Color in Image & Video Processing Applications

Theo Gevers
Joost van de Weijer

Image/Video Applications

Image Segmentation:

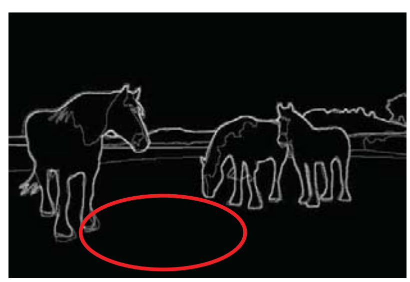

human segmentation

Video retrieval:

video sequences

Why use Color ?

photometric invariance

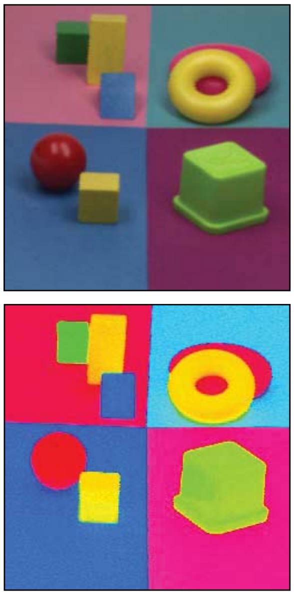

discriminative power

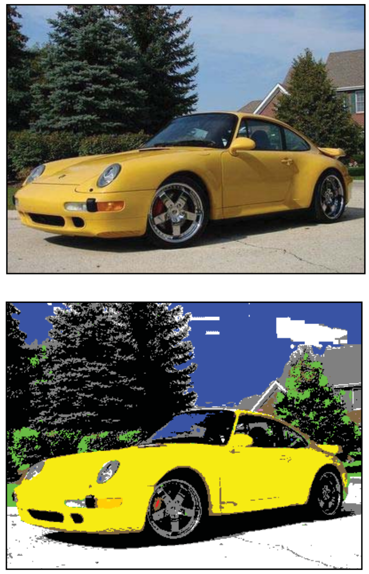

saliency detection

overview

PART I (low-level)
Joost van de Weijer

1. ***Reflection Models***
   - Dichromatic reflection model
   - Color Spaces
2. ***Color Differential Structure***
   - Color Edges
   - Photometric Invariant Edge Detection
3. ***Saliency and Color Boosting***
   - Itti and Koch model
   - Color boosted
4. ***Color Constancy***
   - At the pixel
   - Low-level
   - High-level

PART II (high-level)
Theo Gevers

1. ***Interest point detection***
   - Harris Laplace
   - Color boosted
2. ***Descriptors***
   - SIFT
   - Extension to color
3. ***Object recognition (VOC/TRECVid)***
   - Dense and point sampling
   - Code book generation
   - Results
4. ***Applications***
   - Tracking in video
   - Object replacement
   - Emotion recognition
   - Head pose estimation

