Example-Based Color Image Manipulation and Enhancement

Tania Pouli
Erik Reinhard

Contact: tpouli@mpi-inf.mpg.de

SIGGRAPH2012

Introduction

SIGGRAPH2012

Color transfer was first heard of about 10 years ago, proposed as a technique to make rendered images look more natural by adjusting their color content on the basis of an example image. Now, color transfer is not a single algorithm but a range of methods and techniques that aim to make one image look more like another in terms of color content. Moreover, it turns out that there is utility far beyond these humble beginnings.

Introduction

SIGGRAPH2012

In this short course, we begin by giving an overview of the techniques that are currently available. We will then by means of many examples and comparisons on both images and video show when these algorithms are expected to produce their best results. We will also show how to choose appropriate examples to steer the results and show applications of color transfer that include making night-time images from day-time images, color correcting stereo pairs, and color matching photographs as a pre-processing to panorama stitching.

Introduction

SIGGRAPH2012

This course will allow researchers in this area to understand where there may be further opportunities for algorithmic improvements, and it would allow practitioners in creative industries, which include photography, movies and games, to understand how to make the most of these algorithms.

Audience

This course is for anyone who likes messing with images. More specifically, the course is intended for anyone creating or editing 2D content, including technical directors, artists, photographers, game designers and movie makers. We also aim this course at researchers in image processing who may spot opportunities to develop further techniques or find utility of the work presented here to enhance their own algorithms. We will provide a good balance between explaining the algorithmic details and giving practical hands-on experience.

Tania Pouli

Tania Pouli is a post-doc at the Max Planck Institute for Informatics. She received a B.Sc. in from the University of Bath in 2003 and a Ph.D. from the University of Bristol in 2011. She was also a post-doctoral researcher at the University of Bristol working on psychophysical experiments related to mobile projectors. Her current research focuses on statistical analysis of high dynamic range images, resampling algorithms, novel image-based editing applications as well as color imaging in general and color transfer specifically.

Erik Reinhard

SIGGRAPH2012

Erik Reinhard has been active in computer graphics research since 1993, and in high dynamic range imaging since 2001. He founded ACM Transactions on Applied Perception, and was Editor-in-Chief between 2003 and 2009. Erik is lead author of `High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting' and 'Color Imaging: Fundamentals and Applications'. He was program co-chair for the Eurographics Rendering Symposium 2011. He was keynote speaker for Eurographics 2010, the Computational Color Imaging Workshop 2011, and CGIV 2012. Finally, he was one of the first to pioneer color transfer in 2001.

Syllabus

SIGGRAPH2012

- Introduction - Pouli, 10 mins
- Example methods - Pouli, 20 mins
- User control - Pouli, 10 mins
- Color Space Statistics - Reinhard, 15 mins
- Applications - Reinhard, 15 mins
- Discussion - Both, 5 mins

Introduction

Introduction

- What is color transfer?
- What can you do with it?
- Per-channel and 3D methods
- Some examples

Appearance Transfer

- Most image manipulation happens by parameter specification or dragging sliders

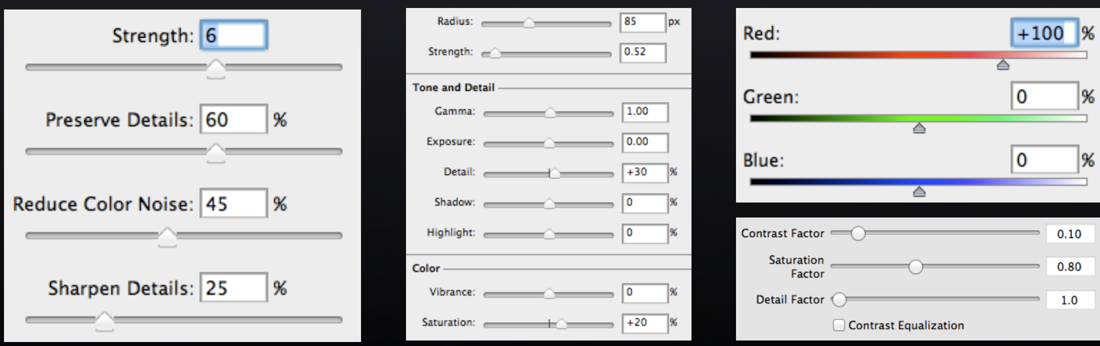

- Requires knowledge of the effect of each parameter

Appearance Transfer

SIGGRAPH2012

- A more intuitive approach would be to **transfer properties** from another image
- E.g. color, texture, painterly styles etc.
- A lot of research over the last decade has focused on transferring **color** between images

Color Transfer

SIGGRAPH2012

- Basic idea:
  - Find an image with desired colors (target)
  - Pick some way of describing the color characteristics of the original (source) image and the target
  - Manipulate the source color descriptor so it approximates that of the target

Colour Transfer

SIGGRAPH2012

Source

Target

Result

Color Transfer

SIGGRAPH2012

Key issue

- Basic idea:
  - Find an image with desired colors (target)
  - Pick some way of describing the color characteristics of the original (source) image and the target
  - Manipulate the source color descriptor so it approximates that of the target

Color Transfer

SIGGRAPH2012

f( , ) = 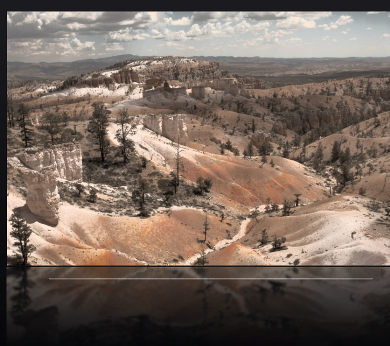

Color Transfer

How do we design a good color descriptor?

Colour Transfer

Constraints:

- The colour palette of the target should be transferred to the source
- No artefacts

Freedom:

- No ground truth

1D vs 3D

SIGGRAPH2012

3 x 1D problems

Each colour channel is manipulated separately

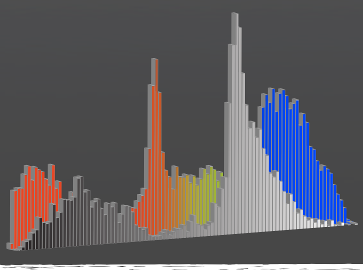

1 x 3D problem

The image is treated as a 3D dataset

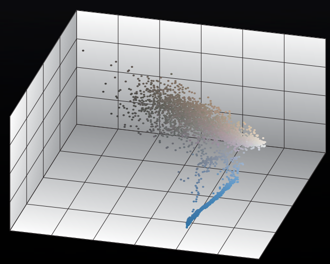

1D vs 3D

SIGGRAPH2012

3 x 1D problems

Each colour channel is manipulated separately

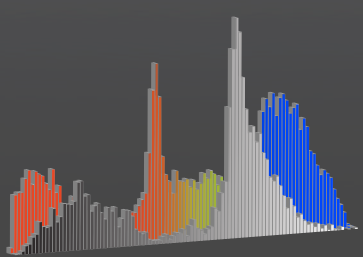

- 3 easier problems BUT...
- The choice of colour space matters
- Channel cross-talk

1D vs 3D

SIGGRAPH2012

- A lot more data
- More complex algorithms
- Often requires optimization to reduce artefacts

1 x 3D problem

The image is treated as a 3D dataset

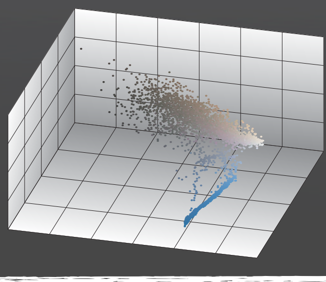

Example Methods

Per-Channel Methods

SIGGRAPH2012

- Color spaces / PCA / Decorrelation
- Means and standard deviation
- Histogram matching
- Histogram reshaping
- High dynamic range color transfer

Previous work

SIGGRAPH2012

3x1D problems

Reinhard et al. 2001

Xiao & Ma 2009

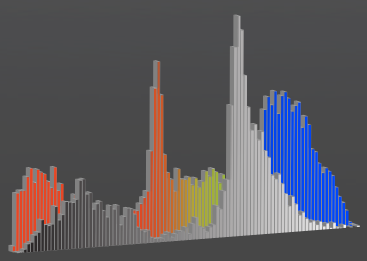

Previous work

SIGGRAPH2012

3x1D problems

Reinhard et al. 2001 → Linear shifting and scaling

Xiao & Ma 2009 → Full histogram matching

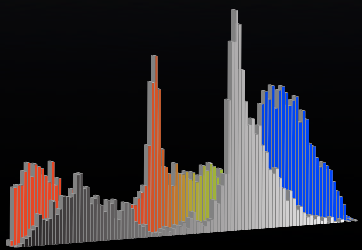

Color Transfer in Lαβ

SIGGRAPH2012

- Convert source I_s and target I_t to Lαβ
- Linear shifting and scaling using mean and standard deviation:

$$I' = I_s - \mu_s$$

$$I'' = \frac{\sigma_t}{\sigma_s} I'$$

$$I_o = I'' + \mu_t$$

- Each channel manipulated separately

Reinhard et al, 'Color Transfer between Images', IEEE Computer Graphics and Applications 21(5), pp. 34-41, 2001

