
  - Lossy JPEG Ultrasound media and network
  - Lossless RLE Ultrasound (occasionally)
  - Lossless JPEG in long term archives
- Growth areas
  - MPEG-2 for visible light video
  - Lossy and progressive J2K for remote access
  - Lossless J2K on CD and DVD for all modalities
  - Lossless or lossy J2K in archives

What about JPEG-XR ?

- Microsoft initiative
  - began as Windows Media Photo
  - renamed HD Photo
- Better than JPEG, simpler than J2K
  - support for more data types (32 bit, floats)
  - DCT-like transform but lossless
  - loss is confined to quantization step
- Proposed to ISO/IEC 29199-2 / ITU T.832
- Work in progress – DIS stage 2009/01
- Friendly IP - available to open source initiatives
- When/why/whether to add DICOM Transfer Syntax ?

Further in the future

- New lossless schemes ?
  - unlikely – may have reached practical limit
  - proprietary claims unsubstantiated
- Visually lossless threshold metric
  - adaptive (per image schemes)
  - discussed by Kim et al (JPEG WG1N4996)
  - encoder not decoder, so no DICOM impact
- JPEG's AIC (Advanced Image Coding)
  - just JPEG XR, or more than that ?

Compressed Transfer Syntax

- “Header” (non-pixel data) same as uncompressed
- List of sorted data elements
- Other large data elements not compressed
- Only Pixel Data (7FE0,0010) is different
  - uncompressed – 8 or 16 bit words
  - compressed – encapsulated “bit stream”
- Signaled by undefined Value Length (0xffffffff)
- Always Explicit Value Representation

Encapsulation Mechanism

- Still-frame schemes (JPEG, J2K)
  - conceptually compress one frame at a time
  - encode compressed frame in “fragments”
  - one *or more* fragments per frame
  - fragments of fixed size allow buffering
  - frames may not span fragments
  - last fragment padded to even length
  - optional “table of contents” prior to fragments
  - encoded as Sequence Items within (7FE0,0010)

Table A.4-2
EXAMPLES OF ELEMENTS FOR AN ENCODED TWO-FRAME IMAGE DEFINED AS A SEQUENCE OF THREE FRAGMENTS WITH BASIC TABLE ITEM VALUES

Pixel Data Element Tag	**Value Representation**		**Data Element Length**	**Data Element**					
				Basic Offset Table with Item Value			First Fragment (Frame 1) of Pixel Data		
				Item Tag	Item Length	Item Value	Item Tag	Item Length	Item Value
(7FE0, 0010) with VR of OB	OB	0000H Reserved	FFFF FFFFH undefined length	(FFFE, E000)	0000 0008H	0000 0000H 0000 0646H	(FFFE, E000)	0000 02C8H	Compressed Fragment
4 bytes	2 bytes	2 bytes	4 bytes	4 bytes	4 bytes	0008H bytes	4 bytes	4 bytes	02C8H bytes

Data Element Continued							
Second Fragment (Frame 1) of Pixel Data			Third Fragment (Frame 2) of Pixel Data			Sequence Delimiter Item	
Item Tag	Item Length	Item Value	Item Tag	Item Length	Item Value	Sequence Delimiter Tag	Item Length
(FFFE, E000)	0000 036EH	Compressed Fragment	(FFFE, E000)	0000 0BC8H	Compressed Fragment	(FFFE, E0DD)	0000 0000H
4 bytes	4 bytes	036EH bytes	4 bytes	4 bytes	0BC8H bytes	4 bytes	4 bytes

Encapsulation Mechanism

- Multi-frame schemes (MPEG2)
  - compress entire bit stream
  - codec handles frames
  - codec uses inter-frame motion prediction
  - everything in just one fragment
  - size limit on fragment – separate instances
  - no “table of contents” (codec may include)

Compression “history”

- Transfer Syntax -> lossy compressed or not
- What if decompressed then stored or transmitted ?
- Lossy Image Compression flag (true or false)
- Derivation Description
  - lossy compressed image is “derived”
  - unless it came off the “sensor” lossy compressed
  - needs a new SOP Instance UID
- Lossy Image Compression Method
- Lossy Image Compression Ratio
  - relative to what - bits stored or allocated ?
- Attributes optional in old image objects, mandatory in new
- Mandatory from a regulatory perspective

Compression Ratio confusion

- Ratio of what relative to what ?
- Number of bits on disk (16), or
- Number of meaningful bits (e.g., 12)

to

- Number of compressed “bits per pixel”

- E.g. 1 bpp - express as 16:1 or 12:1 ?

Various Encoding Issues

- Signed or unsigned
  - JPEG assumes signed; JPEG-LS, J2K have explicit support
- Rescale slope and intercept
  - may need to “fix” during compression if pixel range changed
  - especially with 12 bit JPEG
- Pixel padding value & range
  - detect and replace to fit bit depth
  - “bleed through” if changed by lossy compression
- LUT index values cannot be lossy compressed
  - both color palette and VOI LUT
  - may get away with it if monotonically increasing
- Overlays in pixel data high bits
  - remove and put in Overlay Data (large & not compressed)

Compression on Network

- A network connection is established (TCP/IP)
- An Association is established
- Involves proposal and acceptance (negotiation) of

Abstract Syntax == SOP Class (“type” of object)

combined with

Transfer Syntax (uncompressed, type of compression)

Association Negotiation

- Must always propose uncompressed
  - unless only lossy form available
  - maximizes chance of exchange
- Propose one SOP Class, multiple Transfer Syntaxes
  - SCP (receiver) gets to choose which to use
- Propose separate SOP Class+Transfer Syntax pairs
  - SCU (sender) discovers what SCP supports
  - SCU gets to choose which to use
- Both
  - SCU discovers what SCP supports and what it prefers

