Lossy Compression of Natural HDR Content based on Multi-Component Transform Optimization

Miguel Hernández[⋆], Victor Sanchez[⋆],
Francesc Aulí-Llinàs[†] and Joan Serra-Sagristà[†]

[⋆] WARWICK THE UNIVERSITY OF WARWICK

[†] UAB Universitat Autònoma de Barcelona

EPSRC Engineering and Physical Sciences Research Council

IEEE DMIAF 2016
MA1.5 :: July 4, 2016

Compression

Uncompressed size of a 90 min movie:

Compression

Uncompressed size of a 90 min movie:

- HD: 1920×1080 @ 30 fps, 10-bit: 1.2 TB (222 MB/s)

Compression

Uncompressed size of a 90 min movie:

- HD: 1920×1080 @ 30 fps, 10-bit: 1.2 TB (222 MB/s)
- 4K: 3840×2160 @ 50 fps, 10-bit: 7.6 TB (1483 MB/s)

Compression

Uncompressed size of a 90 min movie:

- HD: 1920×1080 @ 30 fps, 10-bit: 1.2 TB (222 MB/s)
- 4K: 3840×2160 @ 50 fps, 10-bit: 7.6 TB (1483 MB/s)
- 8K: 7680×4320 @ 50 fps, **12**-bit: 36.7 TB (7119 MB/s)

Compression

Uncompressed size of a 90 min movie:

- HD: 1920×1080 @ 30 fps, 10-bit: 1.2 TB (222 MB/s)
- 4K: 3840×2160 @ 50 fps, 10-bit: 7.6 TB (1483 MB/s)
- 8K: 7680×4320 @ 50 fps, **12**-bit: 36.7 TB (7119 MB/s)

Unrealistic values

Compression

Uncompressed size of a 90 min movie:

- HD: 1920×1080 @ 30 fps, 10-bit: 1.2 TB (222 MB/s)
- 4K: 3840×2160 @ 50 fps, 10-bit: 7.6 TB (1483 MB/s)
- 8K: 7680×4320 @ 50 fps, **12**-bit: 36.7 TB (7119 MB/s)

Unrealistic values... because of compression!

Compression

Uncompressed size of a 90 min movie:

- HD: 1920×1080 @ 30 fps, 10-bit: 1.2 TB (222 MB/s)
- 4K: 3840×2160 @ 50 fps, 10-bit: 7.6 TB (1483 MB/s)
- 8K: 7680×4320 @ 50 fps, **12**-bit: 36.7 TB (7119 MB/s)

Unrealistic values... because of compression!