Some Results

SIGGRAPH2012

Some Results

SIGGRAPH2012

Some Results

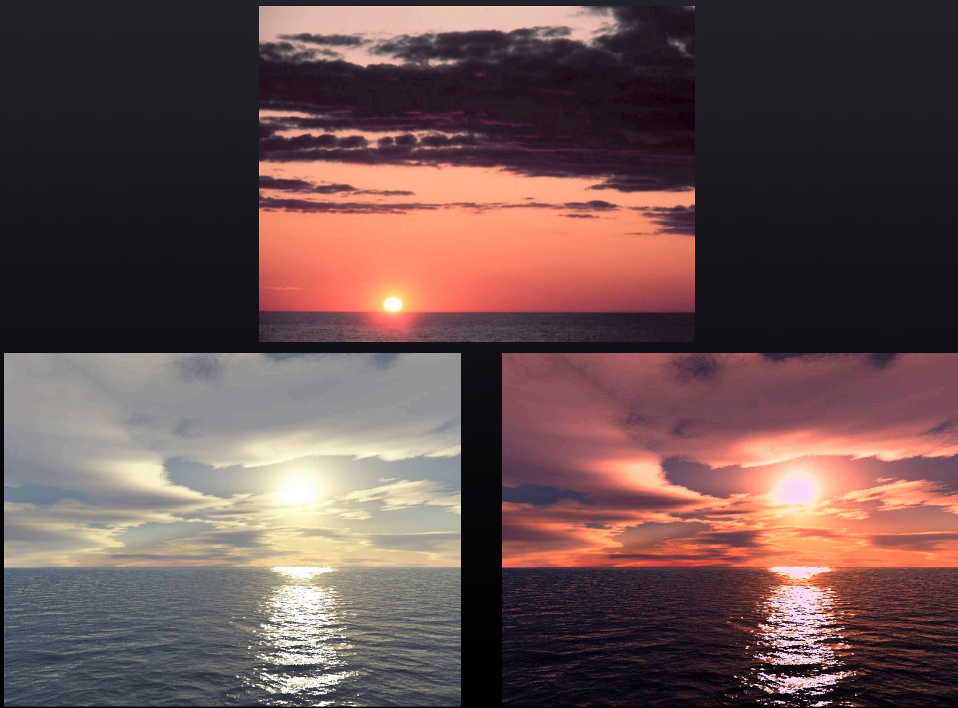

Colour Transfer in Lαβ

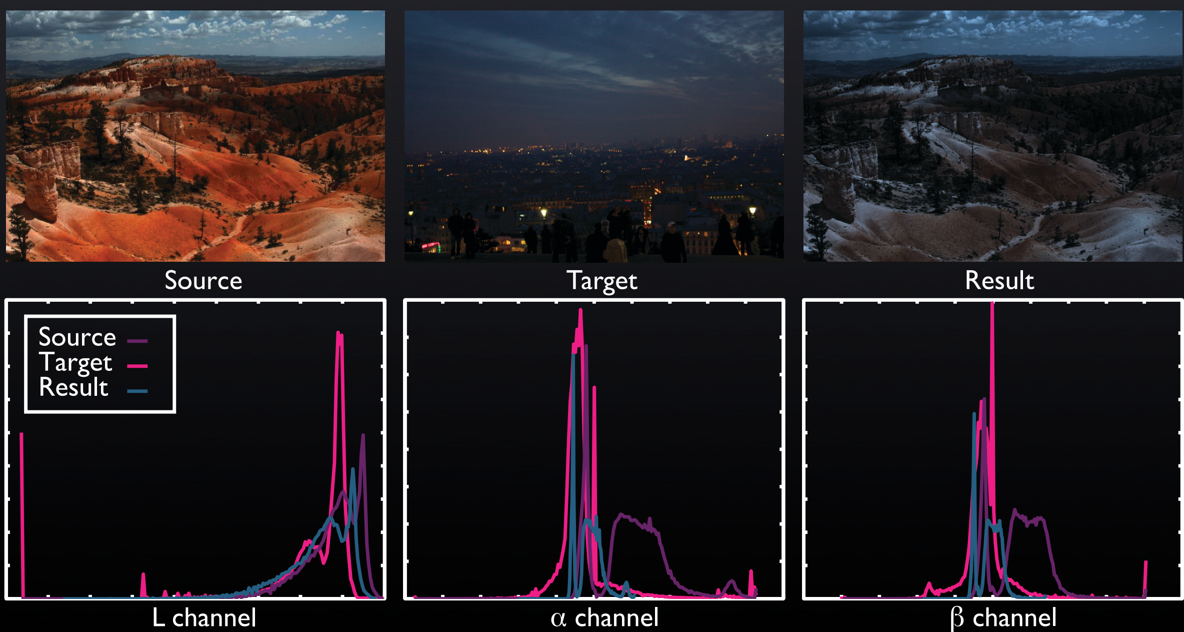

Colour Transfer in Lαβ

- If source and target are very different the transfer may lead to unpredictable results
- No control on amount of matching

Histogram Reshaping

- Histogram scale-space approach
- Colours are matched by manipulating histograms at **different scales** achieving **progressive matches**
- This is done in **CIELab** colour space

Tania Pouli, Erik Reinhard, 'Progressive Color Transfer for Images of Arbitrary Dynamic Range', Computers and Graphics 35(1), pp. 67-80, 2011