Compression on Media

- No interactive negotiation
- Choice has to be defined a priori
- Recipient must support all choices
- Application + image type + compression
- Defined in *Media Application Profile*
- E.g., cardiac XA + 8 bit gray + lossless
- General purpose CD profile – no compression
- General purpose DVD – JPEG *or* J2K (separate)
- IHE PDI DVD – both JPEG *and* J2K

Private Transfer Syntaxes

- Suppose you have a “better” codec ?
- Not in DICOM
- Can negotiate on network
  - if both ends support it – use it
  - if recipient doesn’t support it – use standard
- Allows for progress and proprietary “added value”
  - without sacrificing DICOM object and protocol
- Cannot be used on media

Beyond Store-and-Forward

- DICOM network transfer is like email
  - request what to send
  - decide what to send
  - send it entirely
- No streaming
- No selective retrieval
  - of sub-regions
  - of contrast range
  - of quality (lossy)
- No “interactivity”

Adding Interactivity - JPIP

- JPIP – JPEG Interactive Protocol
- ISO/IEC 15444-9 / ITU T.808
- Part of JPEG 2000 family of standards
- A means of requesting
  - what images and frame and components
  - what sub-regions
  - what quality

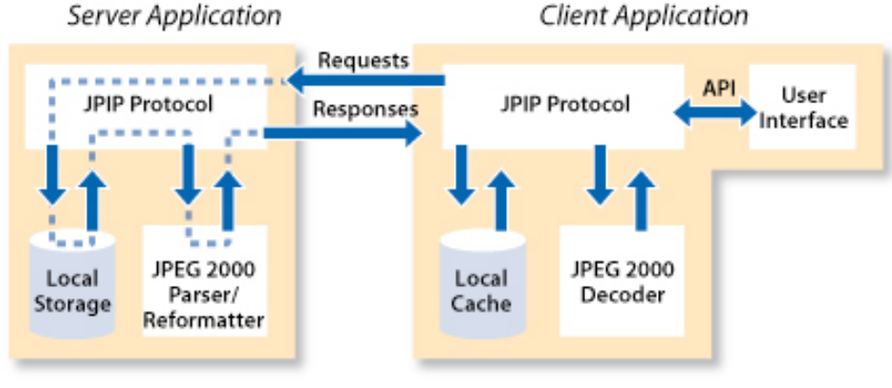

http://www.aware.com/

DICOM and JPIP

- How to add an interactive capability ?
- Without replicating functions of JPIP ?
- How to re-use JPIP implementations ?
- New attribute
  - Pixel Data Provider URL (0028,7FE0)
  - *instead of* Pixel Data (7FE0,0010)
  - can use with *any* image storage SOP Class
  - only permitted for *new Transfer Syntax*
  - not on media (obviously)

JPIP in practice

- Not the panacea that was hoped for
- Organizing the JPEG 2000 image for efficient JPIP retrieval slow and complex
- JPIP server does not guarantee to return what was requested
  - targeted for consumer web with server load problems
  - cannot guarantee requested quality
- DICOM has image size limitations
  - Row and Columns are 16 bit values (2^16-1)
  - impacts pathology Whole Slide Imaging (WSI) application

Adding Interactivity - WADO

- Web Access to DICOM Objects
- DICOM PS 3.18 / ISO 17432
- Uses HTTP GET with URL parameters
  - single transaction
  - select object (by UID)
  - select return format (DICOM, JPEG, etc.)
  - select sub-region
- Primary use is “web-front end to PACS”
  - e.g., to supply consumer images to EMR
- Supported in IHE XDS-I
- Web Service based version under development

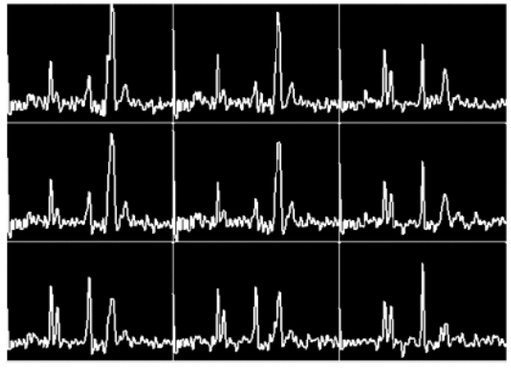

Not an image ?

- Large header attributes in images
  - lookup tables, overlays, bitmap shutter
  - not compressed
  - e.g., mammo 16:1 with overlay – will be as large as compressed pixel data !
  - so don't use overlays !
- Waveforms (ECGs), MR spectroscopy, SR
  - Deflate Transfer Syntax (same as zip)
  - compresses entire object
- Audio
  - within MPEG video stream
  - as a separate waveform – Deflate Transfer Syntax

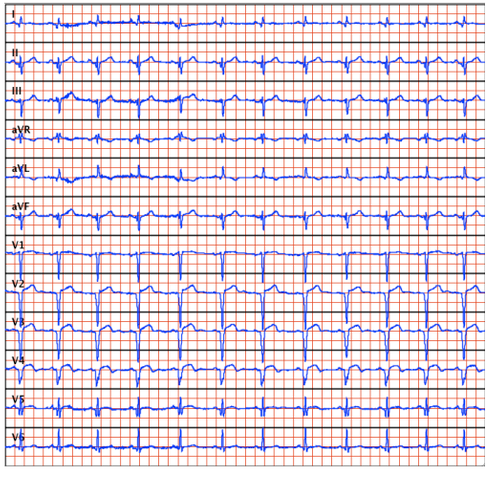

Summary

- DICOM has a lot of compression support
- State of the art lossless & lossy schemes
- Tracks ISO/IEC/ITU standards
- Supports still frame, volume and video
- Doing it right requires attention to detail
- Network, media, web and interactive
- Non-image object support