↑bit-depth (HDR) ⇒ ↑need for compression

MCT Optimization

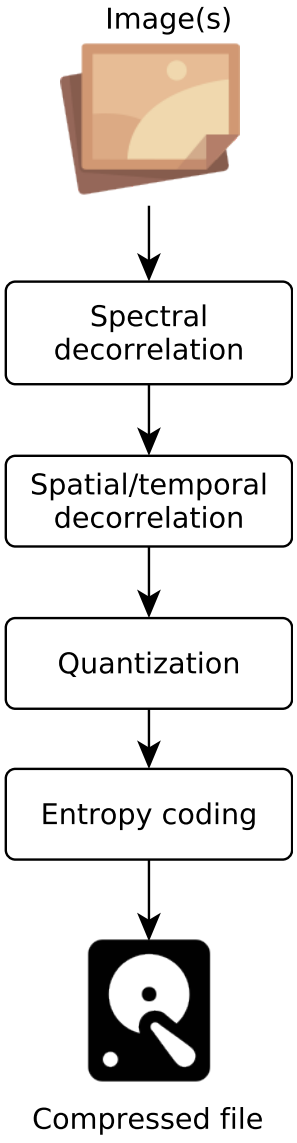

MCT Optimization

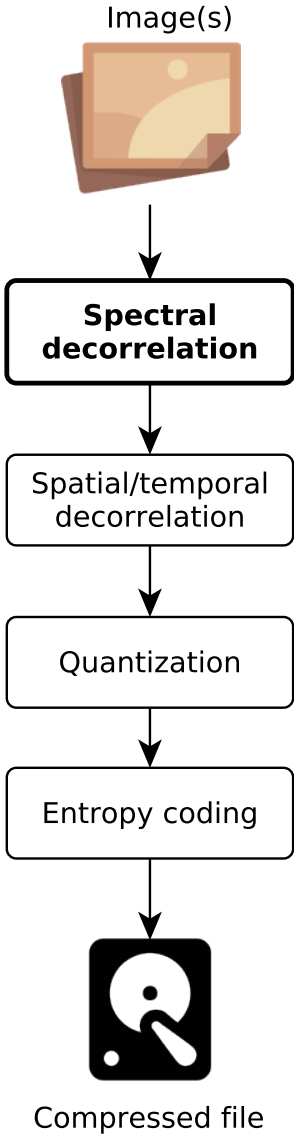

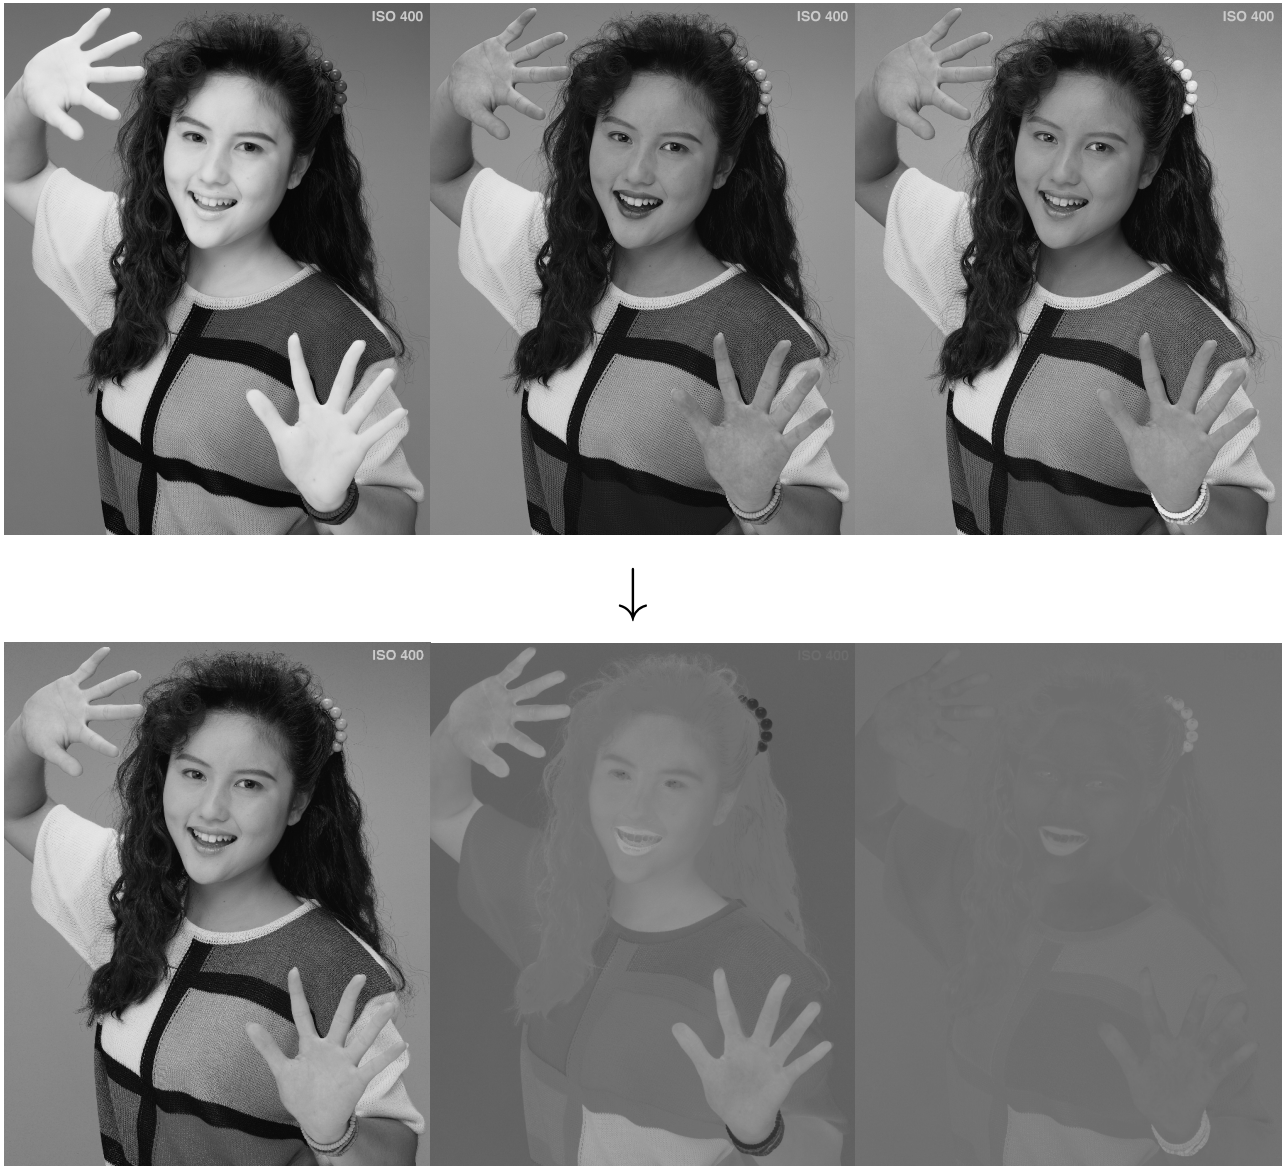

MCT Optimization

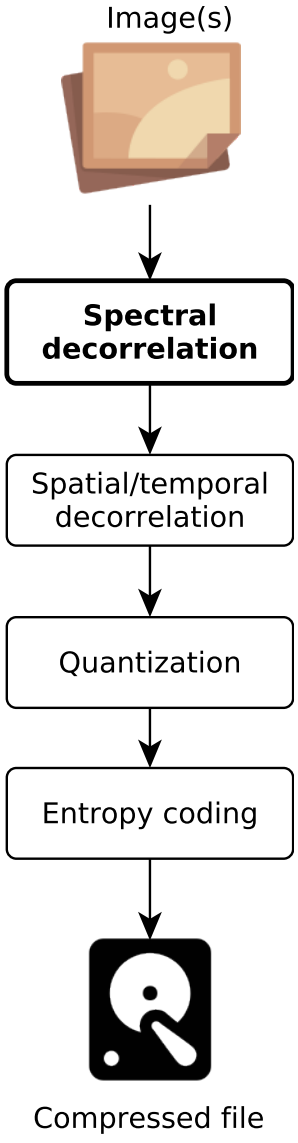

Linear **Multi-Component Transforms** (MCTs):

- decorrelate (color) components
- compact energy
- low complexity

$\Rightarrow$ improve coding performance

MCT Optimization

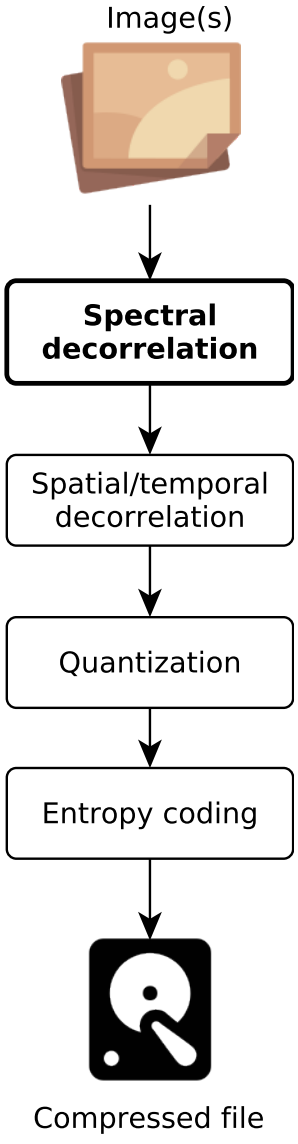

Linear **Multi-Component Transforms** (MCTs):

- decorrelate (color) components
- compact energy
- low complexity

$\Rightarrow$ improve coding performance

JPEG2000 + MCT Optimization: good on LDR

- Natural LDR: +4% PSNR
- Microscopy LDR: +6% PSNR

MCT Optimization

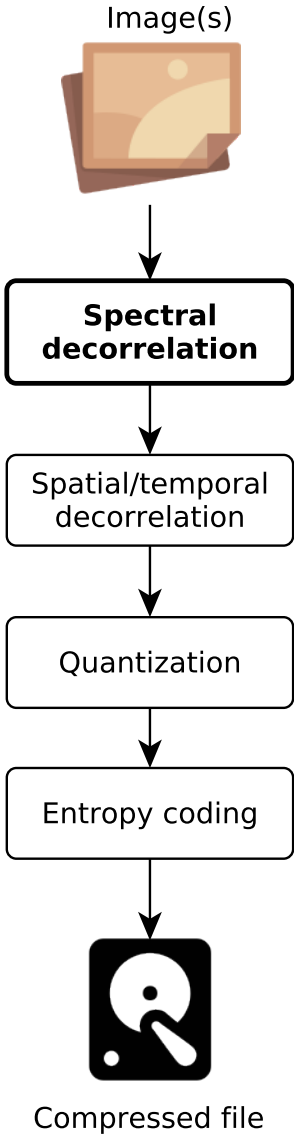

Linear **Multi-Component Transforms** (MCTs):

- decorrelate (color) components
- compact energy
- low complexity

$\Rightarrow$ improve coding performance

JPEG2000 + MCT Optimization: good on LDR

- Natural LDR: +4% PSNR
- Microscopy LDR: +6% PSNR

MCT Optimization for HDR contents?

Contents

Contents

Linear MCTs

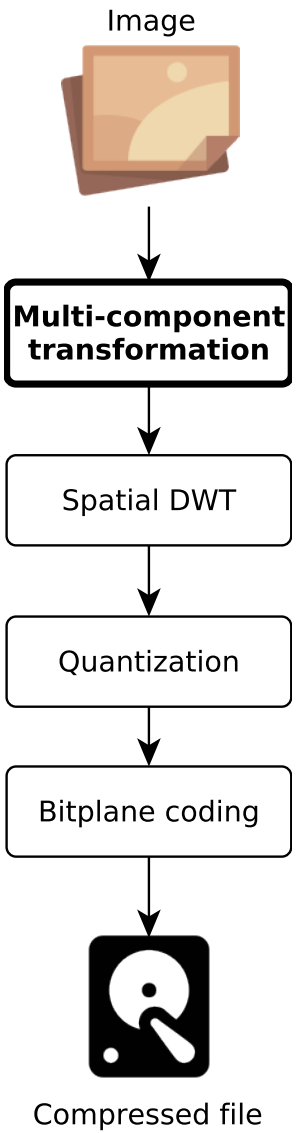

Focus on JPEG2000

Linear MCTs

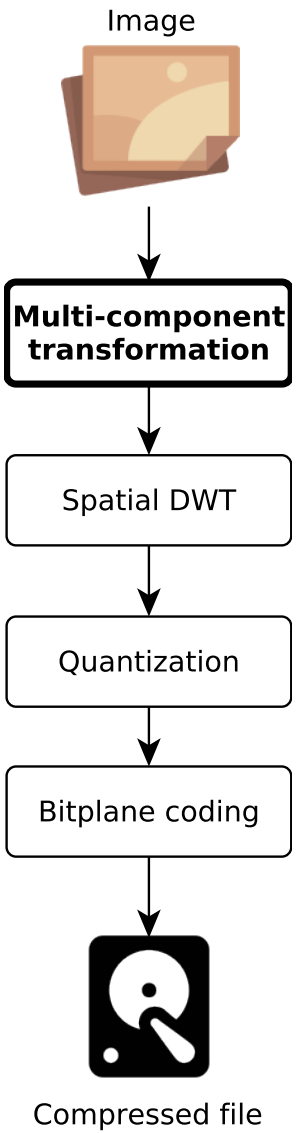

Focus on JPEG2000

Focus on **linear** MCTs:

$$\begin{pmatrix} u_i \\ v_i \\ w_i \end{pmatrix} = \mathbf{M} \begin{pmatrix} r_i \\ g_i \\ b_i \end{pmatrix} \quad \forall \text{ pixel } (r_i, g_i, b_i),$$

Linear MCTs

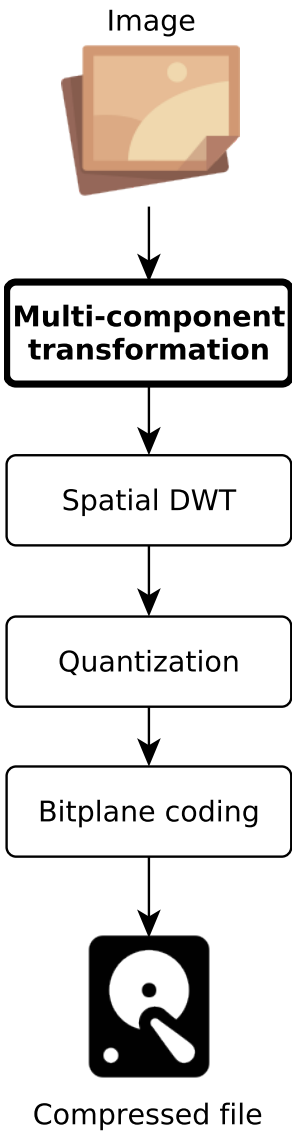

Focus on JPEG2000

Focus on **linear** MCTs:

$$\begin{pmatrix} u_i \\ v_i \\ w_i \end{pmatrix} = \mathbf{M} \begin{pmatrix} r_i \\ g_i \\ b_i \end{pmatrix} \quad \forall \text{ pixel } (r_i, g_i, b_i),$$

where

$$\mathbf{M} = \begin{pmatrix} m_{1,1} & m_{1,2} & m_{1,3} \\ m_{2,1} & m_{2,2} & m_{2,3} \\ m_{3,1} & m_{3,2} & m_{3,3} \end{pmatrix} \in M_{3\times 3}(\mathbb{R})$$