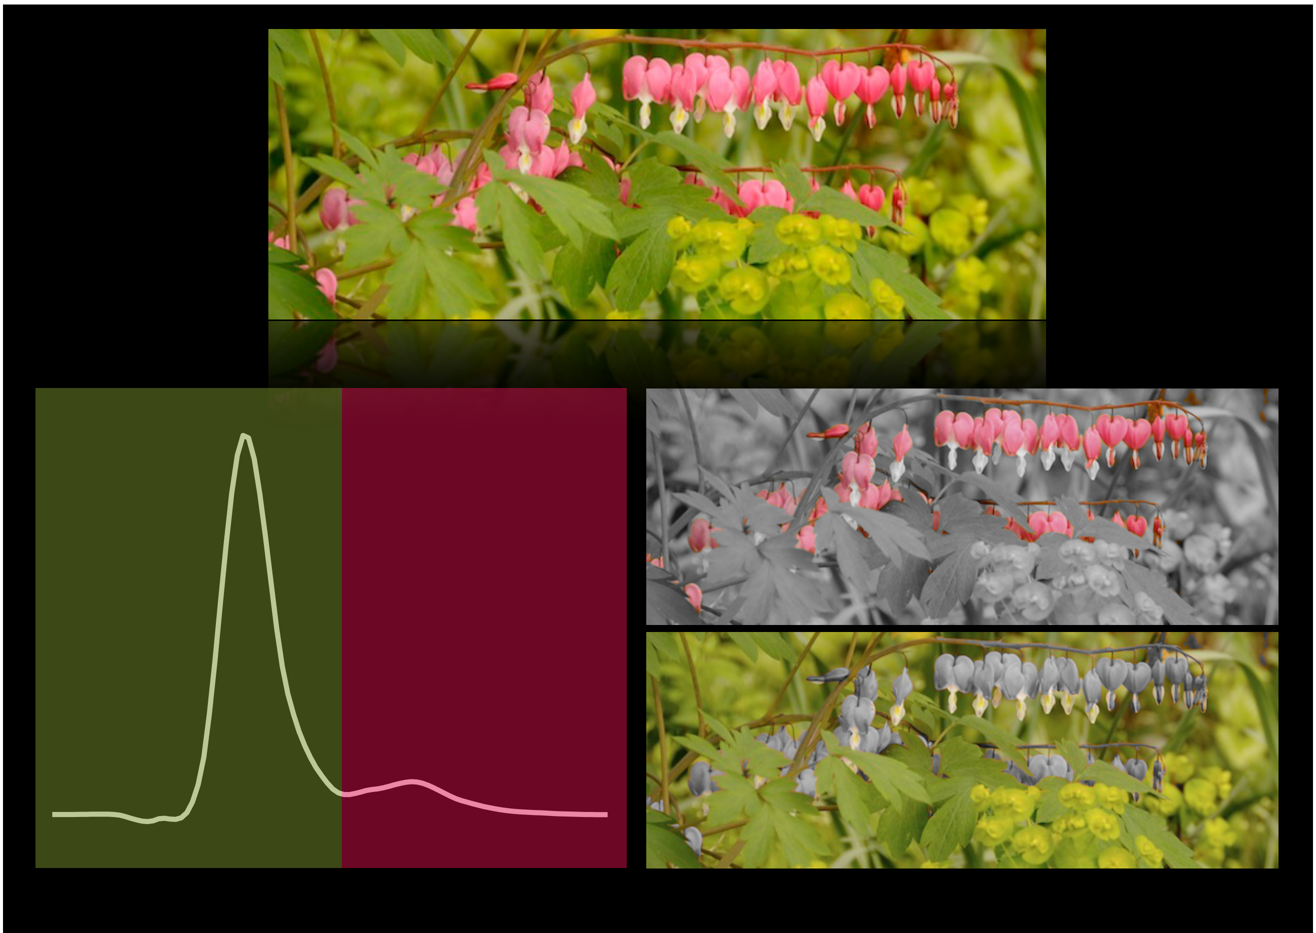

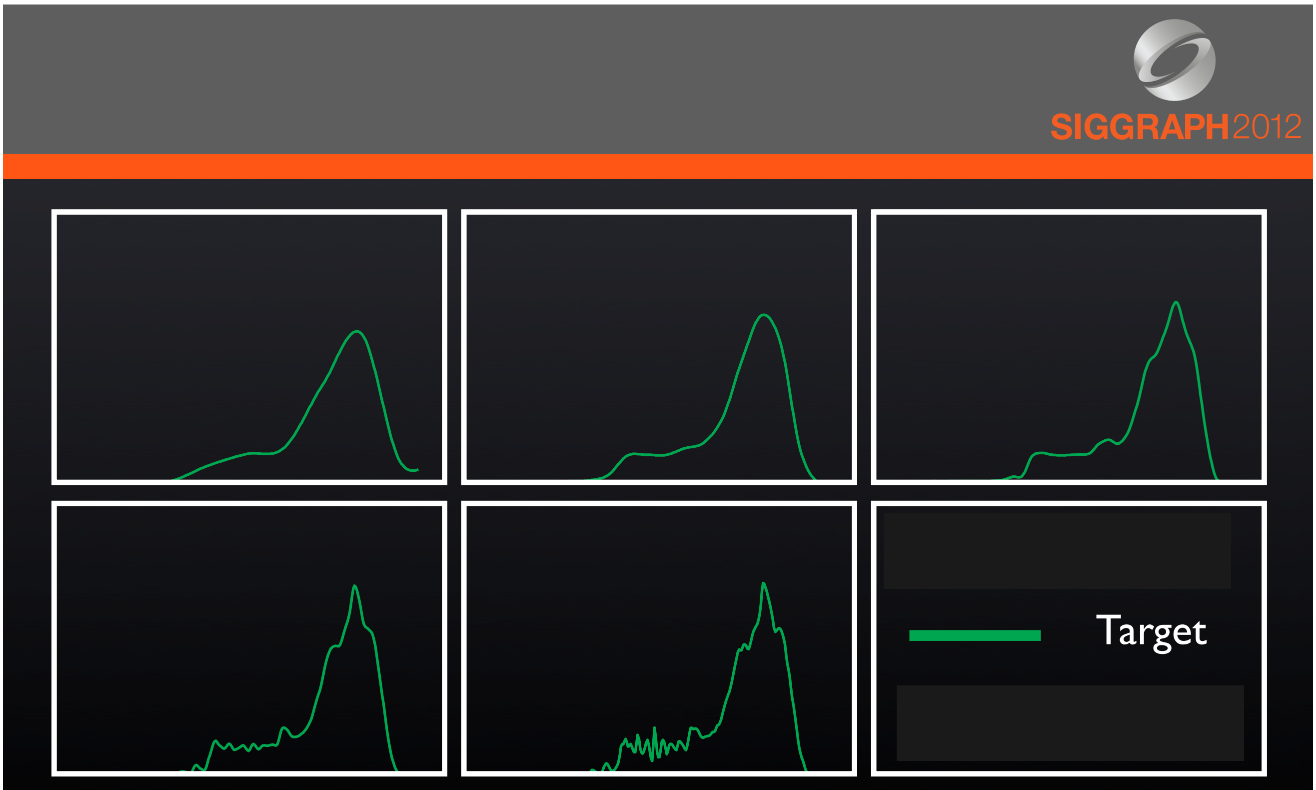
SIGGRAPH2012
Target

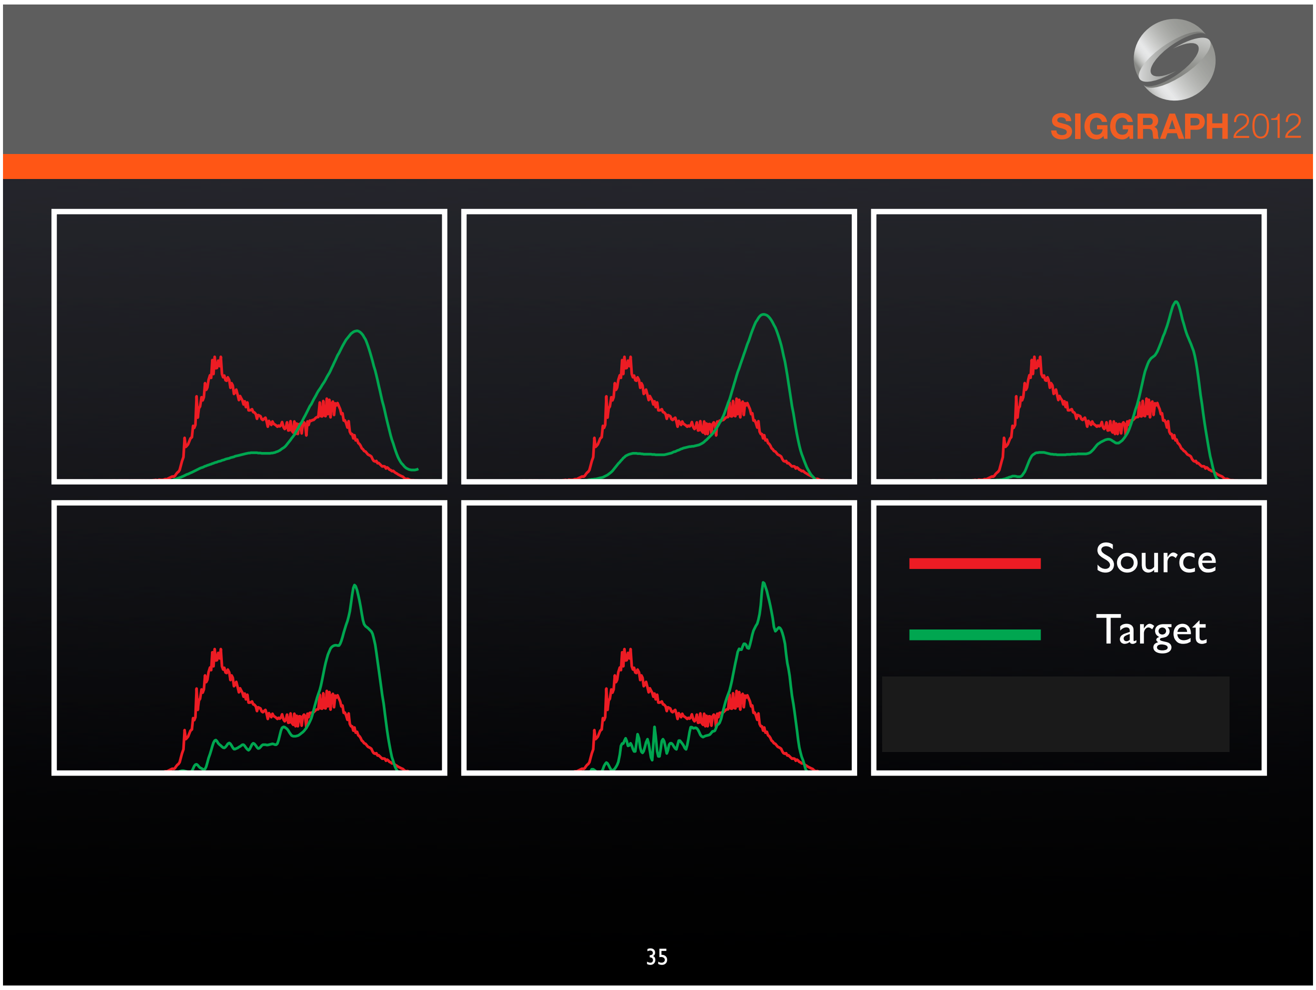
SIGGRAPH2012
Source
Target
35

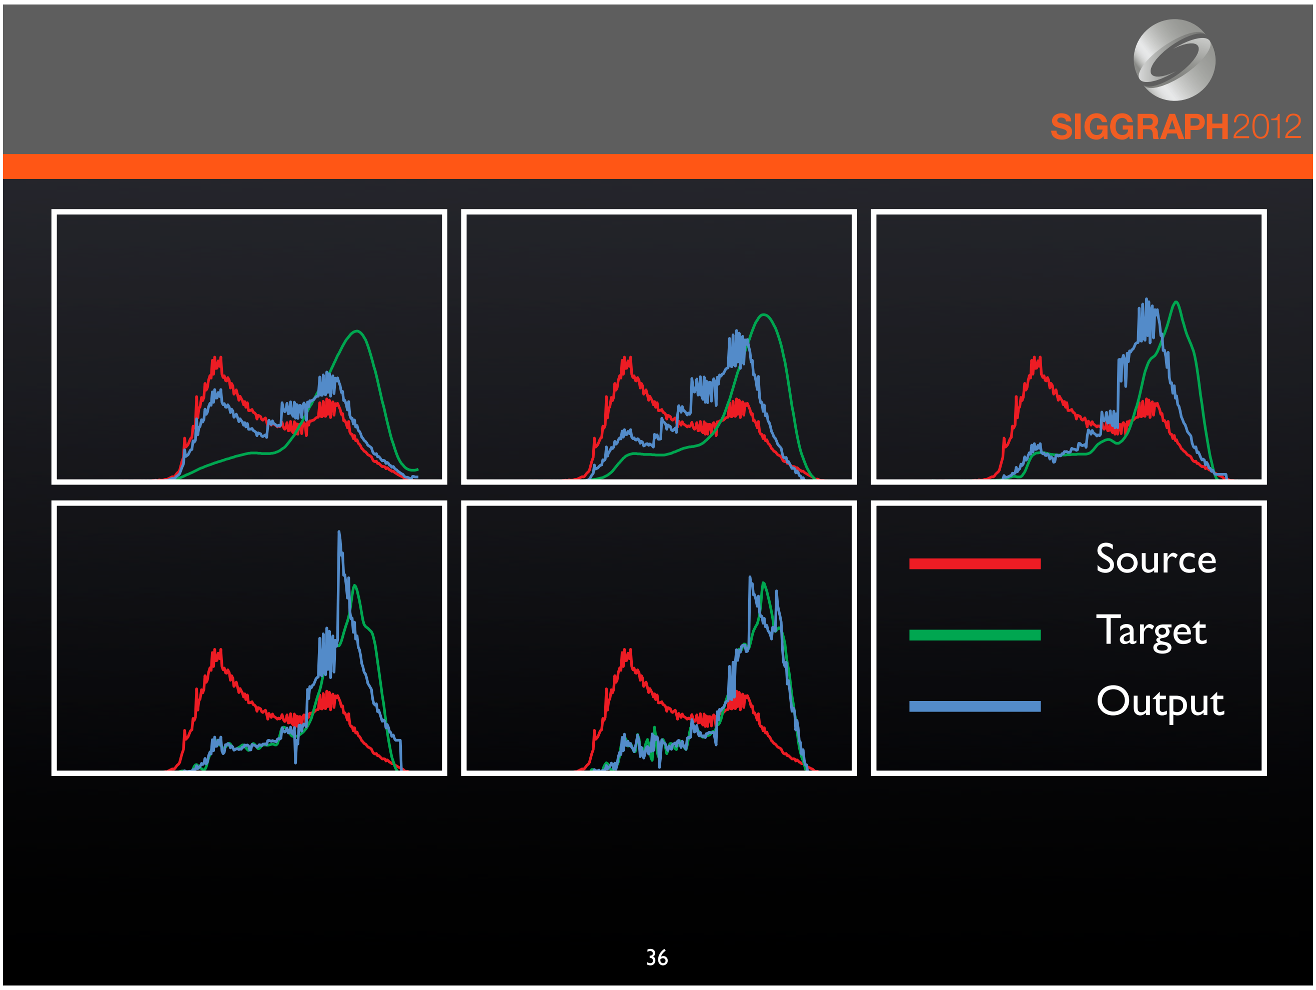
SIGGRAPH2012
Source
Target
Output
36

Histogram Reshaping

SIGGRAPH2012

$\mu_{s,1}$ $\mu_{s,2}$ $\mu_{s,3}$

$\sigma_{s,1}$ $\sigma_{s,2}$ $\sigma_{s,3}$

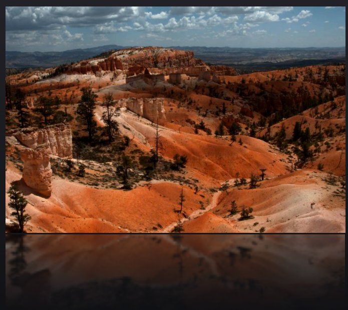

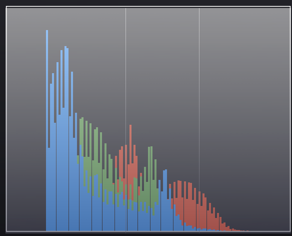

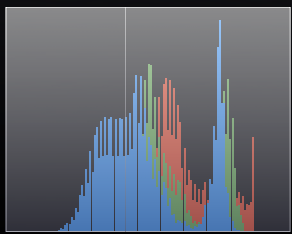

$\mu_{t,1}$ $\mu_{t,2}$ $\mu_{t,3}$

$\sigma_{t,1}$ $\sigma_{t,2}$ $\sigma_{t,3}$

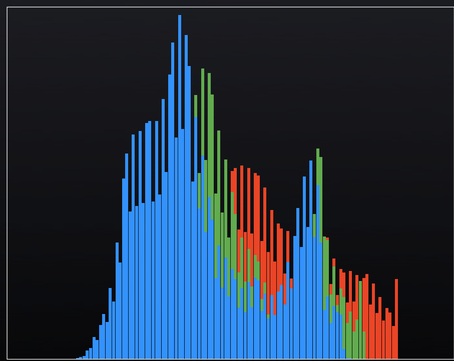

$$h_{o,i} = (h_{s,i} - \mu_{s,i})\frac{\sigma_{t,i}}{\sigma_{s,i}} + \mu_{t,i}$$

Histogram Reshaping

SIGGRAPH2012

Weights vs Scales

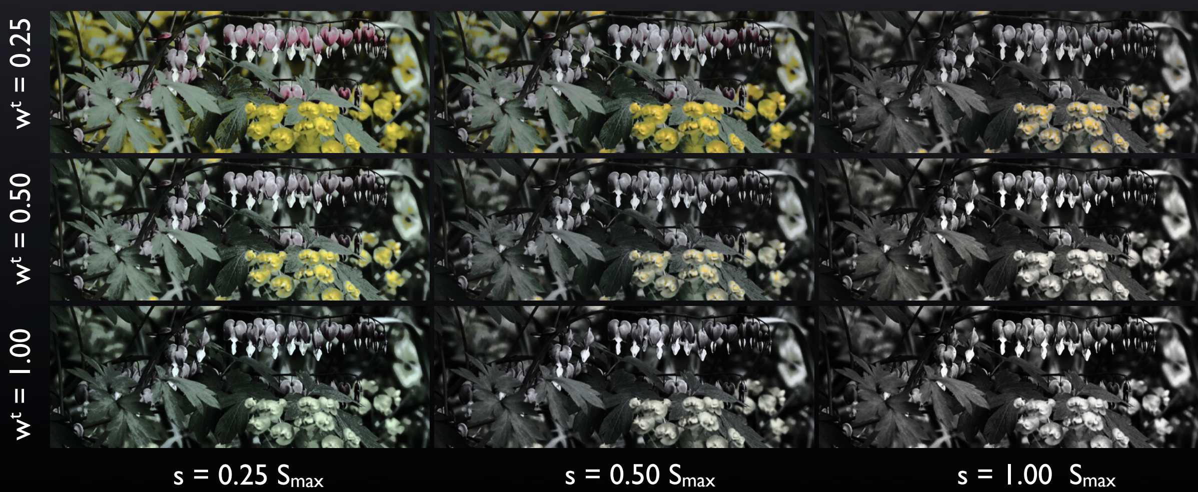

Weights vs Scales

SIGGRAPH2012

Weights vs Scales

SIGGRAPH2012

Weights vs Scales

Histogram Reshaping

Histogram Reshaping

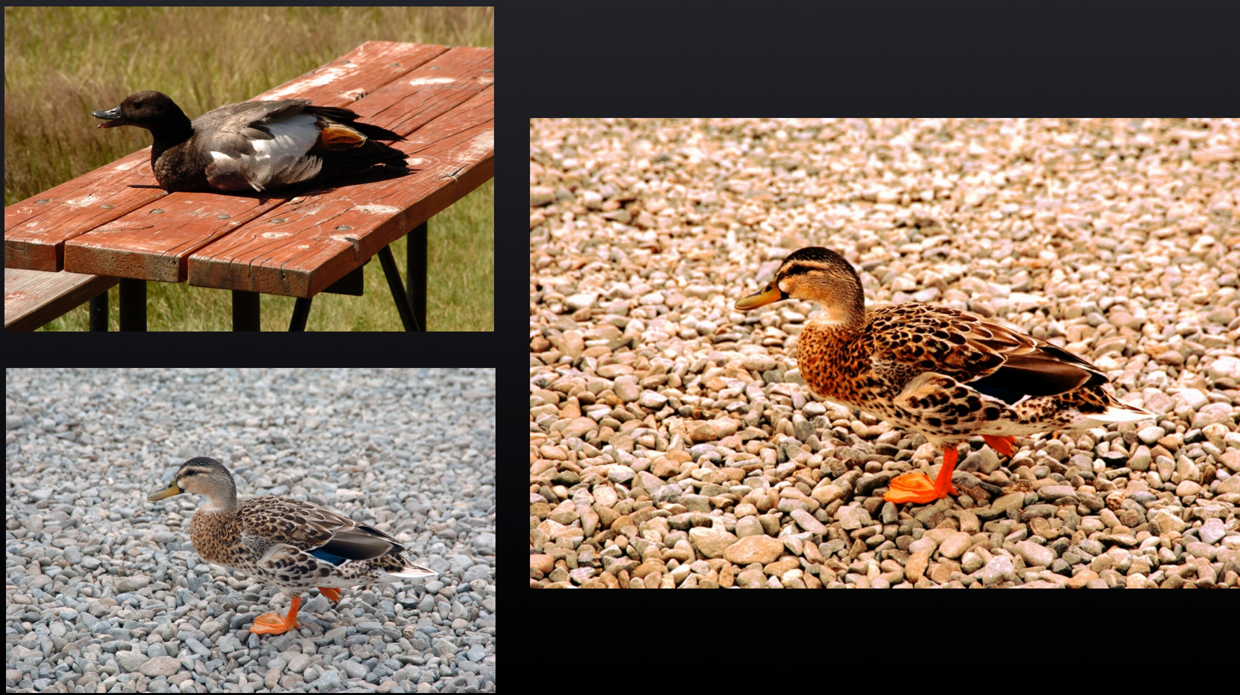

3D Transfers

- 1D approaches are simpler and often faster
- Not always robust
- The relations between channels might carry useful information
- We can treat the image as a 3D point cloud instead

3D Histogram Matching

SIGGRAPH2012

- In 1D, a histogram can be matched with another by inverting the cumulative distribution function

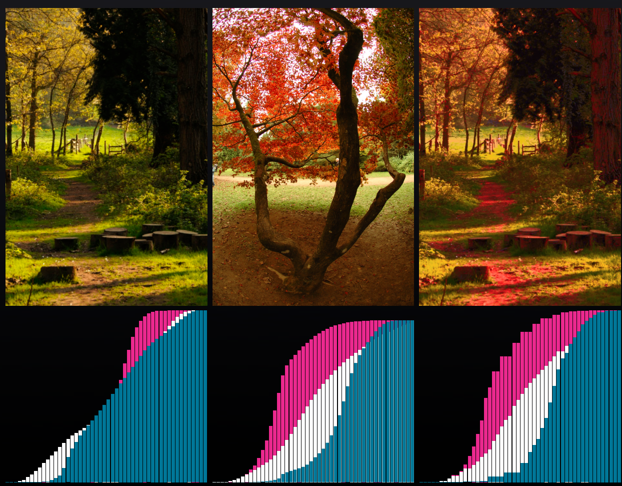

3D Histogram Matching

SIGGRAPH2012

- A 3D cdf can't be directly inverted so simplifications are necessary

ND-Distribution Transfer

- One solution is to iteratively match a series of 1D projections (marginals) of the original 3D distribution
- As the distribution is progressively shaped, the 1-D marginals are updated, eventually matching the target distribution

ND-Distribution Transfer

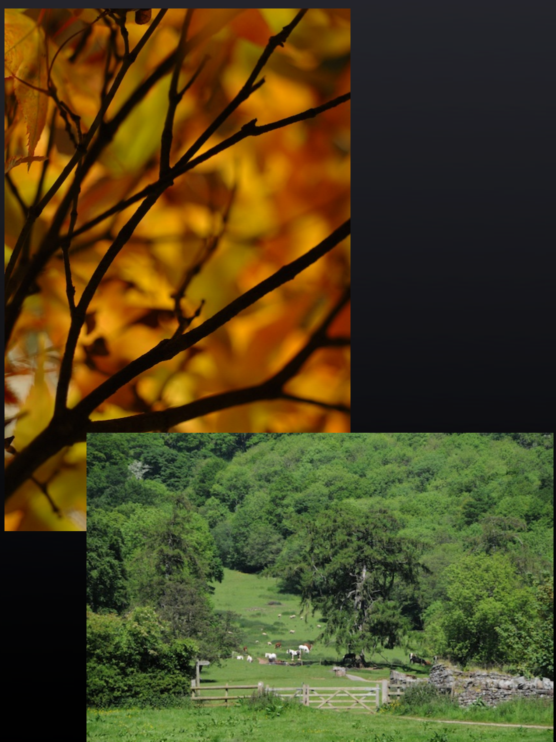

