

High dynamic range imaging


Digital Visual Effects
Yung-Yu Chuang

with slides by Frdo Durand, Brian Curless, Steve Seitz, Paul Debevec and Alexei Efros

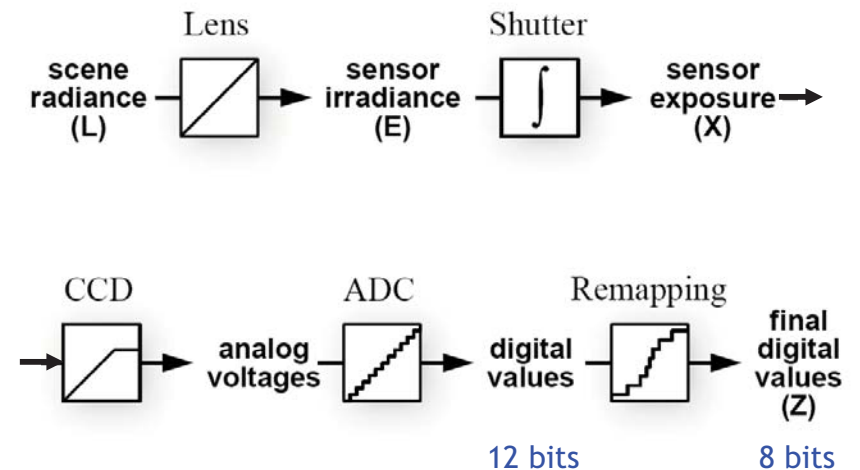
Camera is an imperfect device

- Camera is an imperfect device for measuring the radiance distribution of a scene because it cannot capture the full spectral content and dynamic range.
- Limitations in sensor design prevent cameras from capturing all information passed by lens.

Camera pipeline




Camera pipeline

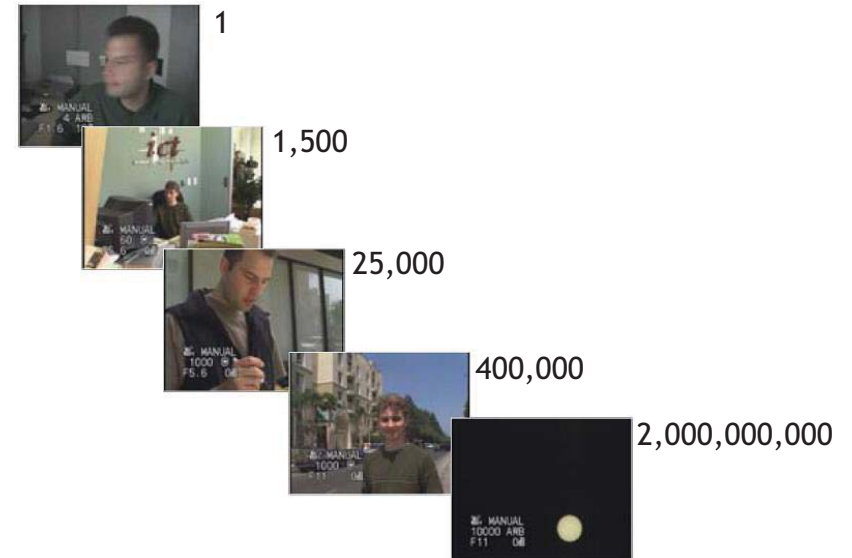


Real-world response functions


In general, the response function is not provided by camera makers who consider it part of their proprietary product differentiation. In addition, they are beyond the standard gamma curves.



The world is high dynamic range




The world is high dynamic range




Real world dynamic range

- Eye can adapt from $\sim 10^{-6}$ to 10^6 cd/m²
- Often 1 : 100,000 in a scene
- Typical 1:50, max 1:500 for pictures




Short exposure

DigiVFX



Long exposure

DigiVFX



Camera is not a photometer


DigiVFX

- Limited dynamic range
 - ⇒ Perhaps use multiple exposures?
- Unknown, nonlinear response
 - ⇒ Not possible to convert pixel values to radiance
- Solution:
 - Recover response curve from multiple exposures, then reconstruct the **radiance map**

Varying exposure

DigiVFX

- Ways to change exposure
 - Shutter speed
 - Aperture
 - Neutral density filters



Shutter speed

DigiVFX


- Note: shutter times usually obey a power series - each “stop” is a factor of 2
- $\frac{1}{4}$, $\frac{1}{8}$, $\frac{1}{15}$, $\frac{1}{30}$, $\frac{1}{60}$, $\frac{1}{125}$, $\frac{1}{250}$, $\frac{1}{500}$, $\frac{1}{1000}$ sec

Usually really is:

$\frac{1}{4}$, $\frac{1}{8}$, $\frac{1}{16}$, $\frac{1}{32}$, $\frac{1}{64}$, $\frac{1}{128}$, $\frac{1}{256}$, $\frac{1}{512}$, $\frac{1}{1024}$ sec

Varying shutter speeds

DigiVFX



HDRI capturing from multiple exposures


DigiVFX

- Capture images with multiple exposures
- Image alignment (even if you use tripod, it is suggested to run alignment)
- Response curve recovery
- Ghost/flare removal

Image alignment

DigiVFX

- We will introduce a fast and easy-to-implement method for this task, called Median Threshold Bitmap (MTB) alignment technique.
- Consider only integral translations. It is enough empirically.
- The inputs are N grayscale images. (You can either use the green channel or convert into grayscale by $Y = (54R + 183G + 19B) / 256$)
- MTB is a binary image formed by thresholding the input image using the median of intensities.