Existing Approaches

Approaches to MCTs:

- Static MCTs
- Data-dependent MCTs

Existing Approaches

Approaches to MCTs:

- Static MCTs – designed for LDR
- Data-dependent MCTs

Existing Approaches

Approaches to MCTs:

- Static MCTs – designed for LDR
- Data-dependent MCTs – neglect compressor details

Existing Approaches

Approaches to MCTs:

- Static MCTs – designed for LDR
- Data-dependent MCTs – neglect compressor details
- Our MCT optimization framework

Our approach

Don't minimize correlation, mutual information

Our approach

Don't ~~minimize correlation, mutual information~~

Do { minimize distortion of reconstructed images

Our approach

Don't ~~minimize correlation, mutual information~~

Do
- minimize distortion of reconstructed images
- use real compression algorithm

Our approach

Don't ~~minimize correlation, mutual information~~

Do
- minimize distortion of reconstructed images
- use real compression algorithm
- state as optimization problem

$$\underset{MCT}{\text{argmin}} \ \textsc{EvaluateMSE}(MCT)$$

Our approach

Don't ~~minimize correlation, mutual information~~

Do
- minimize distortion of reconstructed images
- use real compression algorithm
- state as optimization problem

$$\underset{MCT}{\operatorname{argmin}} \ \textsc{EvaluateMSE}(MCT)$$

Cost function

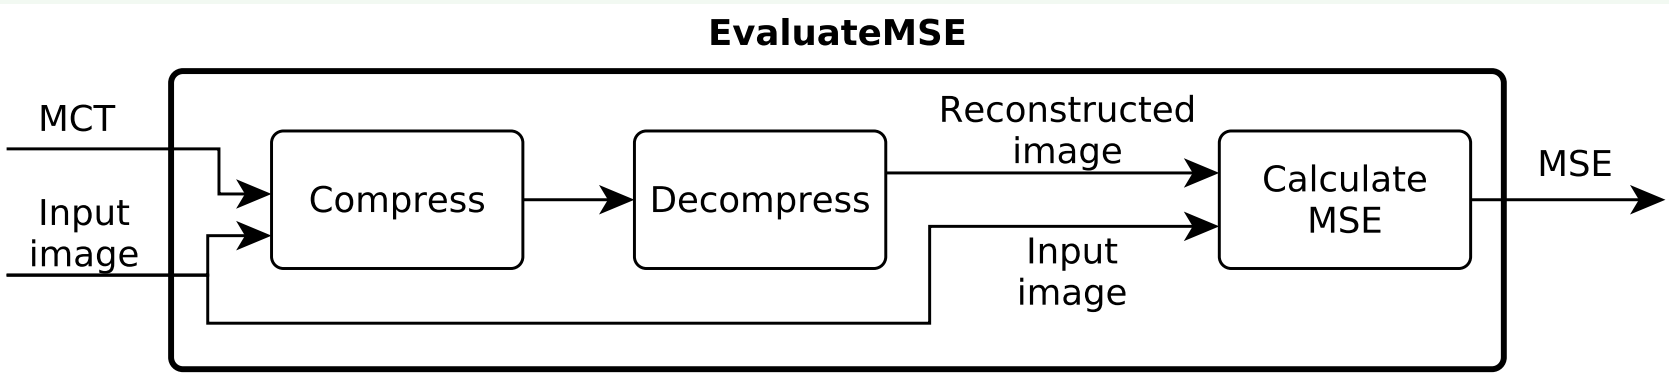

MCT optimization

Analytical EvaluateMSE: intractable ⇒ numerical optimization

MCT optimization

Analytical EvaluateMSE: intractable ⇒ numerical optimization

High-level description

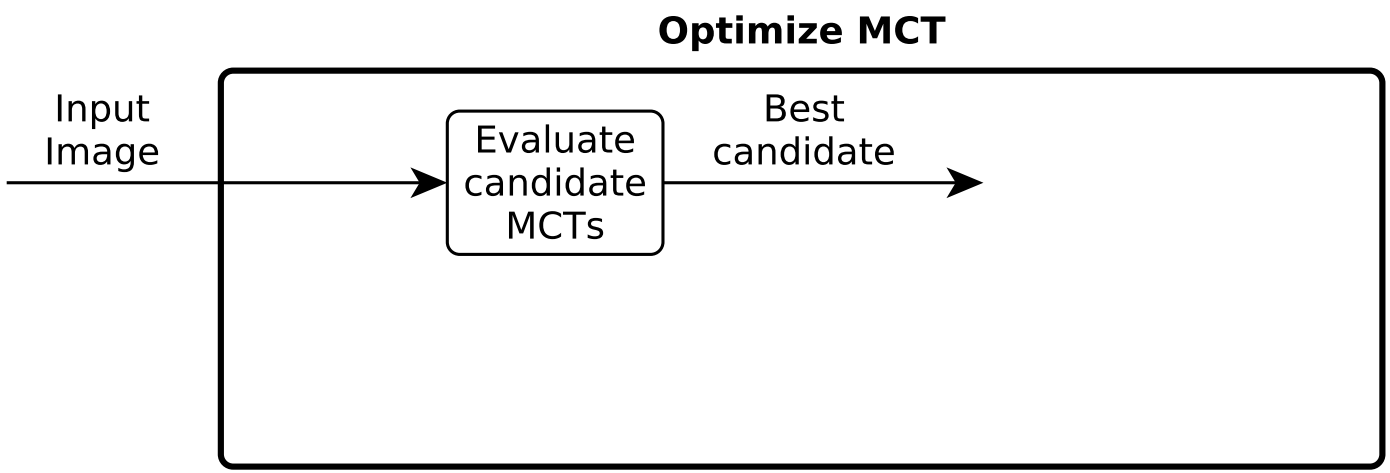

1. Find best candidate

MCT optimization

Analytical EvaluateMSE: intractable ⇒ numerical optimization

High-level description

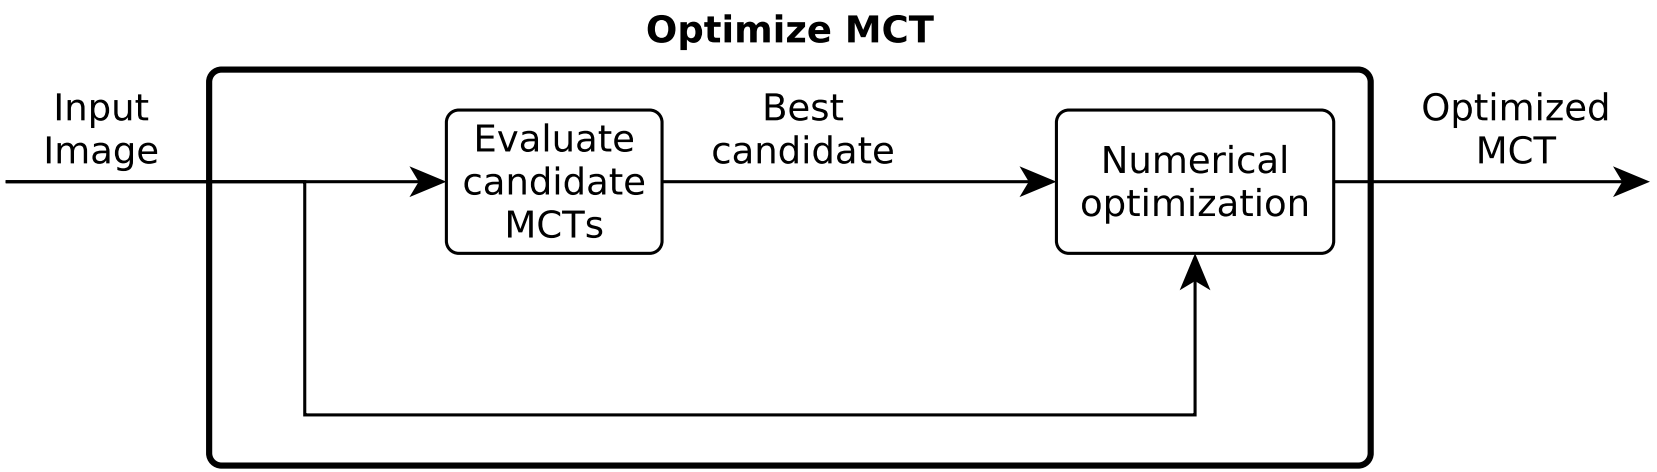

1. Find best candidate
2. Apply Powell's method

Contents

Original MCT Optimization designed for:

- 8-bit integer images
- 3 components (R,G,B)

Original MCT Optimization designed for:

- 8-bit integer images
- 3 components (R,G,B)

How to apply on HDR data?

Original MCT Optimization designed for:

- 8-bit integer images
- 3 components (R,G,B)

How to apply on HDR data?