Color Transfer Metric

SIGGRAPH2012

- Question: How do we **quantitatively** evaluate a color transfer result?
- So far one available solution
  - Color transfer can be defined as a tradeoff between achieving a **color distribution** close to the target
  - and maintaining the **gradient distribution** of the source

Color Transfer Metric

SIGGRAPH2012

By extending the mean squared error definition, one possible color transfer 'success' metric can be formulated as:

$$\text{Error} = \text{MSE}\,(H_o - H_t) + \text{MSE}\,(G_o - G_s)$$

H : Color Histogram
G: Gradient Histogram

Xiao & Ma, Gradient Preserving Color Transfer, Pacific Graphics 2009

Color Transfer Metric

SIGGRAPH2012

- This metric is suitable in some cases
- Quantitative evaluation of a color transfer result is still an open problem though

User Control

User Control

SIGGRAPH2012

- All methods discussed are automatic
- The algorithm decides by itself which color should map to which
- In many cases more control is desirable

53

Masks

SIGGRAPH2012

54

Masks

SIGGRAPH2012

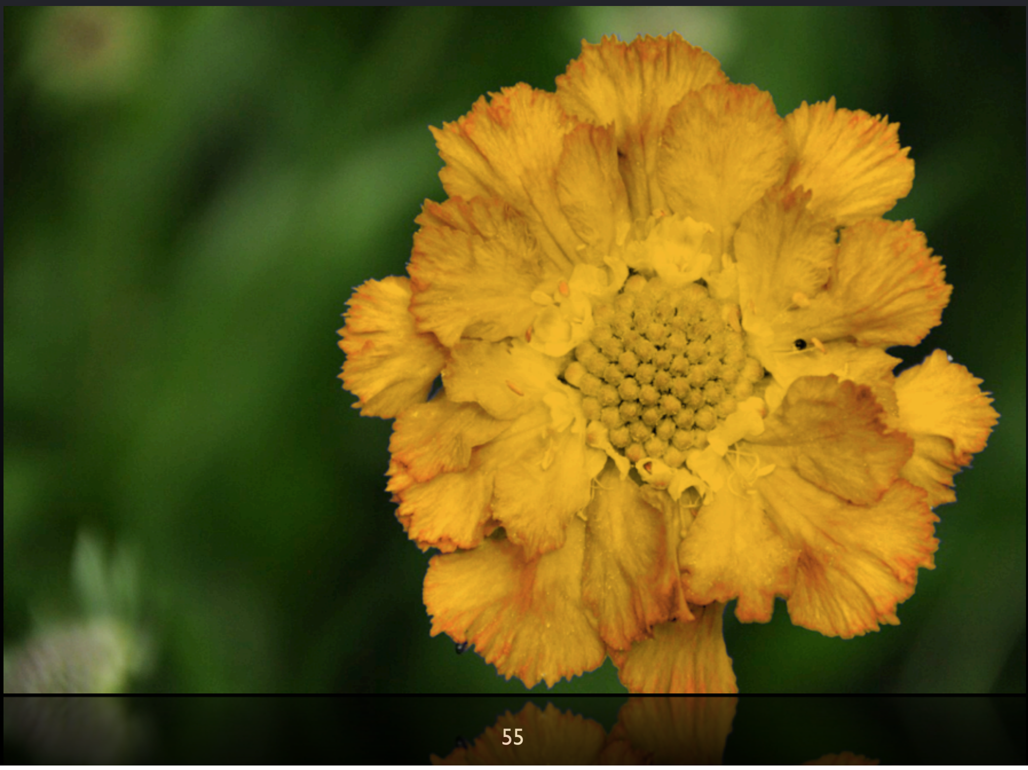

Swatches

SIGGRAPH2012

- Select matching image regions and compute means and standard deviations in each
- In color space this gives a cluster per region
- Each pixel in the image now has a distance to each of the clusters

Swatches

SIGGRAPH2012

- To transfer the color of a pixel, the new pixel assignment takes distance to each cluster in color space as well as the size of the cluster into account
- Size of clusters is given by standard deviations

Swatches

SIGGRAPH2012

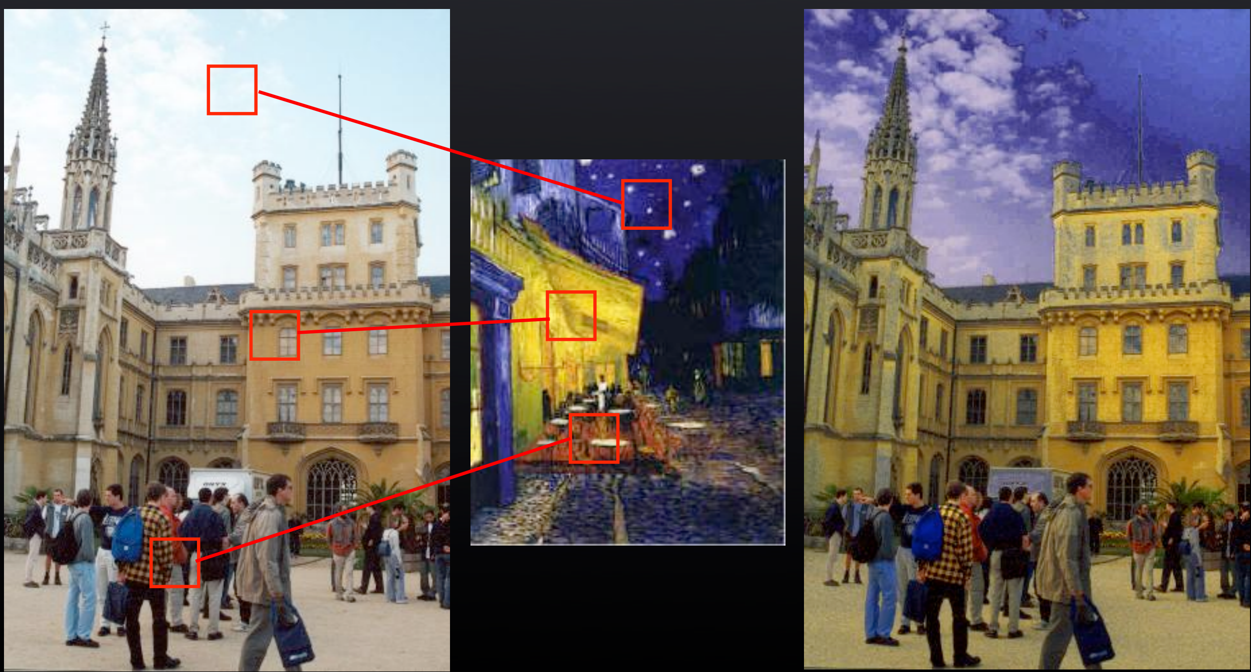

Strokes

SIGGRAPH2012

- The idea of swatches can easily be extended to strokes indicating correspondence between source and target