Why is MTB better than gradient?


DigiVFX

- Edge-detection filters are dependent on image exposures
- Taking the difference of two edge bitmaps would not give a good indication of where the edges are misaligned.

Search for the optimal offset


DigiVFX

- Try all possible offsets.
- Gradient descent
- Multiscale technique
- $\log(\text{max_offset})$ levels
- Try 9 possibilities for the top level
- Scale by 2 when passing down; try its 9 neighbors



Threshold noise

DigiVFX



Efficiency considerations

DigiVFX

- XOR for taking difference
- AND with exclusion maps
- Bit counting by table lookup

Results

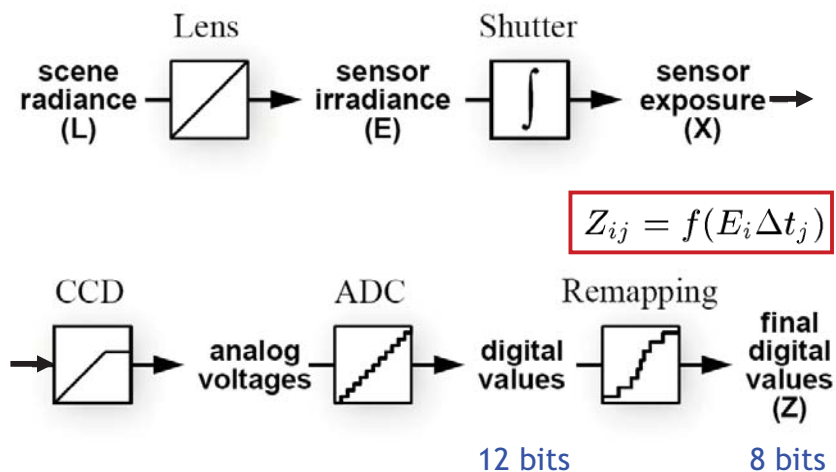
DigiVFX

Success rate = 84%. 10% failure due to rotation. 3% for excessive motion and 3% for too much high-frequency content.



Recovering response curve



DigiVFX



Recovering response curve

DigiVFX



- We want to obtain the inverse of the response curve



Recovering response curve

DigiVFX



Image series



Recovering response curve

DigiVFX


Image series



Idea behind the math

DigiVFX


plot of $g(Z_{ij})$ from three pixels observed in five images, assuming unit radiance at each pixel




Idea behind the math

DigiVFX

plot of $g(Z_{ij})$ from three pixels observed in five images, assuming unit radiance at each pixel



Idea behind the math



Basic idea

- Design an objective function
- Optimize it

Math for recovering response curve

$$Z_{ij} = f(E_i \Delta t_j)$$

f is monotonic, it is invertible

$$\ln f^{-1}(Z_{ij}) = \ln E_i + \ln \Delta t_j$$

let us define function $g = \ln f^{-1}$

$$g(Z_{ij}) = \ln E_i + \ln \Delta t_j$$

minimize the following

$$\mathcal{O} = \sum_{i=1}^N \sum_{j=1}^P [g(Z_{ij}) - \ln E_i - \ln \Delta t_j]^2 + \lambda \sum_{z=Z_{min}+1}^{Z_{max}-1} g''(z)^2$$
$$g''(z) = g(z-1) - 2g(z) + g(z+1)$$

$$g''(z) = g(z-1) - 2g(z) + g(z+1)$$

Recovering response curve



- The solution can be only up to a scale, add a constraint

$$g(Z_{mid}) = 0, \text{ where } Z_{mid} = \frac{1}{2}(Z_{min} + Z_{max})$$

- Add a hat weighting function

$$w(z) = \begin{cases} z - Z_{min} & \text{for } z \leq \frac{1}{2}(Z_{min} + Z_{max}) \\ Z_{max} - z & \text{for } z > \frac{1}{2}(Z_{min} + Z_{max}) \end{cases}$$

$$\mathcal{O} = \sum_{i=1}^N \sum_{j=1}^P \{w(Z_{ij}) [g(Z_{ij}) - \ln E_i - \ln \Delta t_j]\}^2 +$$

$$\lambda \sum_{z=Z_{min}+1}^{Z_{max}-1} [w(z)g''(z)]^2$$

Recovering response curve



- We want $N(P-1) > (Z_{max} - Z_{min})$
If $P=11$, $N \sim 25$ (typically 50 is used)
- We prefer that selected pixels are well distributed and sampled from constant regions. They picked points by hand.
- It is an overdetermined system of linear equations and can be solved using SVD