- Color-graded HDR
- Scene radiance
- Bayer CFA mosaics

Color-graded HDR

- 0.005-4000 cd/m^2 Rec.2020
- RGB, 1920×1080
- 16-bit integers

Color-graded HDR

- 0.005-4000 cd/m^2 Rec.2020
- RGB, 1920×1080
- 16-bit integers

High-level description

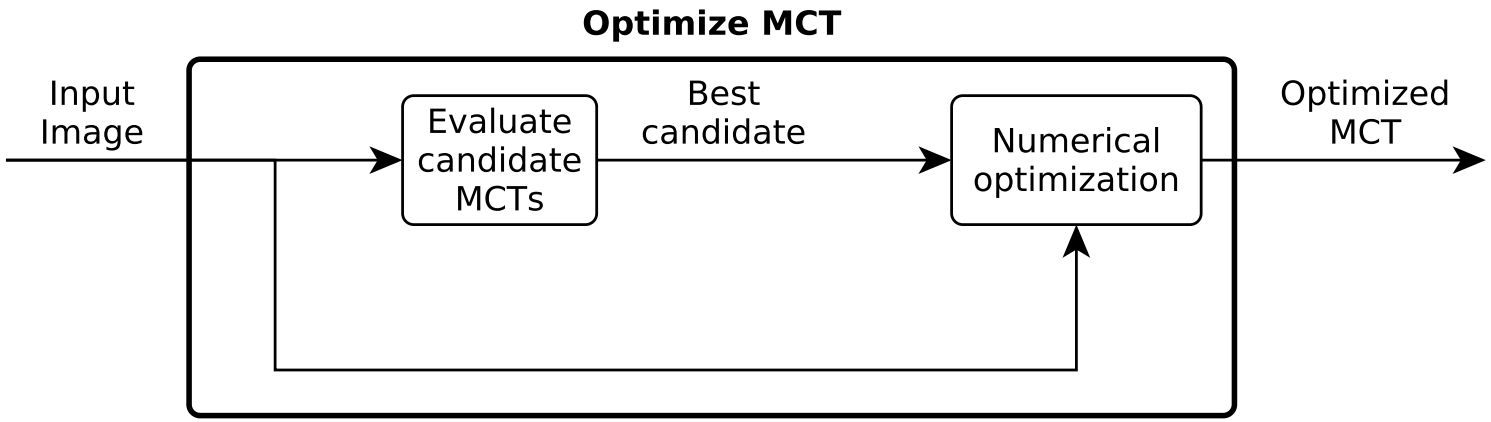

Color-graded HDR

- 0.005-4000 cd/m^2 Rec.2020
- RGB, 1920×1080
- 16-bit integers

Modifications:

Cost function

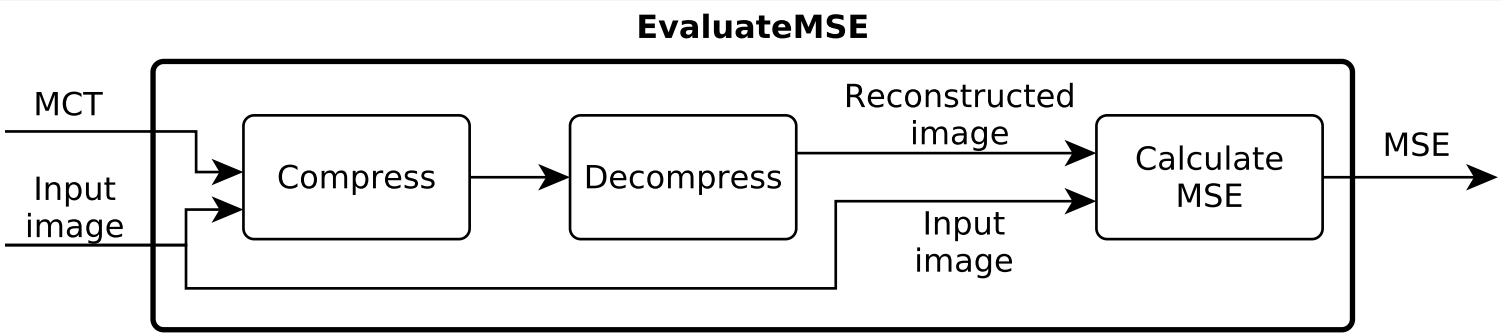

Color-graded HDR

- 0.005-4000 cd/m^2 Rec.2020
- RGB, 1920×1080
- 16-bit integers

Modifications:

Cost function

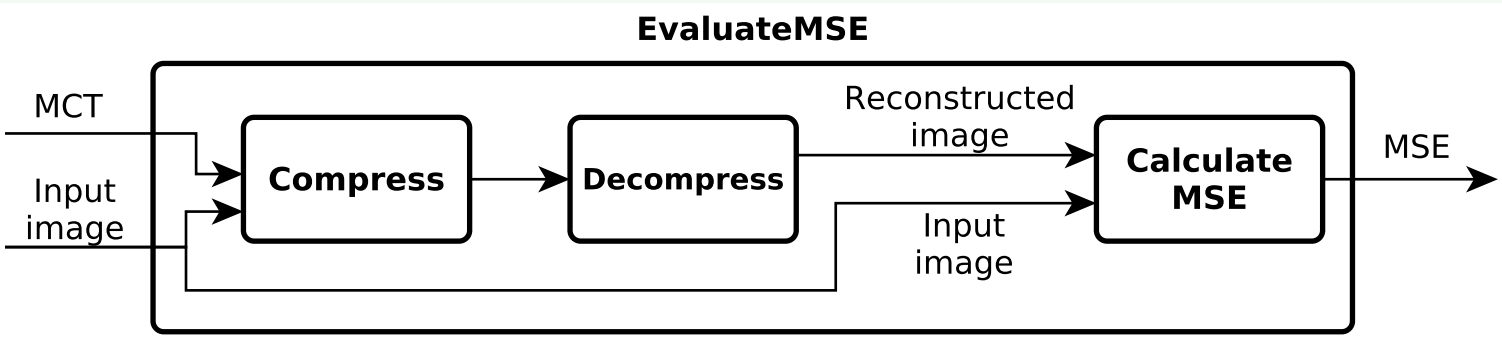

Color-graded HDR

- 0.005-4000 cd/m^2 Rec.2020
- RGB, 1920×1080
- 16-bit integers

Modifications:

Cost function

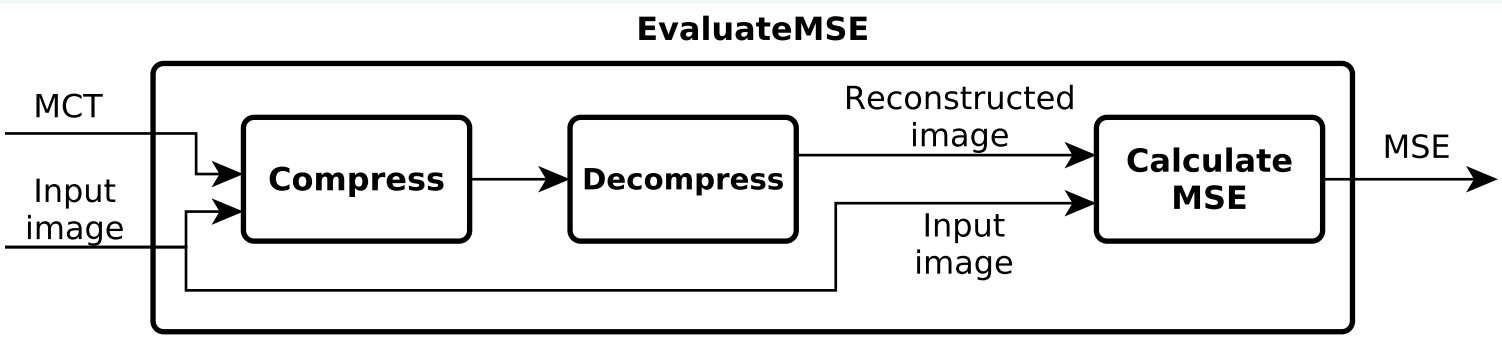

✓ JPEG2000 supports 16-bit images

Scene Radiance HDR

- Raw video footage
- Needs to be “developed”
- 16-bit floats

Scene Radiance HDR

- Raw video footage
- Needs to be “developed”
- 16-bit floats

Scene Radiance HDR

- Raw video footage
- Needs to be “developed”
- 16-bit *floats*
  - 1 bit sign
  - 5 bit exponent
  - 11 bit mantissa

Scene Radiance HDR

- Raw video footage
- Needs to be “developed”
- 16-bit *floats*
  - 1 bit sign
  - 5 bit exponent
  - 11 bit mantissa