Strokes

SIGGRAPH2012

- Strokes can also be used to colorize (or recolorize) images
- Based on the premise that nearby pixels are likely to have similar colors
- Given a grayscale image marked with strokes, the colors are propagated through the image using optimization

A. Levin D. Lischinski and Y. Weiss
Colorization using Optimization.
SIGGRAPH, ACM Transactions on Graphics, 2004.

Strokes

SIGGRAPH2012

Greyscale input

User generated strokes

A. Levin D. Lischinski and Y. Weiss
Colorization using Optimization.
SIGGRAPH, ACM Transactions on Graphics, 2004.

Strokes

SIGGRAPH2012

Resulting colorization

Ground truth

A. Levin D. Lischinski and Y. Weiss
Colorization using Optimization.
SIGGRAPH, ACM Transactions on Graphics, 2004.

Color Space Statistics

Color Space Issues

SIGGRAPH2012

- If we treat each channel separately, choice of color space matters
- If the three channels are correlated, values in one channel are good predictors of the other two
- So if we change the values of the one channel, we might affect the others as well

Choice of Colour Space

SIGGRAPH2012

Input images | Output images

RGB

Lαβ

Red - Green

SIGGRAPH2012

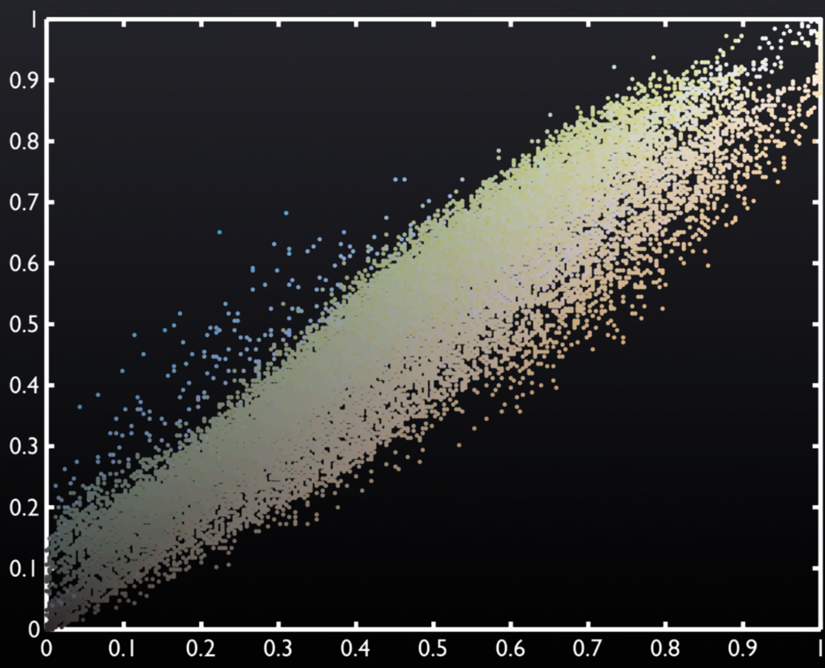

Red vs Green

Red - Blue

SIGGRAPH2012

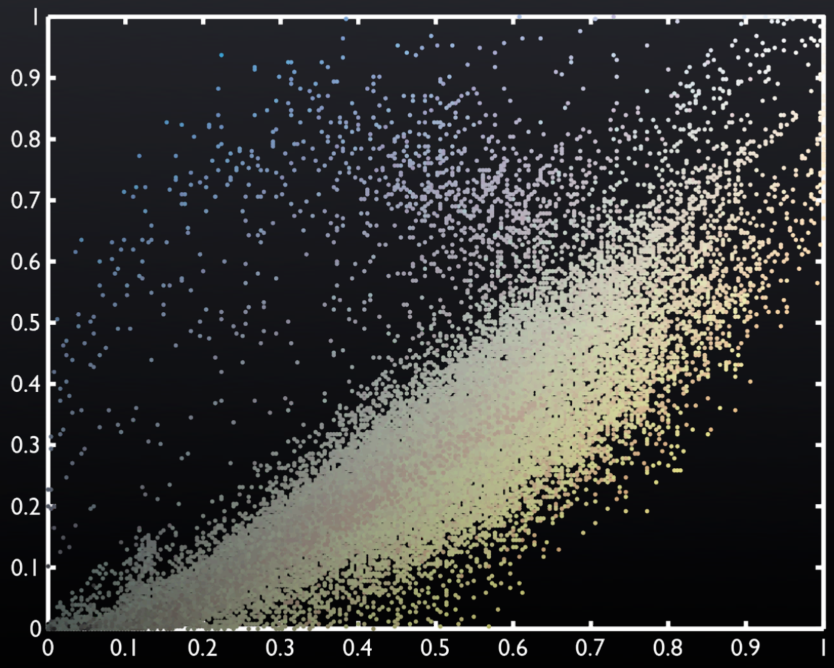

Red vs Blue

Green - Blue

SIGGRAPH2012

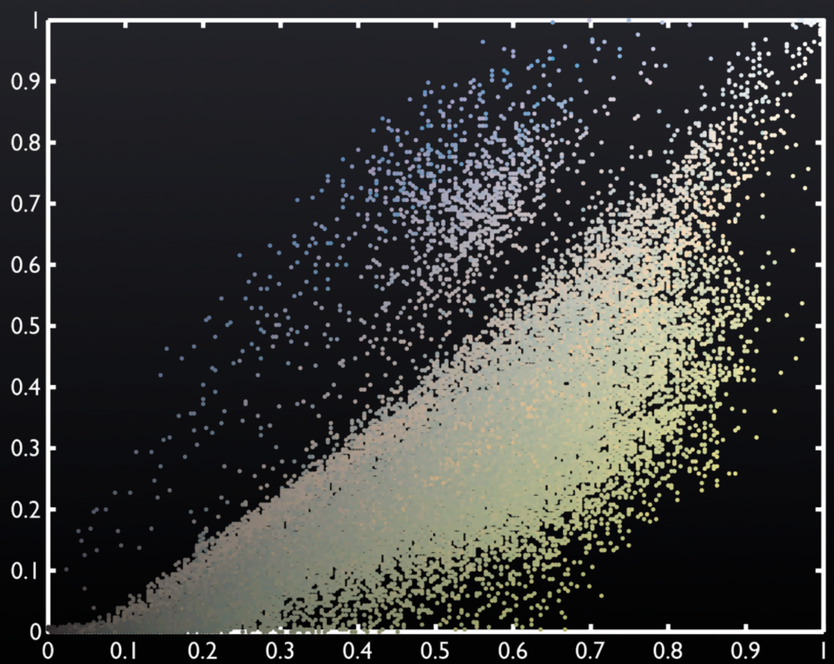

Green vs Blue

Channel Correlation

SIGGRAPH2012

- To manipulate each channel separately without artefacts we need to reduce correlation/dependence...

Image Statistics

SIGGRAPH2012

- Take spectral images of natural scenes
- Convert to LMS
- Run Principal Components Analysis (PCA)
- Analyse axes

- Result: a colour opponent space!

Opponency in the HVS

SIGGRAPH2012

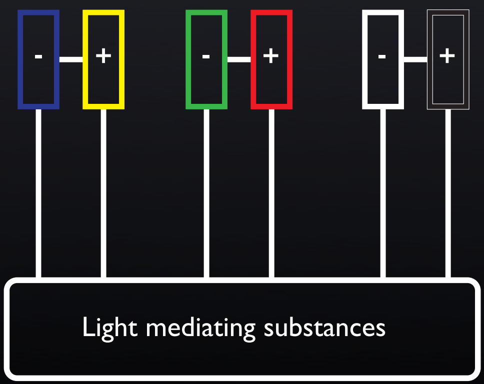

- But this is what the ganglion cells transmit:
- Signal from the cones is recombined in pairs of **opponent colours**
  - Light / Dark - L
  - Red / Green - α
  - Blue / Yellow - β

Laβ colour space

SIGGRAPH2012

- This opponency is particularly good at decorrelating natural scenes
- **Ruderman et al.** (1998) derived the Laβ colour space by decomposing natural images using **Principal Component Analysis** (PCA)