How to optimize?



$$\mathcal{O} = \sum_{i=1}^N \sum_{j=1}^P \{w(Z_{ij}) [g(Z_{ij}) - \ln E_i - \ln \Delta t_j]\}^2 +$$

$$\lambda \sum_{z=Z_{min}+1}^{Z_{max}-1} [w(z)g''(z)]^2$$

1. Set partial derivatives to zero



$$\mathcal{O} = \sum_{i=1}^N \sum_{j=1}^P \{w(Z_{ij}) [g(Z_{ij}) - \ln E_i - \ln \Delta t_j]\}^2 +$$

$$\lambda \sum_{z=Z_{min}+1}^{Z_{max}-1} [w(z)g''(z)]^2$$

How to optimize?



$$\mathcal{O} = \sum_{i=1}^N \sum_{j=1}^P \{w(Z_{ij}) [g(Z_{ij}) - \ln E_i - \ln \Delta t_j]\}^2 +$$

$$\lambda \sum_{z=Z_{min}+1}^{Z_{max}-1} [w(z)g''(z)]^2$$

1. Set partial derivatives to zero
- 2.

$$2. \quad \min \sum_{i=1}^N (\mathbf{a}_i \mathbf{x} - \mathbf{b}_i)^2 \rightarrow \text{least-square solution of } \begin{bmatrix} \mathbf{a}_1 \\ \mathbf{a}_2 \\ \vdots \\ \mathbf{a}_N \end{bmatrix} \mathbf{x} = \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \\ \vdots \\ \mathbf{b}_N \end{bmatrix}$$

Sparse linear system



Diagram illustrating the structure of the matrix A in the equation $Ax=b$. The matrix A is partitioned into blocks. The top-left block is labeled 256 and n . The bottom-left block is labeled 1 and 254 . The right side of the equation shows a vector with elements $g(0)$, \vdots , $g(255)$, $\ln E_1$, \vdots , $\ln E_n$. A blue dashed line separates the top 256 elements from the bottom 254 elements. A red dashed line separates the single element 1 from the 254 elements below it.

Questions



- Will $g(127)=0$ always be satisfied? Why or why not?
- How to find the least-square solution for an over-determined system?

Least-square solution for a linear system



$$\begin{matrix} & \mathbf{A} & \mathbf{x} & = & \mathbf{b} \\ m \times n & n & & & m \\ m > n & & & & \end{matrix}$$

They are often mutually incompatible. We instead find \mathbf{x} to minimize the norm $\|\mathbf{Ax} - \mathbf{b}\|$ of the residual vector $\mathbf{Ax} - \mathbf{b}$. If there are multiple solutions, we prefer the one with the minimal length $\|\mathbf{x}\|$.

Least-square solution for a linear system DigiVFX

If we perform SVD on A and rewrite it as

$$A = U \Sigma V^T$$

then $\hat{x} = \boxed{V \Sigma^+ U^T} b$ is the least-square solution.
pseudo inverse

$$\Sigma^+ = \begin{bmatrix} 1/\sigma_1 & & & 0 & \dots & 0 \\ & \ddots & & & & \\ & & 1/\sigma_r & & & \\ & & & 0 & & \\ & & & & \ddots & \\ & & & & & 0 & 0 & \dots & 0 \end{bmatrix}$$

Proof DigiVFX

find x 使 $\|Ax - b\|$ 最小

$$\|Ax - b\| = \|U \Sigma V^T x - b\|$$

$$= \|U (\Sigma V^T x - U^T b)\|$$

U 是 rotation 不动长度

$$= \|\Sigma V^T x - U^T b\|$$

$$\text{令 } y = V^T x \quad c = U^T b$$

则 相当于 找 y 使 $\|\Sigma y - c\|$ 最小

$$\begin{pmatrix} \sigma_1 & \dots & 0 \\ & \ddots & \\ & & \sigma_r & \dots & 0 \\ & & & \ddots & \\ 0 & & & & 0 \end{pmatrix} \begin{pmatrix} y_1 \\ \vdots \\ y_r \\ \vdots \\ y_n \end{pmatrix} = \begin{pmatrix} c_1 \\ \vdots \\ c_r \\ \vdots \\ c_n \end{pmatrix}$$