  $\Rightarrow$ $\downarrow$values: $\uparrow$precision

  $\Rightarrow$ low average intensity

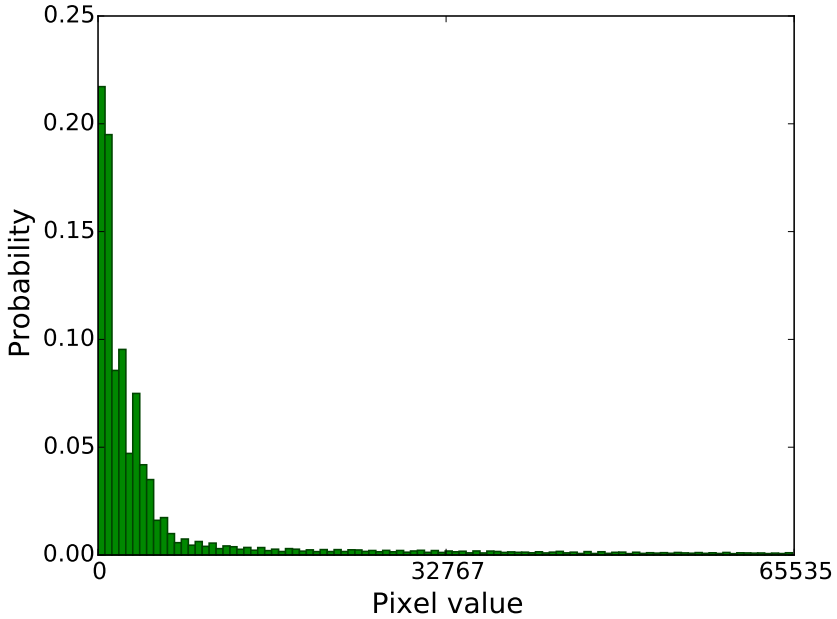

Scene Radiance HDR

- Raw video footage
- Needs to be "developed"
- 16-bit *floats*
  - 1 bit sign
  - 5 bit exponent
  - 11 bit mantissa

  $\Rightarrow$ $\downarrow$values: $\uparrow$precision

  $\Rightarrow$ low average intensity

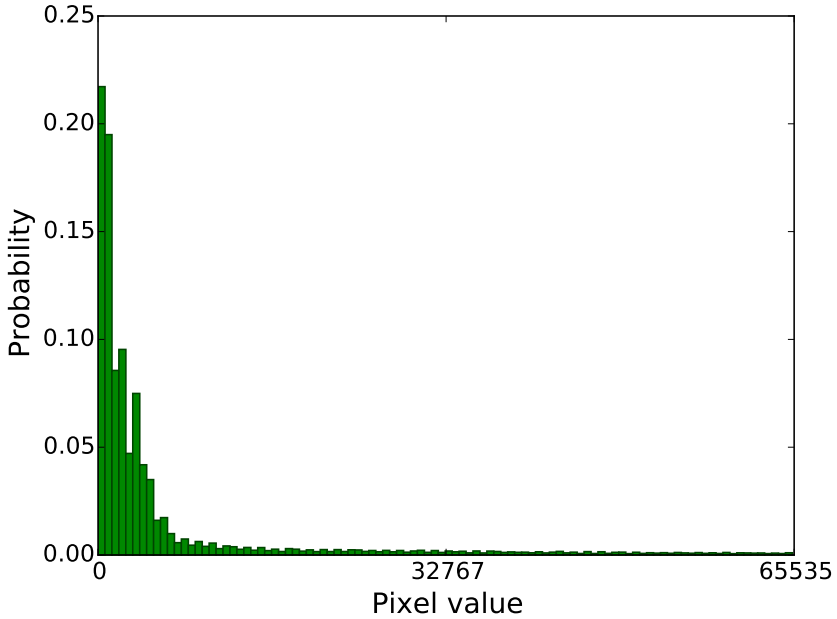

Modifications:

1. [Xu2005] Apply $y = \log_2(x - x_{\min} + 1)$
2. [Xu2005] Uniform quantizer: float $\rightarrow$ integers
   map $[y_{\min}, y_{\max}] \rightarrow [0, 2^{16} - 1]$

Scene Radiance HDR

- Raw video footage
- Needs to be "developed"
- 16-bit *floats*
  - 1 bit sign
  - 5 bit exponent
  - 11 bit mantissa

  $\Rightarrow$ $\downarrow$values: $\uparrow$precision

  $\Rightarrow$ low average intensity

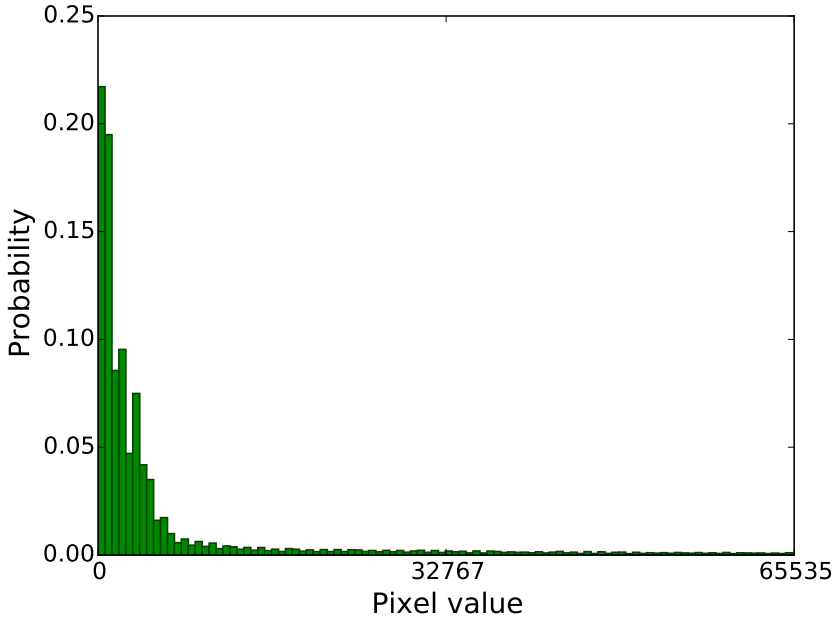

Modifications:

1. [Xu2005] Apply $y = \log_2(x - x_{\min} + 1)$
2. [Xu2005] Uniform quantizer: float $\rightarrow$ integers
   map $[y_{\min}, y_{\max}] \rightarrow [0, 2^{16} - 1]$

✓ Use 16-bit EvaluateMSE

Bayer CFA Mosaics

- Raw camera output
- 10-14-bit integers
- Needs to be “developed”
- R,G,B mosaic

Bayer CFA Mosaics

- Raw camera output
- 10-14-bit integers
- Needs to be “developed”
- R,G,B mosaic

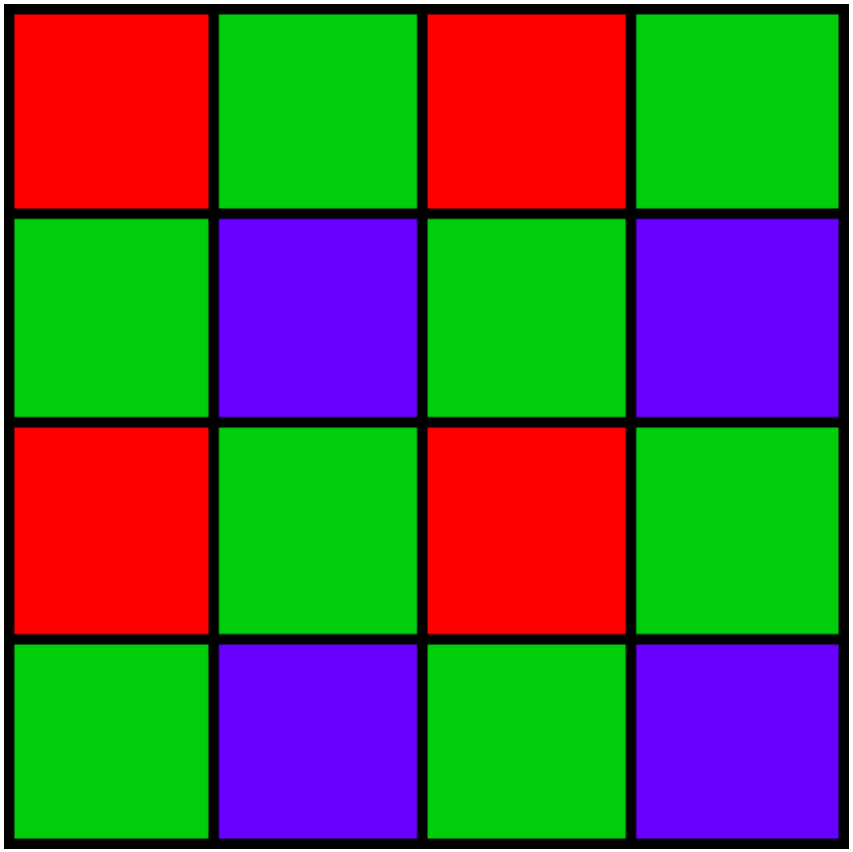

Bayer CFA Mosaics

- Raw camera output
- 10-14-bit integers
- Needs to be "developed"
- R,G,B mosaic

Modifications:

1. Split into 4 components

Following [Koh2003]:

$$\begin{vmatrix} r_{0,0} & G_{0,0} & r_{0,1} & G_{0,1} \\ g_{0,0} & b_{0,0} & g_{0,1} & b_{0,1} \\ r_{1,0} & G_{1,0} & r_{1,1} & G_{1,1} \\ g_{1,0} & b_{1,0} & g_{1,1} & b_{1,1} \end{vmatrix}$$

$$\downarrow$$

$$\begin{vmatrix} r_{0,0} & r_{0,1} \\ r_{1,0} & r_{1,1} \end{vmatrix} \quad \begin{vmatrix} G_{0,0} & G_{0,1} \\ G_{1,0} & G_{1,1} \end{vmatrix}$$

$$\begin{vmatrix} g_{0,0} & g_{0,1} \\ g_{1,0} & g_{1,1} \end{vmatrix} \quad \begin{vmatrix} b_{0,0} & b_{0,1} \\ b_{1,0} & b_{1,1} \end{vmatrix}$$

Bayer CFA Mosaics

- Raw camera output
- 10-14-bit integers
- Needs to be "developed"
- R,G,B mosaic

Modifications:

1. Split into 4 components
2. Optimize 4×4 matrix

$$\begin{pmatrix} m_{0,0} & m_{0,1} & m_{0,2} & m_{0,3} \\ m_{1,0} & m_{1,1} & m_{1,2} & m_{1,3} \\ m_{2,0} & m_{2,1} & m_{2,2} & m_{2,3} \\ m_{3,0} & m_{3,1} & m_{3,2} & m_{3,3} \end{pmatrix}$$

Following [Koh2003]:

$$\begin{vmatrix} r_{0,0} & G_{0,0} & r_{0,1} & G_{0,1} \\ g_{0,0} & b_{0,0} & g_{0,1} & b_{0,1} \\ r_{1,0} & G_{1,0} & r_{1,1} & G_{1,1} \\ g_{1,0} & b_{1,0} & g_{1,1} & b_{1,1} \end{vmatrix}$$

$$\downarrow$$

$$\begin{vmatrix} r_{0,0} & r_{0,1} \\ r_{1,0} & r_{1,1} \end{vmatrix} \quad \begin{vmatrix} G_{0,0} & G_{0,1} \\ G_{1,0} & G_{1,1} \end{vmatrix}$$

$$\begin{vmatrix} g_{0,0} & g_{0,1} \\ g_{1,0} & g_{1,1} \end{vmatrix} \quad \begin{vmatrix} b_{0,0} & b_{0,1} \\ b_{1,0} & b_{1,1} \end{vmatrix}$$

Bayer CFA Mosaics

- Raw camera output
- 10-14-bit integers
- Needs to be "developed"
- R,G,B mosaic

Following [Koh2003]:

$$\begin{vmatrix} r_{0,0} & G_{0,0} & r_{0,1} & G_{0,1} \\ g_{0,0} & b_{0,0} & g_{0,1} & b_{0,1} \\ r_{1,0} & G_{1,0} & r_{1,1} & G_{1,1} \\ g_{1,0} & b_{1,0} & g_{1,1} & b_{1,1} \end{vmatrix}$$

$$\downarrow$$

$$\begin{vmatrix} r_{0,0} & r_{0,1} \\ r_{1,0} & r_{1,1} \end{vmatrix} \quad \begin{vmatrix} G_{0,0} & G_{0,1} \\ G_{1,0} & G_{1,1} \end{vmatrix}$$

$$\begin{vmatrix} g_{0,0} & g_{0,1} \\ g_{1,0} & g_{1,1} \end{vmatrix} \quad \begin{vmatrix} b_{0,0} & b_{0,1} \\ b_{1,0} & b_{1,1} \end{vmatrix}$$

Modifications:

1. Split into 4 components
2. Optimize 4×4 matrix

$$\begin{pmatrix} m_{0,0} & m_{0,1} & m_{0,2} & m_{0,3} \\ m_{1,0} & m_{1,1} & m_{1,2} & m_{1,3} \\ m_{2,0} & m_{2,1} & m_{2,2} & m_{2,3} \\ m_{3,0} & m_{3,1} & m_{3,2} & m_{3,3} \end{pmatrix}$$

✓ JPEG2000 supports 4×4 MCTs

Contents

Color-graded HdM-HDR-2014

- Froehlich et al.'s **Color-graded HdM-HDR-2014**
- 16 sequences, 10 frames/sequence = 160 images
- 1920 $\times$ 1080, 16-bit int, RGB

Color-graded HdM-HDR-2014

RD for Beerfest frame

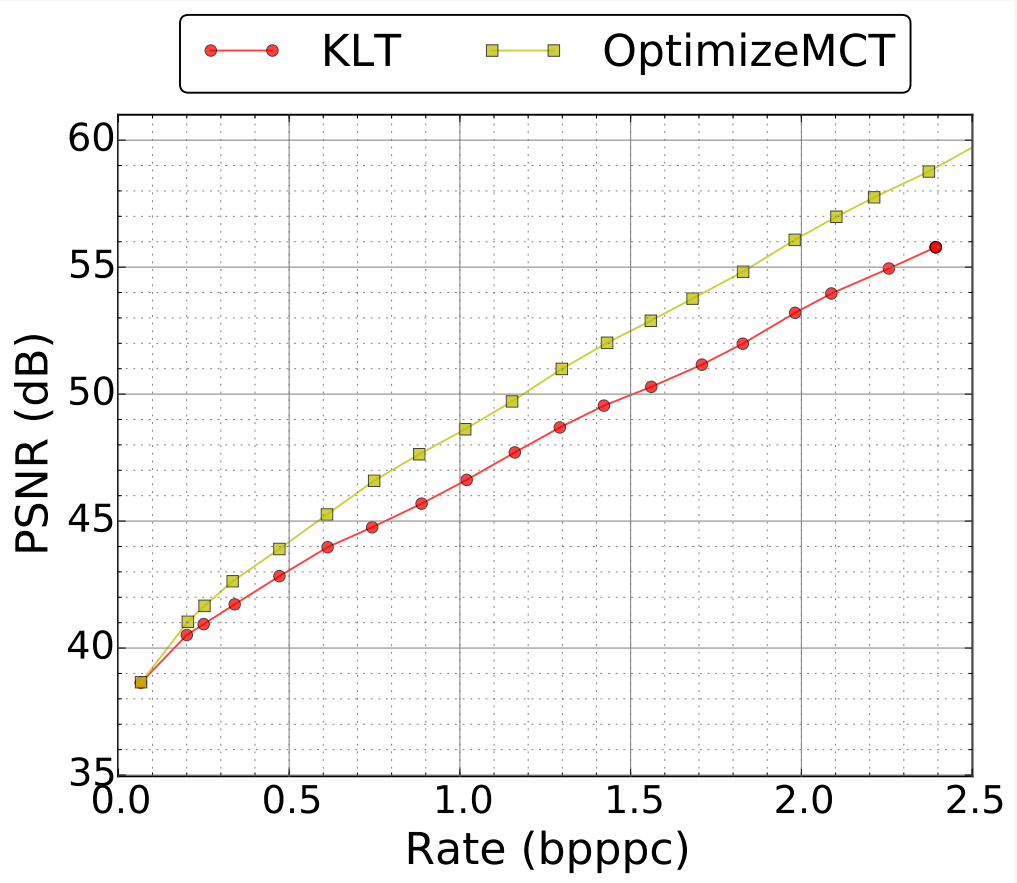

- Gains across all bitrates
- ↑rate ⇒ ↑gain

Color-graded HdM-HDR-2014

Results for medium bitrates (160 images):

Avg. Rate	KLT	No MCT	OST	OptimizeMCT
0.962 bpp				

Color-graded HdM-HDR-2014

Results for medium bitrates (160 images):

Avg. Rate	KLT	No MCT	OST	OptimizeMCT
0.962 bpp	44.86 dB (anchor)	44.17 dB -1.55%	44.03 dB -1.86%	**45.15 dB +0.63%**

Color-graded HdM-HDR-2014

Results for medium bitrates (160 images):

Avg. Rate	KLT	No MCT	OST	OptimizeMCT
0.962 bpp	44.86 dB (anchor)	44.17 dB -1.55%	44.03 dB -1.86%	**45.15 dB +0.63%**

- OptimizeMCT $>$ KLT, OST
- Better for all images

Largest/smallest improvements

Largest improvements (1.44% – 1.77%)



Largest/smallest improvements

Largest improvements (1.44% – 1.77%)



Smallest improvements (0.06% – 0.11%)



Scene radiance HdM-HDR-2014

- Froehlich et al.'s
  Reconstructed scene radiance HdM-HDR-2014
- 16 sequences, 10 frames/sequence = 160 images
- 1920 × 1080, 16-bit float, RGB