L - α

SIGGRAPH2012

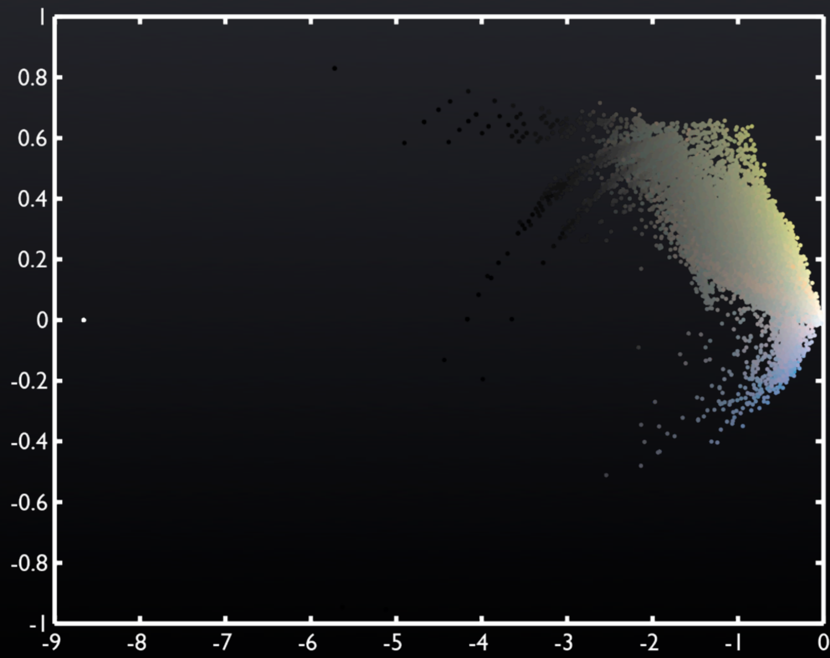

L - β

SIGGRAPH2012

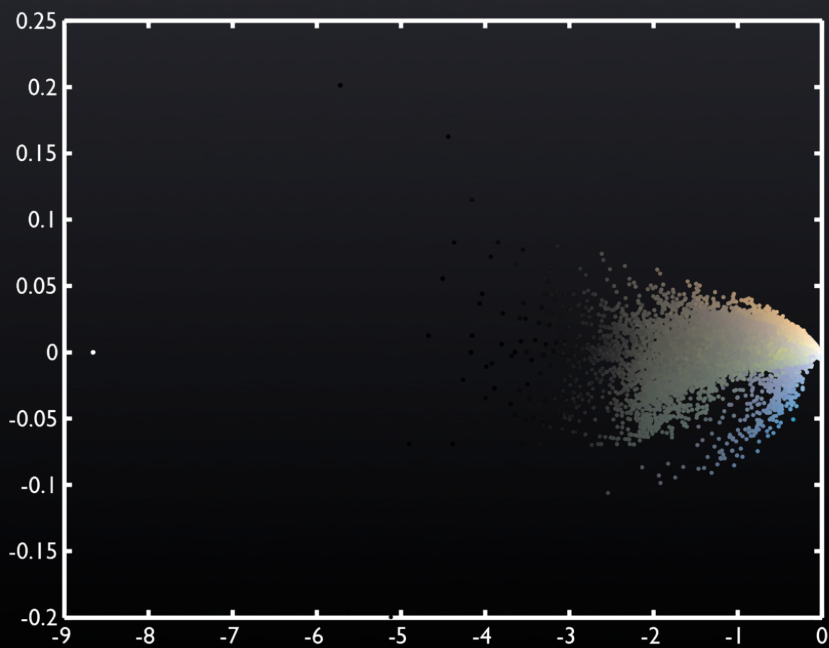

α - β

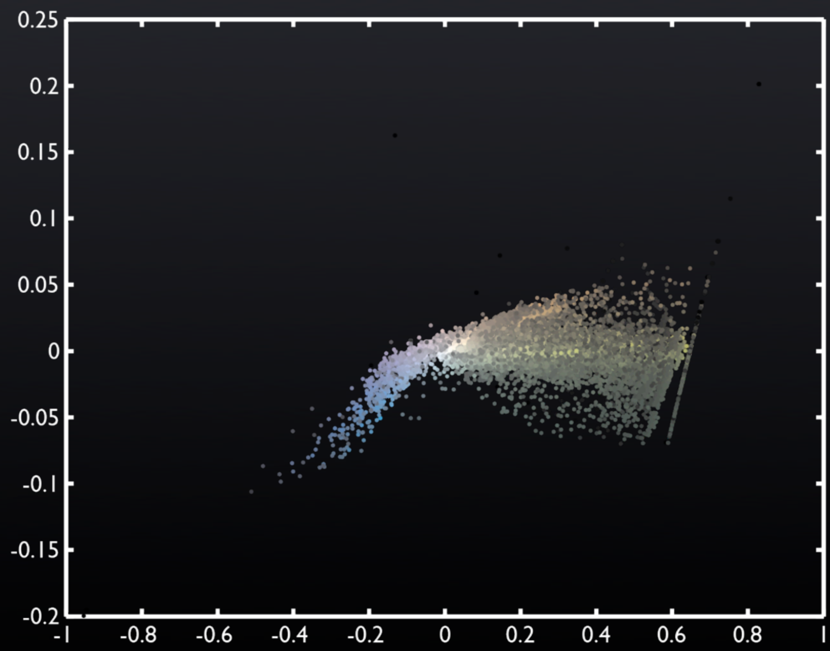

Color Space Choices

SIGGRAPH2012

- Linear shifting and scaling → Lαβ
- Histogram reshaping → CIELAB

- Many other colour spaces have been used in colour transfer

Color Space Choices

SIGGRAPH2012

- Xiao & Ma, 'Gradient-preserving color transfer': **Lαβ**
- Wen et al., 'Example-based multiple local color transfer by strokes': **CIELab**
- Neumann & Neumann, 'Color style transfer techniques using hue, lightness and saturation histogram matching': **LCh***
- Levin et al., 'Colorization using optimization': **Yuv**

Color Space Choices

SIGGRAPH2012

- Many colour transfer algorithms compute **custom colour** spaces for the given input images
- Using **PCA** a set of rotation axes can be computed
- Theoretically should **maximally decorrelate** the input

Time for a Competition

SIGGRAPH2012

So which is the best colour space?

Colour Statistics

SIGGRAPH2012

Overview

- Different scenes types
- Comparison of colour spaces used in colour transfer
- Scene-specific and image specific spaces

Image Sets

SIGGRAPH2012

- We looked at four scene types:

Indoors (IN)

Outdoors (MD)

Night (MN)

Natural (ND)

Colour Statistics

SIGGRAPH2012

- How decorrelated are colour spaces?
- **Covariance** between pairs of channels is a good measure for that
- $$\mathrm{Cov}(I,J) = \sum_{p=1}^{N} \frac{(I(p)-\mu_I)(J(p)-\mu_J)}{N-1}$$

Covariance Results

SIGGRAPH2012

Covariance for different colour spaces for the four scene types (lower is better):

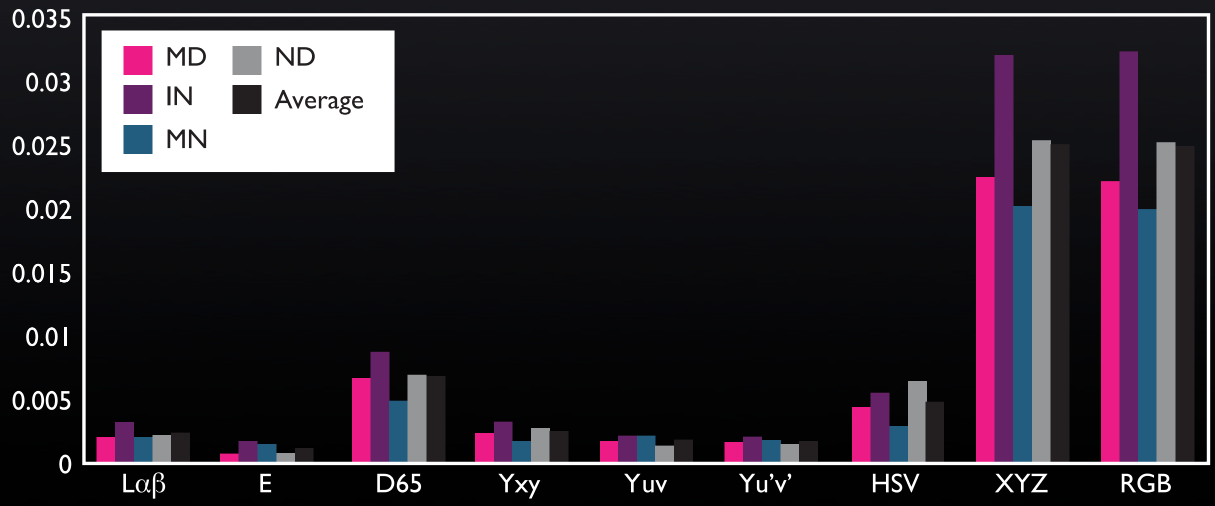

Covariance Ranking

SIGGRAPH2012

Rank	MD	IN	MN	ND
1	CIELab (E)	CIELab (E)	CIELab (E)	CIELab (E)
2	Yu'v'	Yu'v'	Yxy	Yu'v'
3	Yuv	Yuv	Yu'v'	Yuv
4	lab	lab	lab	lab
5	Yxy	Yxy	Yuv	Yxy
6	HSV	HSV	HSV	HSV
7	CIELab (D65)	CIELab (D65)	CIELab (D65)	CIELab (D65)
8	RGB	RGB	RGB	RGB
9	XYZ	XYZ	XYZ	XYZ

Discussion

SIGGRAPH2012

- Results highly consistent across datasets
- CIELAB (E) shows least covariance
- What about PCA-based colour spaces?

PCA-Based Spaces

SIGGRAPH2012

With four datasets, we could compute four dedicated colour spaces using PCA.

We expect lower co-variance between channels, and better performance in colour transfer.

PCA-Based Spaces

SIGGRAPH2012

We use a process similar to Ruderman et al. to compute dedicated spaces for each image category

1. Convert images to LMS cone space
2. NxN patch selected from each image
3. Patches are log compressed to spread data more symmetrically

PCA-Based Spaces

SIGGRAPH2012

4. Centre values around the mean for each channel

$$\mathbf{L} = \log L - < \log L >$$
$$\mathbf{M} = \log M - < \log M >$$
$$\mathbf{S} = \log S - < \log S >$$

5. Data reshaped to an (NxN) vector of log LMS triples

6. PCA on the resulting vector to compute 3 axes defining the custom space

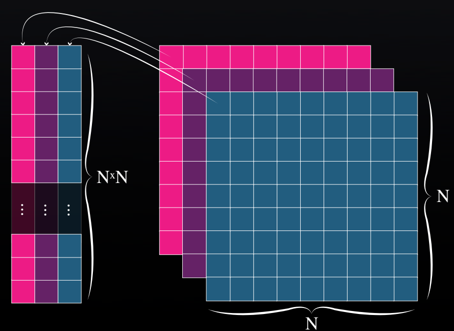

Custom Spaces Covariance

SIGGRAPH2012

Covariance for PCA-based colour spaces for the four scene types (lower is better):

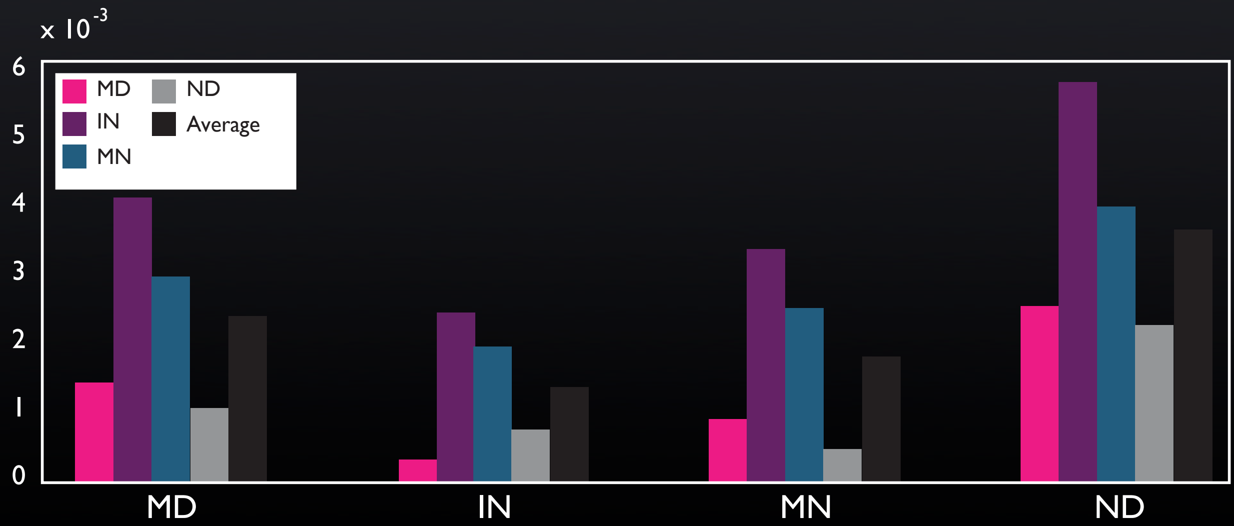

Custom Spaces Ranking

SIGGRAPH2012

Rank	MD	IN	MN	ND
1	**IN**	**IN**	**IN**	**MN**
2	MN	MN	MN	IN
3	MD	MD	MD	MD
4	ND	ND	ND	ND

Discussion

SIGGRAPH2012

- PCA space computed from the Indoors dataset ranks highest for most datasets
- Once more, high consistency across datasets - counter-intuitive!
- Covariance much lower than for standard colour spaces

Pilot Study

SIGGRAPH2012

- Does covariance predict performance of colour transfer algorithms?
- Pilot study:
    - 2 participants --- you are looking at one of them!
    - 1 colour transfer algorithm
    - Within-set colour transfers (4 datasets)
    - Count successful colour transfers for each colour space (15 colour spaces)

Pilot Study

SIGGRAPH2012

Transfer preference using several existing colour spaces (higher is better):