Proof DigiVFX

$$\Rightarrow y_i = \frac{c_i}{\sigma_i} \quad i=1 \dots r \quad y_i = 0 \quad i=r+1 \dots n$$

$$\Rightarrow \tilde{y} = \begin{pmatrix} 1/\sigma_1 & \dots & 0 \\ & \ddots & \\ & & 1/\sigma_r & \dots & 0 \\ & & & \ddots & \\ 0 & & & & 0 \end{pmatrix} \begin{pmatrix} c_1 \\ \vdots \\ c_r \\ \vdots \\ c_n \end{pmatrix} = \Sigma^+ c$$

$$\Rightarrow \tilde{y} = V^T \tilde{x} = \Sigma^+ c = \Sigma^+ U^T b$$

$$\Rightarrow \tilde{x} = V \Sigma^+ U^T b$$

Libraries for SVD DigiVFX

- Matlab
- GSL
- Boost
- LAPACK
- ATLAS

Matlab code



```
%
% gsolve.m - Solve for imaging system response function
%
% Given a set of pixel values observed for several pixels in several
% images with different exposure times, this function returns the
% imaging system's response function g as well as the log film irradiance
% values for the observed pixels.
%
% Assumes:
%
%   Zmin = 0
%   Zmax = 255
%
% Arguments:
%
%   Z(i,j) is the pixel values of pixel location number i in image j
%   B(j)   is the log delta t, or log shutter speed, for image j
%   l      is lambda, the constant that determines the amount of smoothness
%   w(z)   is the weighting function value for pixel value z
%
% Returns:
%
%   g(z)   is the log exposure corresponding to pixel value z
%   lE(i)  is the log film irradiance at pixel location i
%
```

Matlab code



```
function [g,lE]=gsolve(Z,B,l,w)

n = 256;
A = zeros(size(Z,1)*size(Z,2)+n+1,n+size(Z,1));
b = zeros(size(A,1),1);

k = 1;           %% Include the data-fitting equations
for i=1:size(Z,1)
    for j=1:size(Z,2)
        wij = w(Z(i,j)+1);
        A(k,Z(i,j)+1) = wij; A(k,n+i) = -wij; b(k,1) = wij * B(j);
        k=k+1;
    end
end


A(k,129) = 1;    %% Fix the curve by setting its middle value to 0
k=k+1;

for i=1:n-2      %% Include the smoothness equations
    A(k,i)=l*w(i+1); A(k,i+1)=-2*l*w(i+1); A(k,i+2)=l*w(i+1);
    k=k+1;
end

x = A\b;         %% Solve the system using SVD

g = x(1:n);
lE = x(n+1:size(x,1));
```

Recovered response function



Constructing HDR radiance map




$$\ln E_i = g(Z_{ij}) - \ln \Delta t_j$$

combine pixels to reduce noise and obtain a more reliable estimation

$$\ln E_i = \frac{\sum_{j=1}^P w(Z_{ij})(g(Z_{ij}) - \ln \Delta t_j)}{\sum_{j=1}^P w(Z_{ij})}$$


Reconstructed radiance map

DigiVFX



What is this for?

DigiVFX



- Human perception
- Vision/graphics applications

Automatic ghost removal

DigiVFX




before



after

Weighted variance


DigiVFX



Moving objects and high-contrast edges render high variance.

Region masking

DigiVFX



Thresholding; dilation; identify regions;

Best exposure in each region

DigiVFX



Lens flare removal

DigiVFX




before

after


Easier HDR reconstruction

DigiVFX



Easier HDR reconstruction


DigiVFX



Portable floatMap (.pfm)

DigiVFX

- 12 bytes per pixel, 4 for each channel



sign exponent mantissa


Text header similar to Jeff Poskanzer's .ppm image format:

```
PF
768 512
1
<binary image data>
```

Floating Point TIFF similar

Radiance format (.pic, .hdr, .rad)

DigiVFX



$$\begin{aligned}
 (145, 215, 87, 149) &= (145, 215, 87) * 2^{(149-128)} = 1190000 \ 1760000 \ 713000 \\
 (145, 215, 87, 103) &= (145, 215, 87) * 2^{(103-128)} = 0.00000432 \ 0.00000641 \ 0.00000259
 \end{aligned}$$

Ward, Greg. "Real Pixels," in Graphics Gems IV, edited by James Arvo, Academic Press, 1994

ILM's OpenEXR (.exr)

DigiVFX

- 6 bytes per pixel, 2 for each channel, compressed



sign exponent mantissa

- Several lossless compression options, 2:1 typical
- Compatible with the "half" datatype in NVidia's Cg
- Supported natively on GeForce FX and Quadro FX


- Available at <http://www.openexr.net/>

Radiometric self calibration

- Assume that any response function can be modeled as a high-order polynomial

$$X = g(Z) = \sum_{m=0}^M c_m Z^m$$

- No need to know exposure time in advance. Useful for cheap cameras



Mitsunaga and Nayar

- To find the coefficients c_m to minimize the following

$$\mathcal{E} = \sum_{i=1}^N \sum_{j=1}^P \left[\sum_{m=0}^M c_m Z_{ij}^m - R_{j,j+1} \sum_{m=0}^M c_m Z_{i,j+1}^m \right]^2$$

A guess for the ratio of

$$\frac{X_{ij}}{X_{i,j+1}} = \frac{E_i \Delta t_j}{E_i \Delta t_{j+1}} = \frac{\Delta t_j}{\Delta t_{j+1}}$$

Mitsunaga and Nayar

- Again, we can only solve up to a scale. Thus, add a constraint $f(1)=1$. It reduces to M variables.
- How to solve it?