Scene radiance HdM-HDR-2014

RD for Beerfest frame

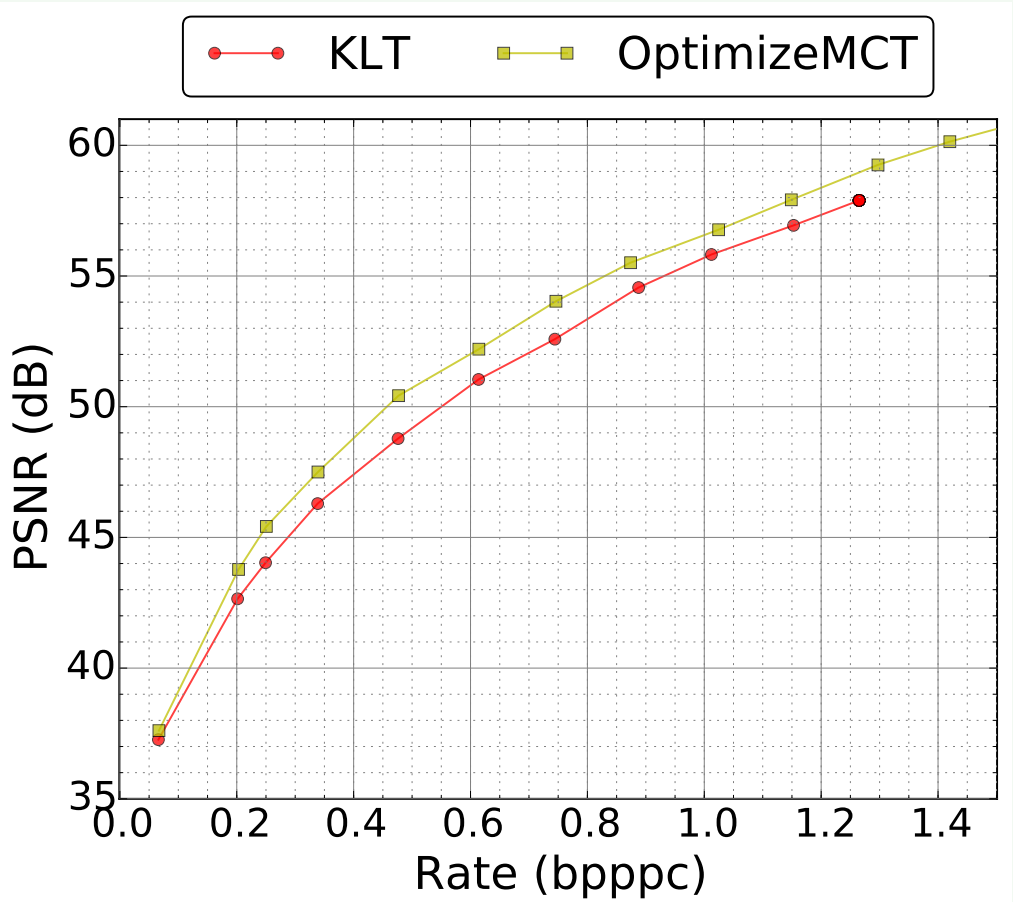

- Gains across all bitrates
- ↑rate ⇒ ↑gain
- Smaller gains than for color-graded

Scene radiance HdM-HDR-2014

Results for medium bitrates (160 images):

Avg. Rate	KLT	No MCT	OST	OptimizeMCT
0.554 bpp				

Scene radiance HdM-HDR-2014

Results for medium bitrates (160 images):

Avg. Rate	KLT	No MCT	OST	OptimizeMCT
0.554 bpp	49.37 dB (anchor)	43.75 dB -11.39%	48.85 dB -1.05%	**49.60 dB +0.47%**

Scene radiance HdM-HDR-2014

Results for medium bitrates (160 images):

Avg. Rate	KLT	No MCT	OST	OptimizeMCT
0.554 bpp	49.37 dB (anchor)	43.75 dB -11.39%	48.85 dB -1.05%	**49.60 dB +0.47%**

- OptimizeMCT $>$ KLT, OST
- Better for all images

Scene radiance HdM-HDR-2014

Largest improvements (0.90% – 1.31%)



Scene radiance HdM-HDR-2014

Largest improvements (0.90% – 1.31%)



Smallest improvements (0.09% – 0.14%)



Bayer CFA

$$\longrightarrow \quad \begin{vmatrix} r_{0,0} & r_{0,1} \\ r_{1,0} & r_{1,1} \end{vmatrix} \quad \begin{vmatrix} G_{0,0} & G_{0,1} \\ G_{1,0} & G_{1,1} \end{vmatrix} \quad \begin{vmatrix} g_{0,0} & g_{0,1} \\ g_{1,0} & g_{1,1} \end{vmatrix} \quad \begin{vmatrix} b_{0,0} & b_{0,1} \\ b_{1,0} & b_{1,1} \end{vmatrix}$$

- 20 Nikon camera models
- 20 real Bayer CFA images
- 3000 $\times$ 2000, 12–14-bit integers, 4 components

Bayer CFA

RD for DS200 Image

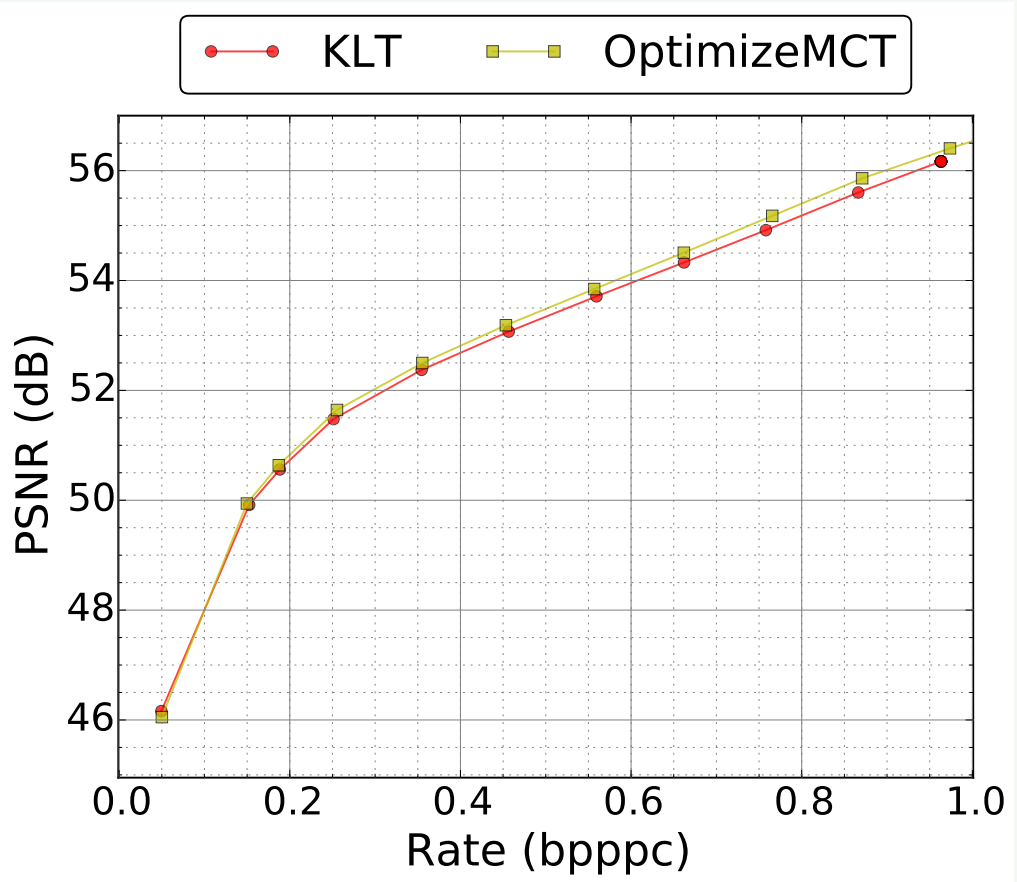

- Gains across all bitrates
- Smaller gains than for color-graded/radiance

Bayer CFA

Average 20 images:

Avg. Rate	KLT	No MCT	OST	OptimizeMCT
0.413 bpp				

Bayer CFA

Average 20 images:

Avg. Rate	KLT	No MCT	OST	OptimizeMCT
0.413 bpp	47.37 dB (anchor)	45.96 dB -2.99%	47.33 dB -0.08%	**47.45 dB +0.17%**

Bayer CFA

Average 20 images:

Avg. Rate	KLT	No MCT	OST	OptimizeMCT
0.413 bpp	47.37 dB (anchor)	45.96 dB -2.99%	47.33 dB -0.08%	**47.45 dB +0.17%**

- OptimizeMCT $>$ KLT, OST
- Better for all images

Conclusions & Future Work

- OptimizeMCT: can be adapted to HDR

Conclusions & Future Work

- OptimizeMCT: can be adapted to HDR
- Improve upon KLT, OST (PSNR, all images)

Set	**OST vs KLT**	**OptimizeMCT vs KLT**
Color graded	-1.86%	**+0.63%**
Scene radiance	-1.05%	**+0.47%**
Bayer CFA	-0.08%	**+0.17%**

Conclusions & Future Work

- OptimizeMCT: can be adapted to HDR
- Improve upon KLT, OST (PSNR, all images)

Set	**OST vs KLT**	**OptimizeMCT vs KLT**
Color graded	-1.86%	**+0.63%**
Scene radiance	-1.05%	**+0.47%**
Bayer CFA	-0.08%	**+0.17%**

- Future:
  - New metrics: HDR-VDP-2
  - New coders: HEVC

Conclusions & Future Work

- OptimizeMCT: can be adapted to HDR
- Improve upon KLT, OST (PSNR, all images)

Set	OST vs KLT	OptimizeMCT vs KLT
Color graded	-1.86%	**+0.63%**
Scene radiance	-1.05%	**+0.47%**
Bayer CFA	-0.08%	**+0.17%**

- Future:
  - New metrics: HDR-VDP-2
  - New coders: HEVC

MCT Optimization

Decisions

? Optimization algorithm

MCT Optimization

Decisions

? Optimization algorithm

- Global vs local?
- Optimization algorithm?