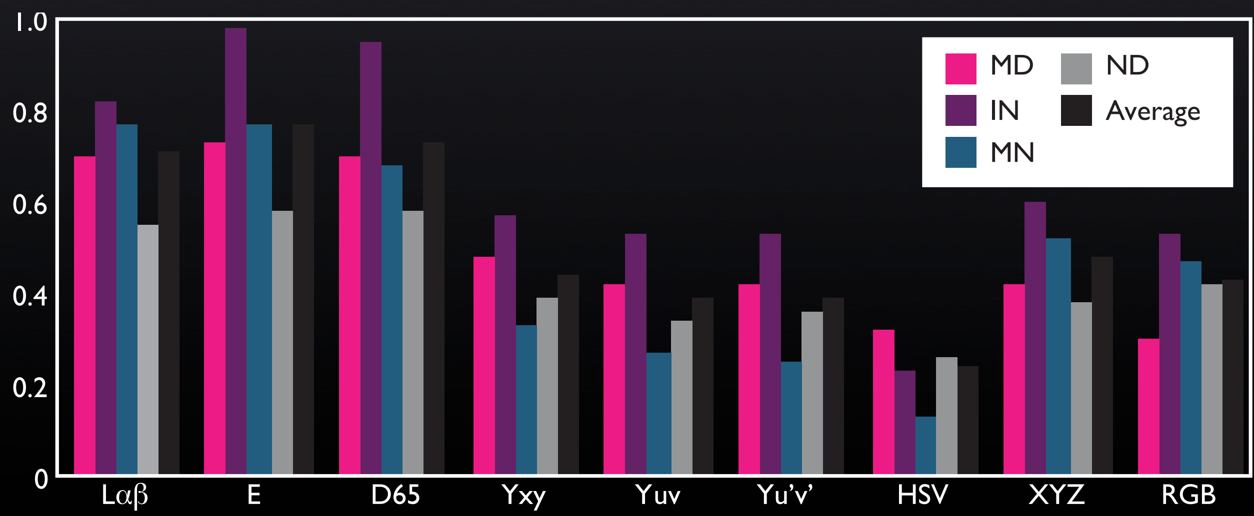

Pilot Study

SIGGRAPH2012

Transfer preference using colour spaces computed from each image ensemble:

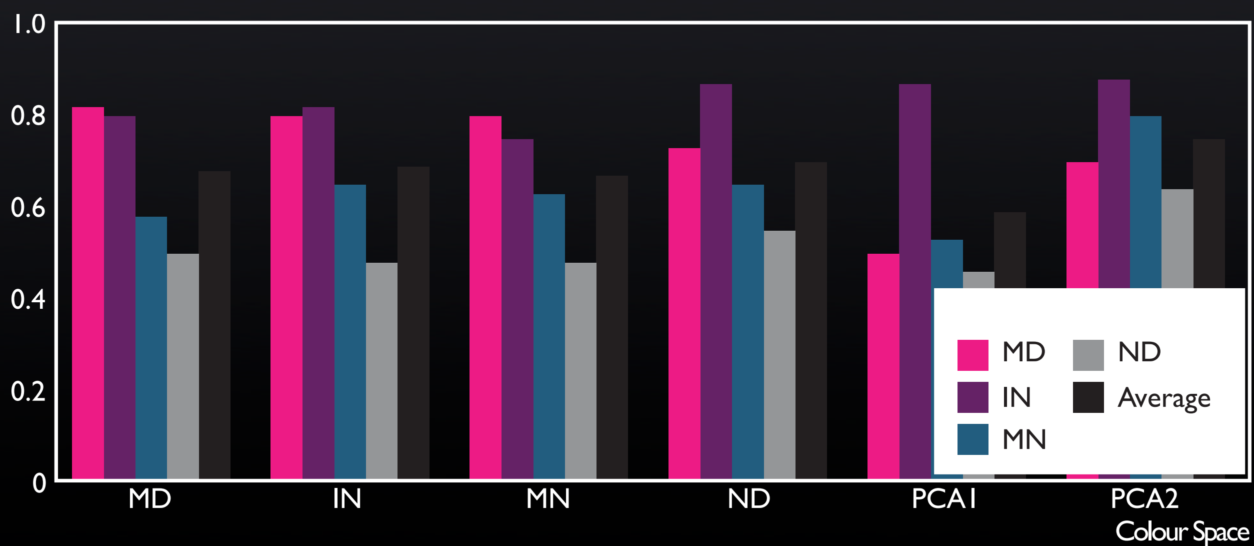

Pilot Study Ranking

SIGGRAPH2012

1. CIELAB (E) - 77%
2. PCA 2 (PCA on target image) - 75%
3. CIELAB (D65) - 73%
4. Lαβ - 71%
5. Ensemble-specific spaces - 67% - 70%

Color Space Conclusions

SIGGRAPH2012

- Average absolute co-variance predicts colour transfer success reasonably well
- Co-variance measure does not take into account gamut problems (Yxy, Yuv, Yu'v', HSV)

Color Space Conclusions

- For colour transfer using three separate channels, one colour space (CIELAB E) appears to outperform all other spaces, including PCA-based spaces
- Caveats: sample size (5 images per dataset in pilot study), study size (2 participants)

Applications

Applications

SIGGRAPH2012

- Color correction for rendered imagery
- Creative tonemapping
- Day to night
- Video
- Image Analogies

Color for rendering

SIGGRAPH2012

- Choosing colors for 3D scenes is hard.
- We can transfer colors from a photograph with similar content to our rendered scene

Color for rendering

SIGGRAPH2012

Color for rendering

SIGGRAPH2012

- This can be taken a step further
- What if we transfer **materials** from an image instead of just colors?

Material Transfer

Input image

Diffuse

Specular

Nguyen et al., 3D Material Style Transfer, Eurographics 2012

Material Transfer

Segmentation

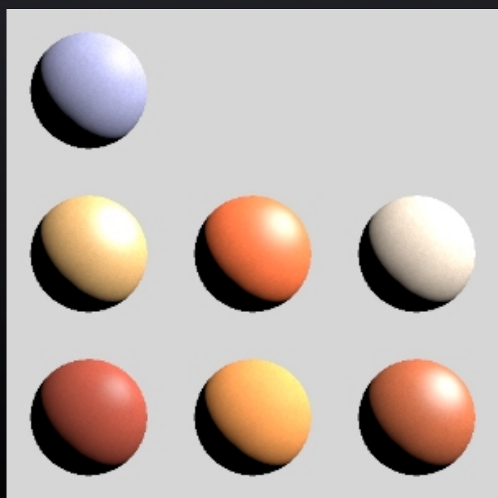

Materials

Nguyen et al., 3D Material Style Transfer, Eurographics 2012

Material Transfer

SIGGRAPH 2012

Nguyen et al., 3D Material Style Transfer, Eurographics 2012

Creative Tonemapping

SIGGRAPH 2012

- Most tonemapping solutions are either automatic...
- ...or they require manual parameter adjustment
- Photoshop solution:

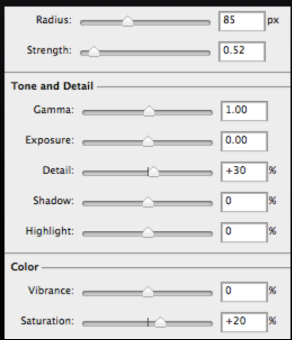

Creative Tonemapping

SIGGRAPH2012

- We can tonemap images using a reference instead
- The target image specifies both **color palette** and **dynamic range**

Creative Tonemapping

SIGGRAPH2012

Source: HDR
Target: LDR
Output: LDR

Source HDR (linear)

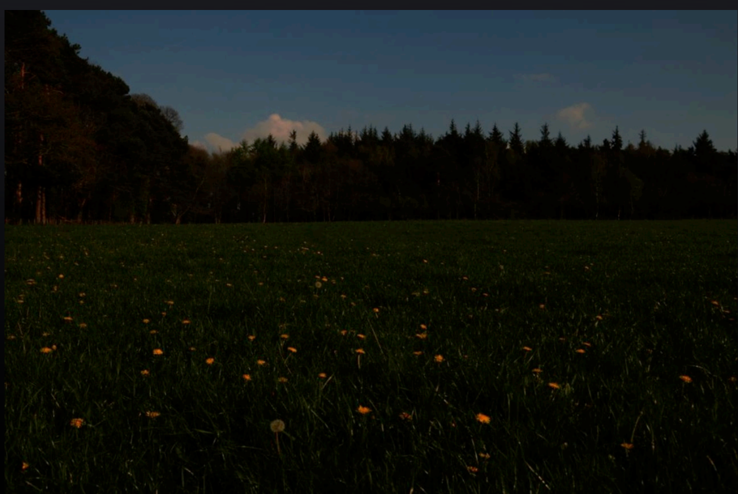

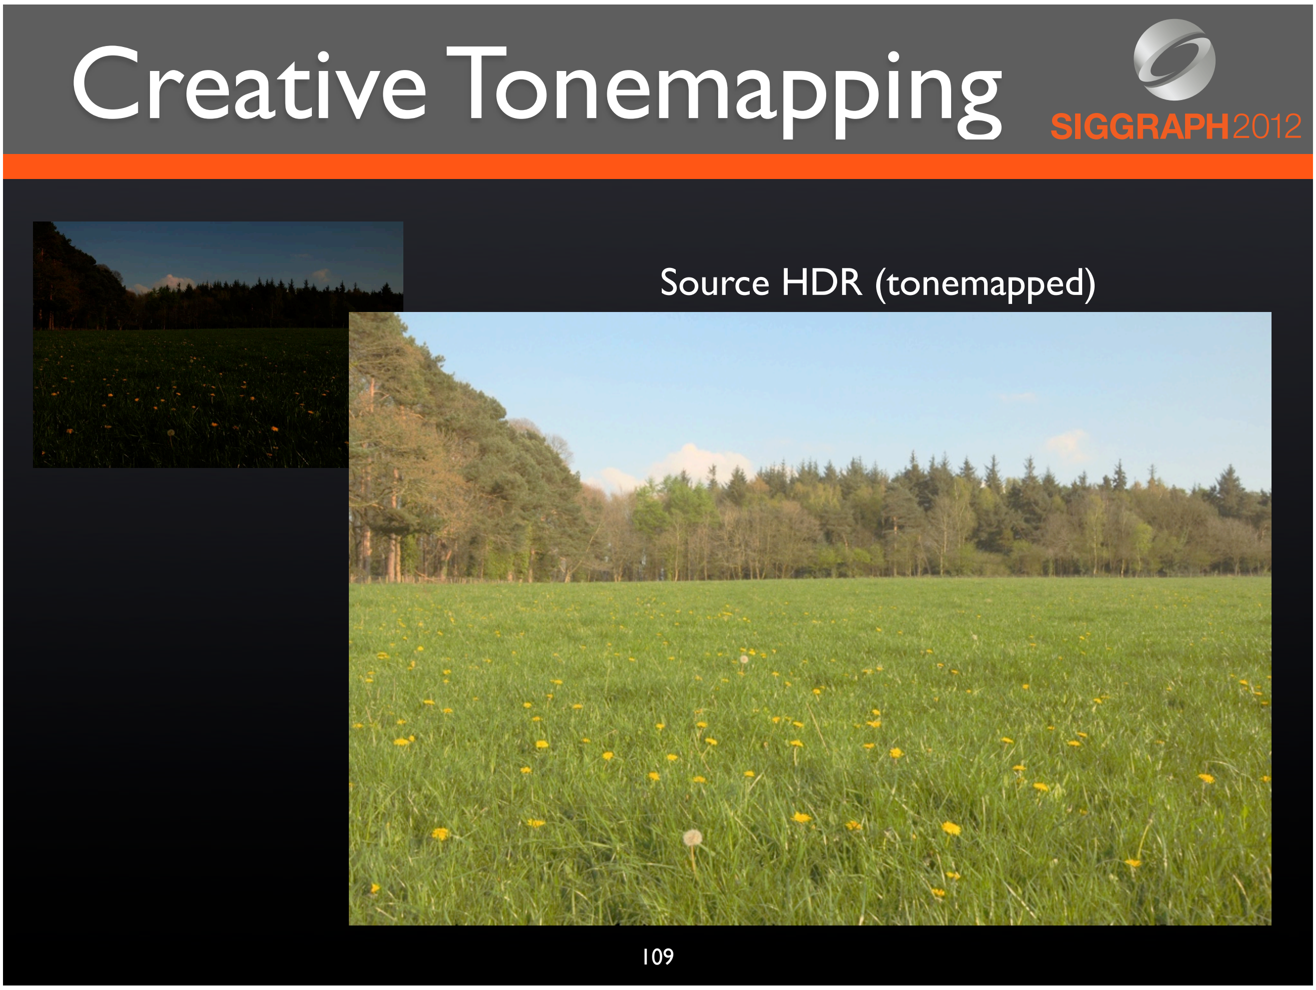
Creative Tonemapping
SIGGRAPH2012
Source HDR (tonemapped)
109