Mitsunaga and Nayar

- We solve the above iteratively and update the exposure ratio accordingly


$$R_{j,j+1}^{(k)} = \frac{1}{N} \sum_{i=1}^N \frac{\sum_{m=0}^M c_m^{(k)} Z_{ij}^m}{\sum_{m=0}^M c_m^{(k)} Z_{i,j+1}^m}$$

- How to determine M ? Solve up to $M=10$ and pick up the one with the minimal error. Notice that you prefer to have the same order for all channels. Use the combined error.

Space of response curves



Space of response curves



Robertson et. al.

$$Z_{ij} = f(E_i \Delta t_j)$$

$$g(Z_{ij}) = f^{-1}(Z_{ij}) = E_i \Delta t_j$$

Given Z_{ij} and Δt_j , the goal is to find both E_i and $g(Z_{ij})$

Maximum likelihood

$$\Pr(E_i, g | Z_{ij}, \Delta t_j) \propto \exp\left(-\frac{1}{2} \sum_{ij} w(Z_{ij}) (g(Z_{ij}) - E_i \Delta t_j)^2\right)$$

$$\hat{g}, \hat{E}_i = \arg \min_{g, E_i} \sum_{ij} w(Z_{ij}) (g(Z_{ij}) - E_i \Delta t_j)^2$$

Robertson et. al.

$$\hat{g}, \hat{E}_i = \arg \min_{g, E_i} \sum_{ij} w(Z_{ij}) (g(Z_{ij}) - E_i \Delta t_j)^2$$

repeat

assuming $g(Z_{ij})$ is known, optimize for E_i

assuming E_i is known, optimize for $g(Z_{ij})$

until converge

$$\hat{g}, \hat{E}_i = \arg \min_{g, E_i} \sum_{ij} w(Z_{ij}) (g(Z_{ij}) - E_i \Delta t_j)^2$$

repeat

assuming $g(Z_{ij})$ is known, optimize for E_i

assuming E_i is known, optimize for $g(Z_{ij})$

until converge

$$\hat{g}, \hat{E}_i = \arg \min_{g, E_i} \sum_{ij} w(Z_{ij}) (g(Z_{ij}) - E_i \Delta t_j)^2$$

repeat

assuming $g(Z_{ij})$ is known, optimize for E_i

assuming E_i is known, optimize for $g(Z_{ij})$

until converge

$$E_i = \frac{\sum_j w(Z_{ij}) g(Z_{ij}) \Delta t_j}{\sum_j w(Z_{ij}) \Delta t_j^2}$$

$$\hat{g}, \hat{E}_i = \arg \min_{g, E_i} \sum_{ij} w(Z_{ij}) (g(Z_{ij}) - E_i \Delta t_j)^2$$

repeat

assuming $g(Z_{ij})$ is known, optimize for E_i

assuming E_i is known, optimize for $g(Z_{ij})$

until converge


$$g(m) = \frac{1}{|E_m|} \sum_{ij \in E_m} E_i \Delta t_j$$

normalize so that
 $g(128) = 1$



Deep learning HDR assembly

DigiVFX




LDR Images

Our Tonemapped HDR Image

Deep learning HDR assembly

DigiVFX



Kang (40.02 dB)


Sen (46.12 dB)

Ours (48.88 dB)

Ground Truth

Camera pipeline


DigiVFX



量化到8 bit


轉換成適合人看的色調

高光區過曝



ExpandNet

DigiVFX

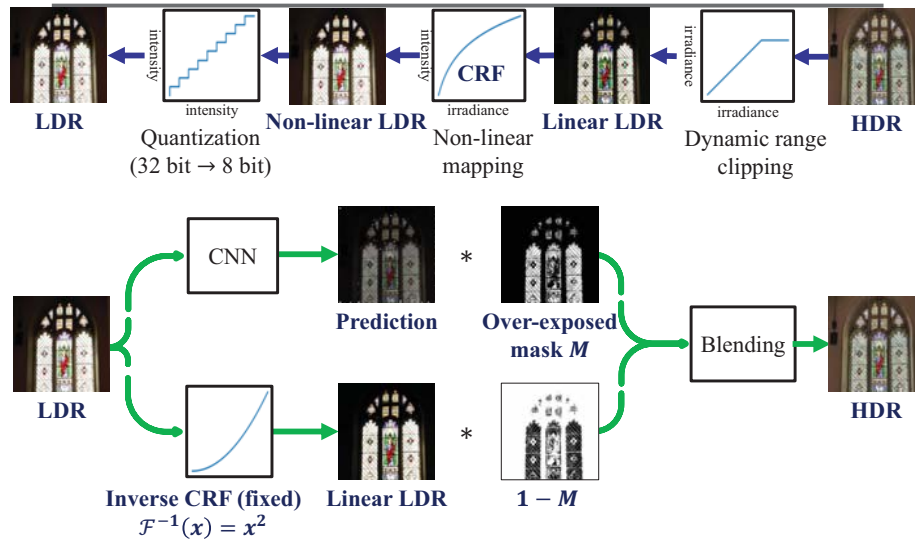


ExpandNet [Marnerides et al., Eurographics'18]