MCT Optimization

Decisions

? Optimization algorithm

- Global vs local?
- Optimization algorithm?

MCT Optimization

Decisions

? Optimization algorithm

- Global vs local?
- Optimization algorithm?

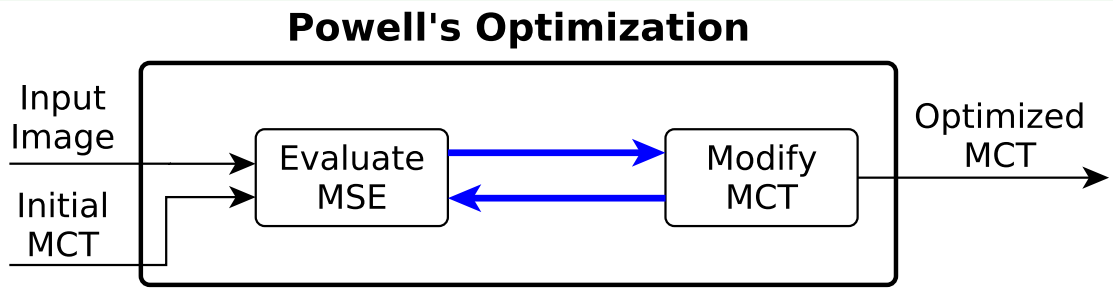

MCT Optimization

Decisions

- ? Optimization algorithm
- ? Initial MCT

PSNR Evolution (4 candidates)

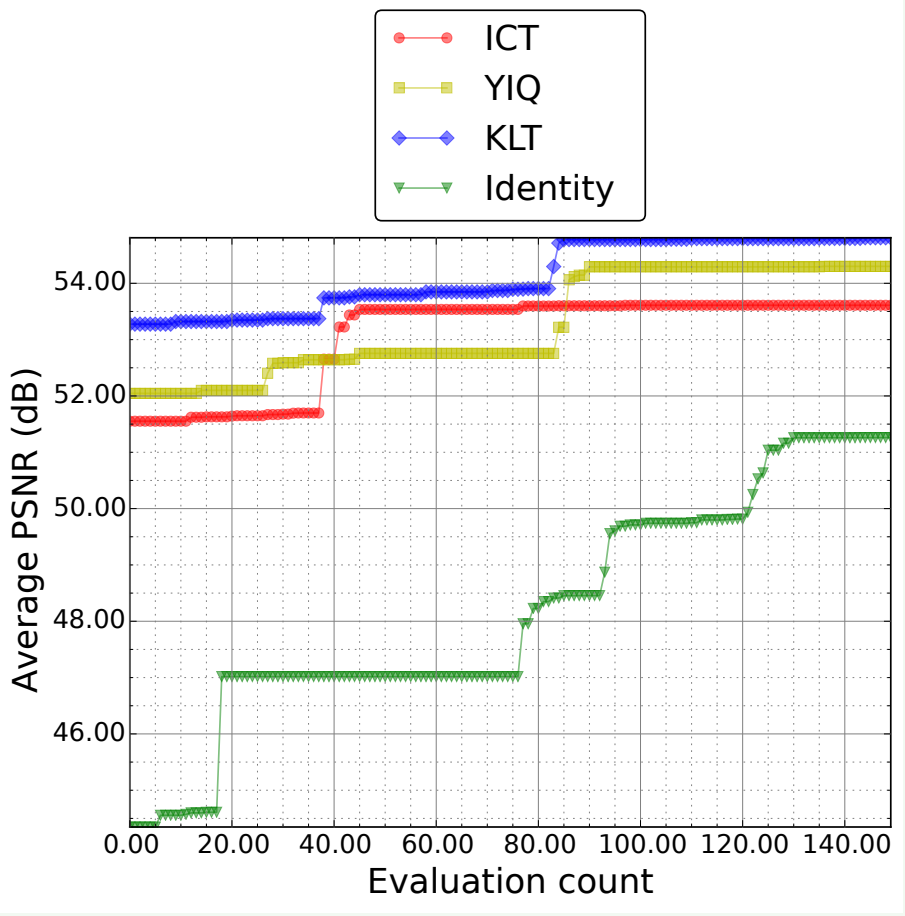

MCT Optimization

Decisions

- ✓ Optimization algorithm
- ✓ Initial MCT

Proposed Algorithm

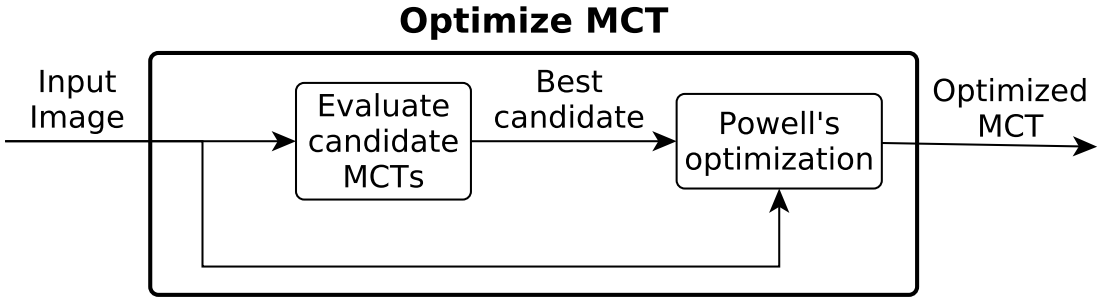

1. Evaluate N=3 candidates
2. Optimize best candidate

MCT Optimization

Decisions

- ✓ Optimization algorithm
- ✓ Initial MCT
- ? Iteration count

Proposed Algorithm

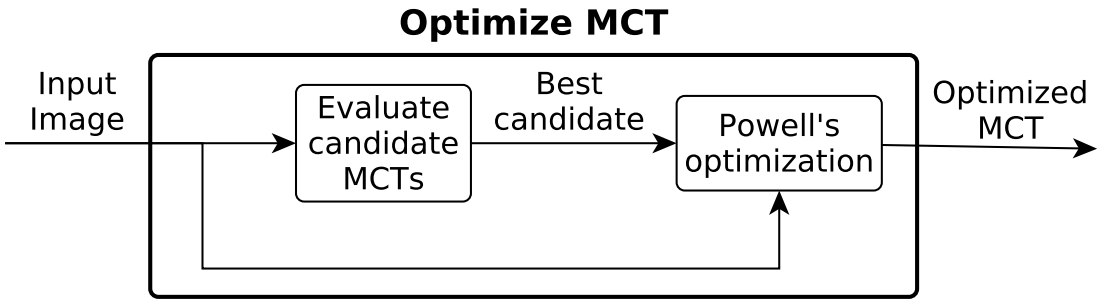

1. Evaluate N=3 candidates
2. Optimize best candidate

MCT Optimization

Decisions

- ✓ Optimization algorithm
- ✓ Initial MCT
- ✓ Iteration count

Proposed Algorithm

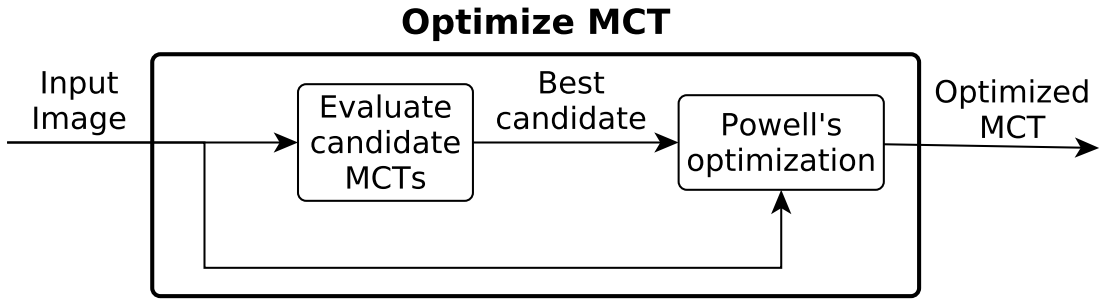

1. Evaluate N candidates
2. Optimize best candidate

Limit to 250–300 iterations