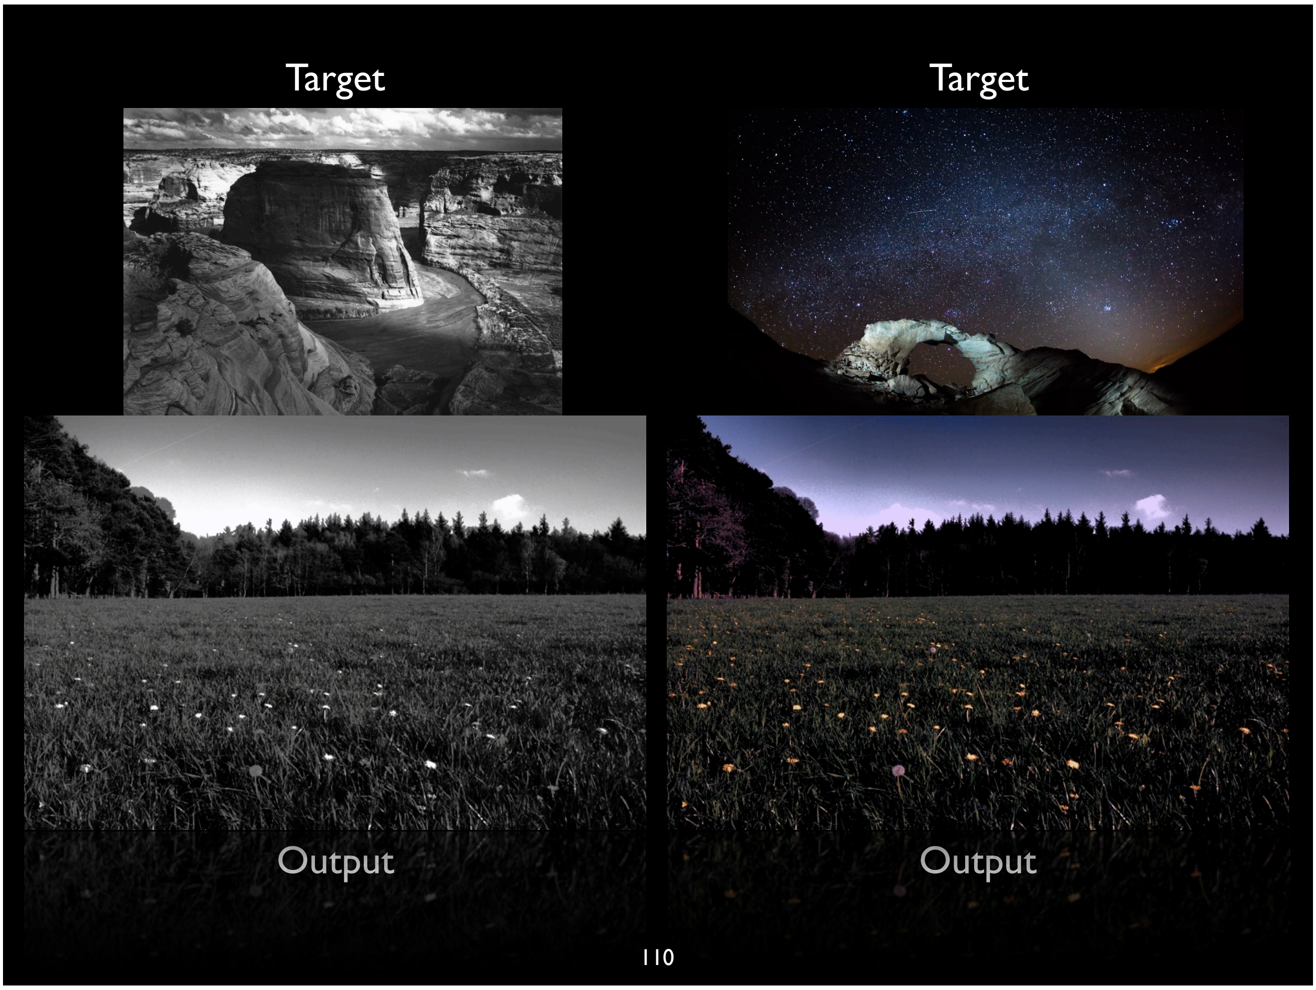
Target
Target
Output
Output
110


Output
Target
111


Contrast Adjustment
SIGGRAPH2012
Tonemapped source
Linear luminance
-
Bilateral (base)
=
Detail
112

Contrast Adjustment

SIGGRAPH2012

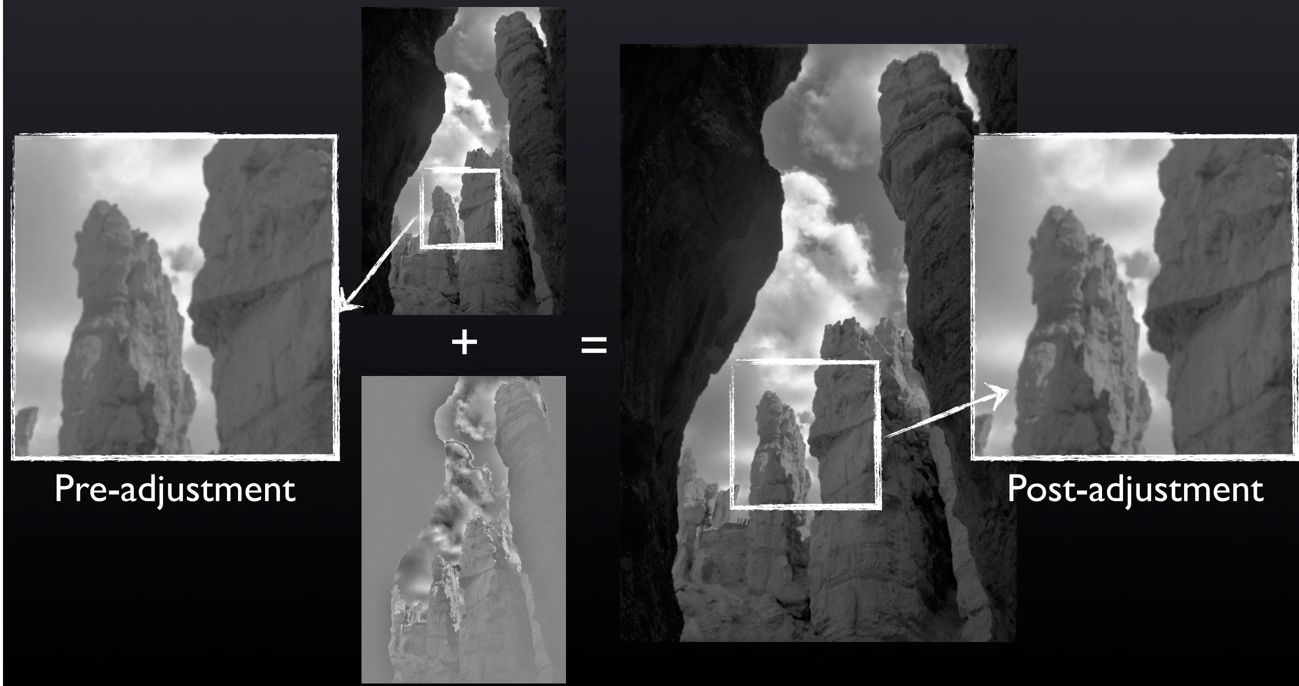

Pre-adjustment

Post-adjustment

HDR Examples

SIGGRAPH2012

Source HDR
(linear)

Target LDR

HDR Examples

SIGGRAPH2012

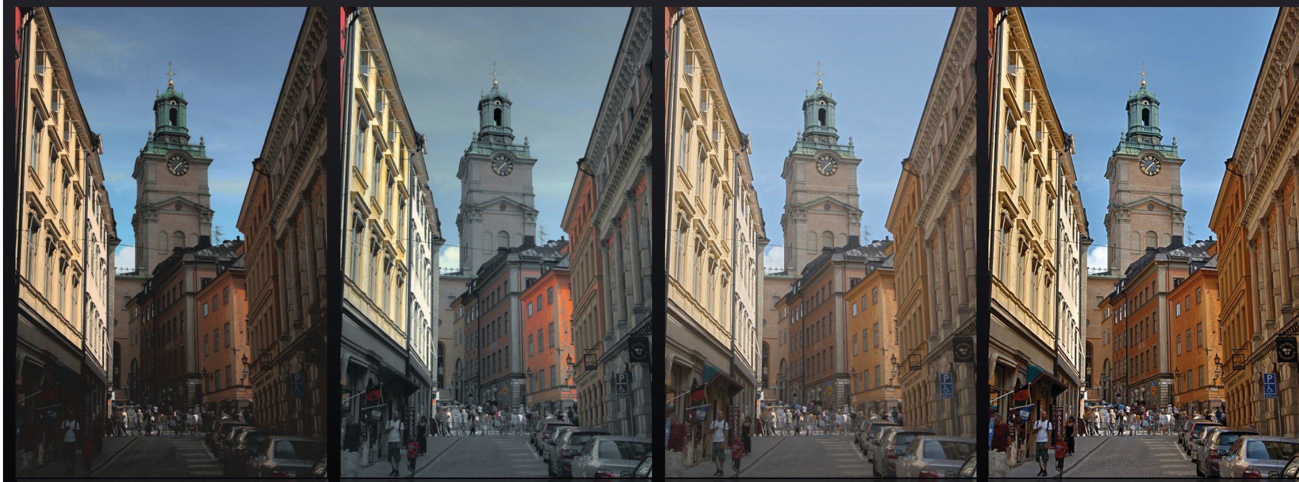

Partial match | Full match | Reinhard et al. (2002) | Li et al. (2005)

Creative Tonemapping

SIGGRAPH2012

Tania Pouli, Erik Reinhard, 'Progressive Color Transfer for Images of Arbitrary Dynamic Range', Computers and Graphics 35(1), pp. 67-80, 2011

Night Images

SIGGRAPH2012

- Color transfer is very good at faking night images
- Night scenes are almost monochromatic
- Makes the transfer easier

Night Images

SIGGRAPH2012

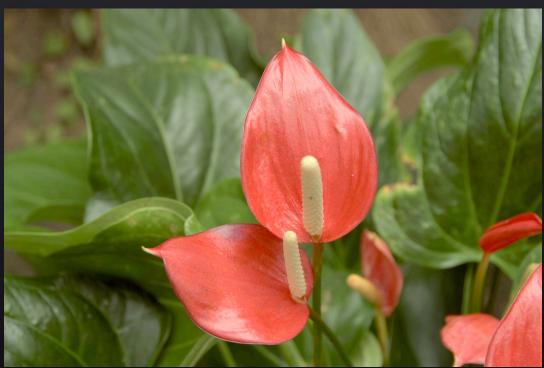

Night Images

SIGGRAPH2012

Video

SIGGRAPH2012

- How about video content?
- Many source frames - 1 target image
- Several issues to solve:
  - Flicker
  - Temporal coherence - same color should be attributed to same object even as they move and change shape

Image Analogies

- Technique to learn filters by presenting an example pair
- Can be used for instance for texture synthesis

A Hertzmann, C Jacobs, N Oliver, B Curless and D Salesin, Image Analogies, SIGGRAPH 2001.

Image Analogies

SIGGRAPH2012

Discussion

Discussion

SIGGRAPH2012

- Color transfer and example based solutions in general offer a more intuitive control mechanism
- Not suitable for everything but already many applications
- How do we quantify a **successful** transfer?