Implicitly learns everything, cannot generalize well

HDRCNN


DigiVFX



HDRCNN [Eilertsen et al., SIGGRAPH ASI/

HDRCNN

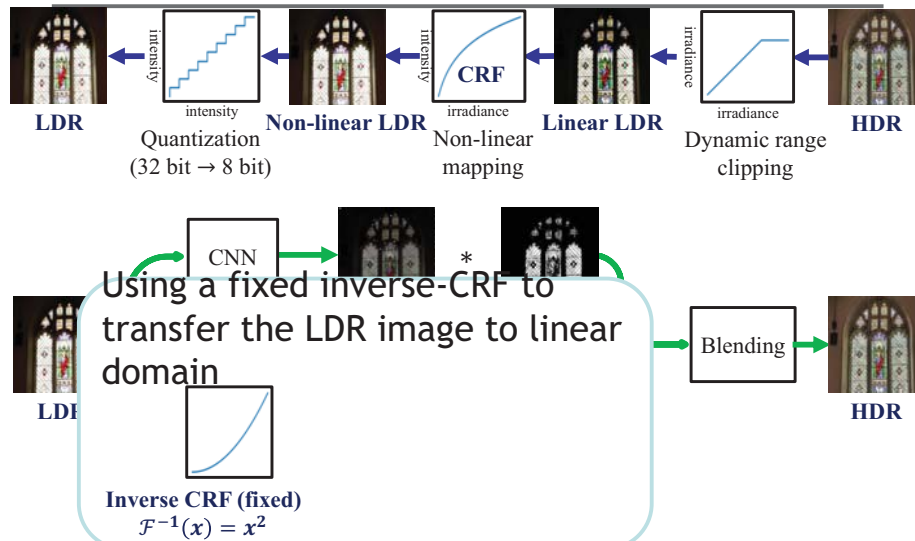
DigiVFX



HDRCNN [Eilertsen et al., SIGGRAPH ASI/

HDRCNN


DigiVFX



HDRCNN [Eilertsen et al., SIGGRAPH ASI/

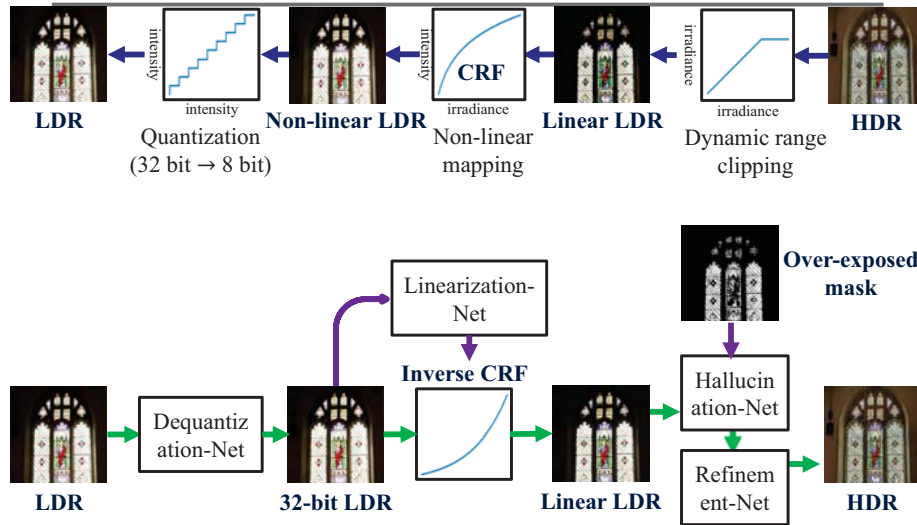
Our approach

DigiVFX



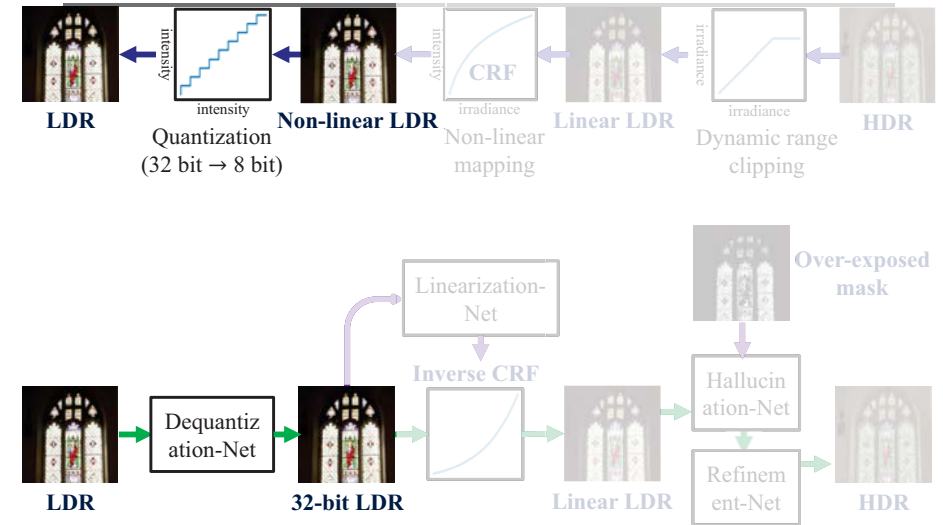
Our approach

DigiVFX




Dequantization-Net

DigiVFX



Dequantization-Net

DigiVFX




Dequantization-Net

DigiVFX



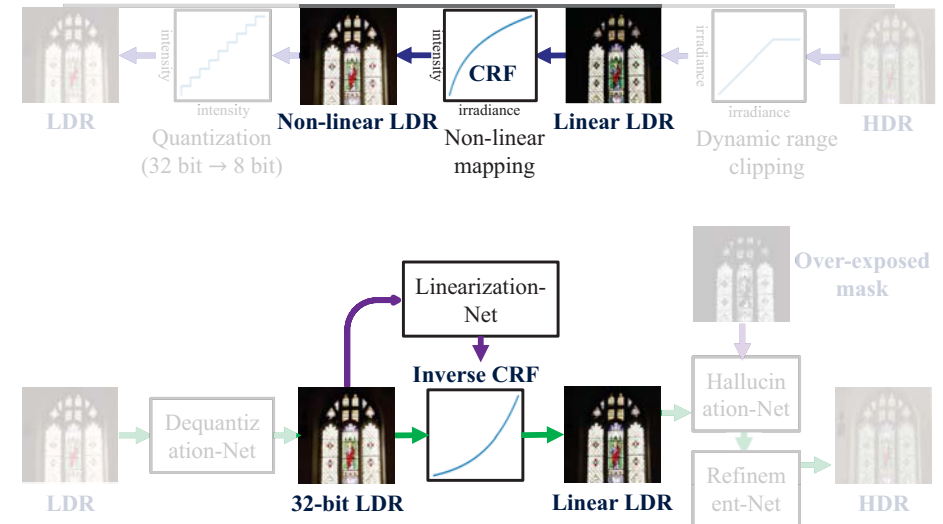
Dequantization-Net

DigiVFX



Linearization-Net

DigiVFX



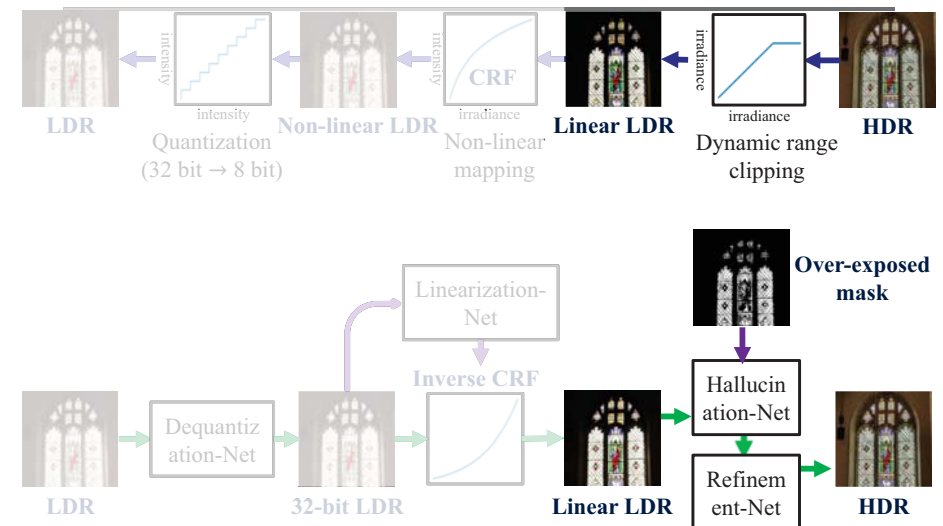
Linearization-Net




(a) Input

Hallucination-Net

DigiVFX



Hallucination-Net




Input

Ours

Ground-truth

Results

DigiVFX




input

label

Results

DigiVFX




input

ours

Results

DigiVFX




HDRCNN

ours

Results

DigiVFX



DrTMO

ours

Results

DigiVFX



ExpandNet

ours

Input

DigiVFX



Result

DigiVFX



Input

DigiVFX



Result

DigiVFX



Input

DigiVFX



Result

DigiVFX



HDR Video

- High Dynamic Range Video

Sing Bing Kang, Matthew Uyttendaele, Simon Winder, Richard Szeliski

SIGGRAPH 2003



HDR Video

High Dynamic Range Video


Sing Bing Kang
Matthew Uyttendaele
Simon Winder
Richard Szeliski

Microsoft Research, Redmond, WA

Assorted pixel



Assorted pixel



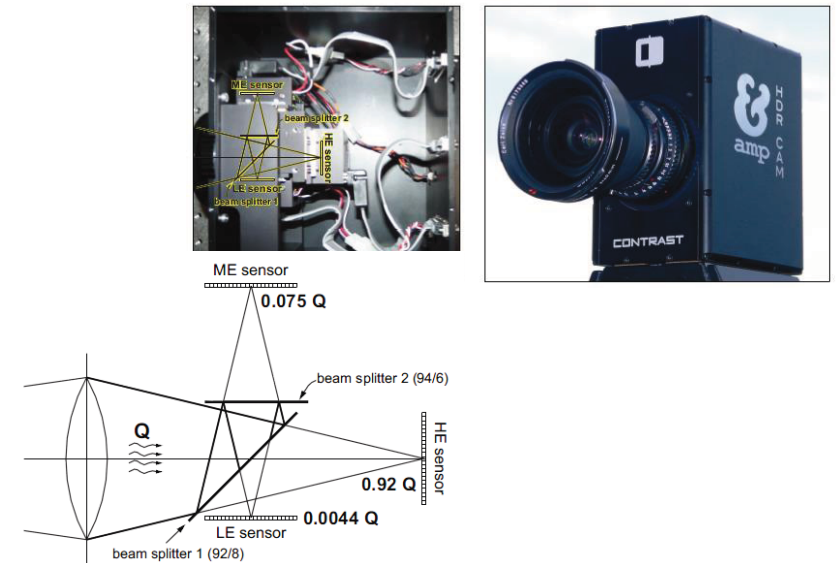
Assorted pixel

DigiVFX



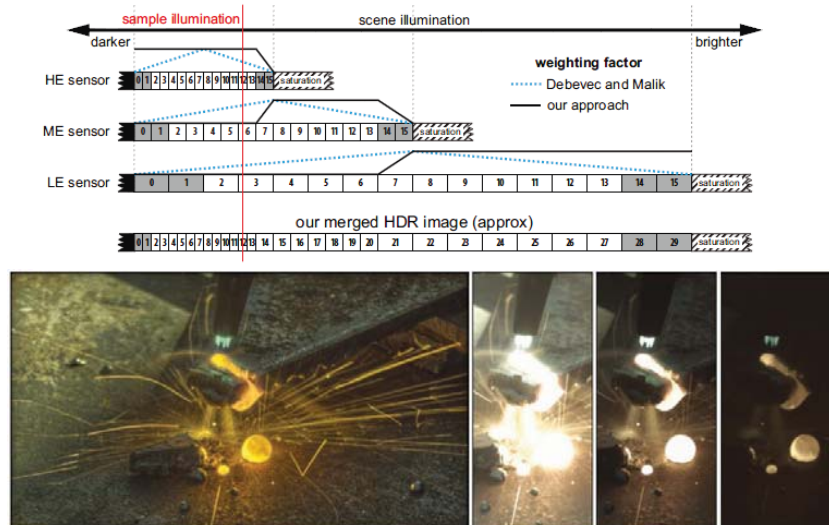
A Versatile HDR Video System

DigiVFX



A Versatile HDR Video System

DigiVFX



A Versatile HDR Video System

DigiVFX

A Versatile HDR Video Production System ACM SIGGRAPH 2011




HDR becomes common practice



- Many cameras has bracket exposure modes
- For example, since iPhone 4, iPhone has HDR option. But, it could be more exposure blending rather than true HDR.

References



References



- Paul E. Debevec, Jitendra Malik, [Recovering High Dynamic Range Radiance Maps from Photographs](#), SIGGRAPH 1997.
- Tomoo Mitsunaga, Shree Nayar, [Radiometric Self Calibration](#), CVPR 1999.
- Mark Robertson, Sean Borman, Robert Stevenson, [Estimation-Theoretic Approach to Dynamic Range Enhancement using Multiple Exposures](#), Journal of Electronic Imaging 2003.
- Michael Grossberg, Shree Nayar, [Determining the Camera Response from Images: What Is Knowable](#), PAMI 2003.
- Michael Grossberg, Shree Nayar, [Modeling the Space of Camera Response Functions](#), PAMI 2004.
- Srinivasa Narasimhan, Shree Nayar, [Enhancing Resolution Along Multiple Imaging Dimensions Using Assorted Pixels](#), PAMI 2005.
- G. Krawczyk, M. Goesele, H.-P. Seidel, [Photometric Calibration of High Dynamic Range Cameras](#), MPI Research Report 2005.
- G. Ward, [Fast Robust Image Registration for Compositing High Dynamic Range Photographs from Hand-held Exposures](#), jgt 2003.