Reflection Models

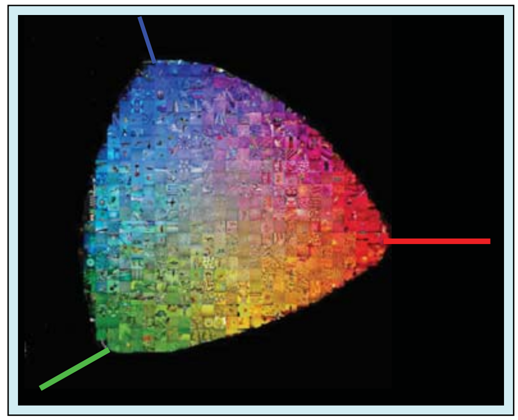

Electromagnetic radiation spectrum

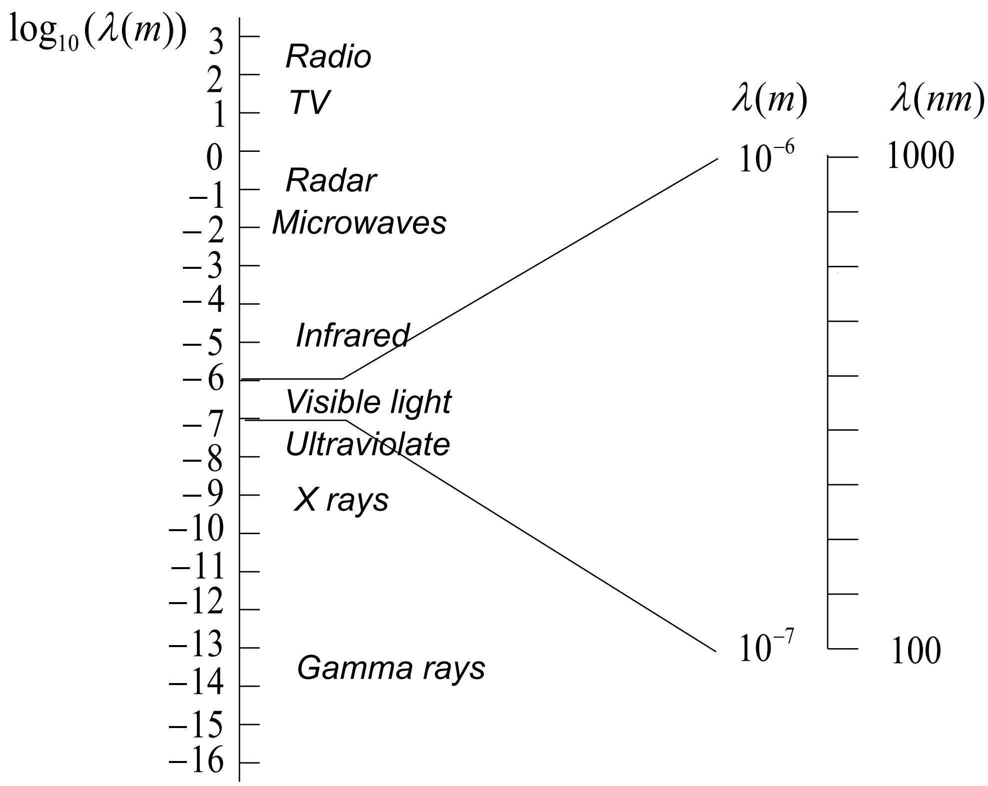

visible light spectrum

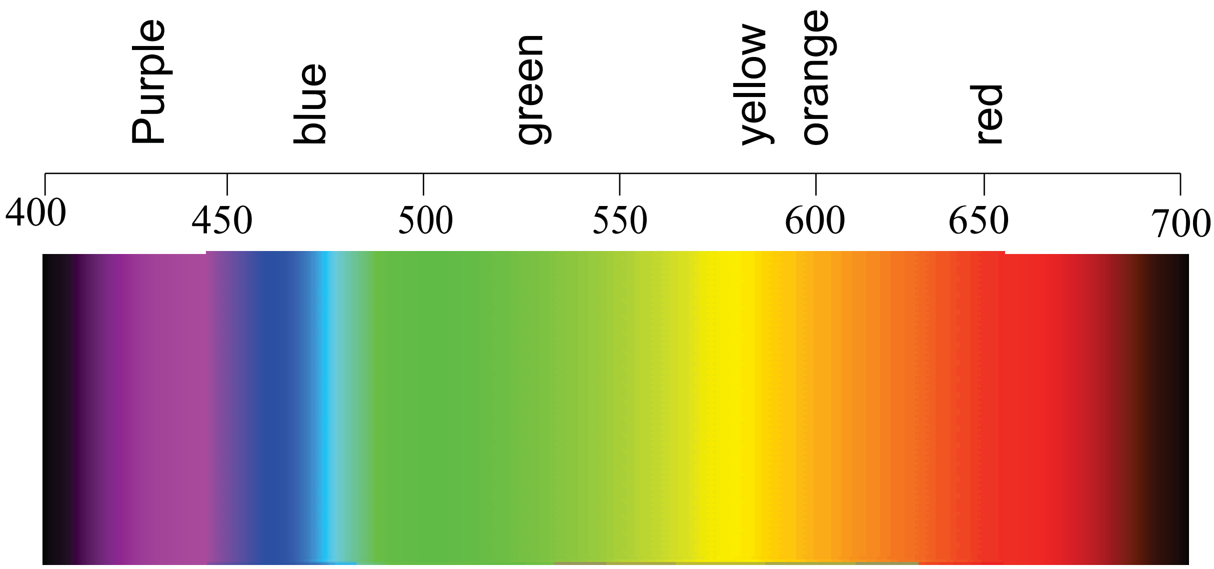

$nanometers \Rightarrow \; 1nm = 10^{-9}m$

Spectra.exe

http://askabiologist.asu.edu/research/seecolor/atable.html

slide credit: R. Baldrich

Surface reflectance

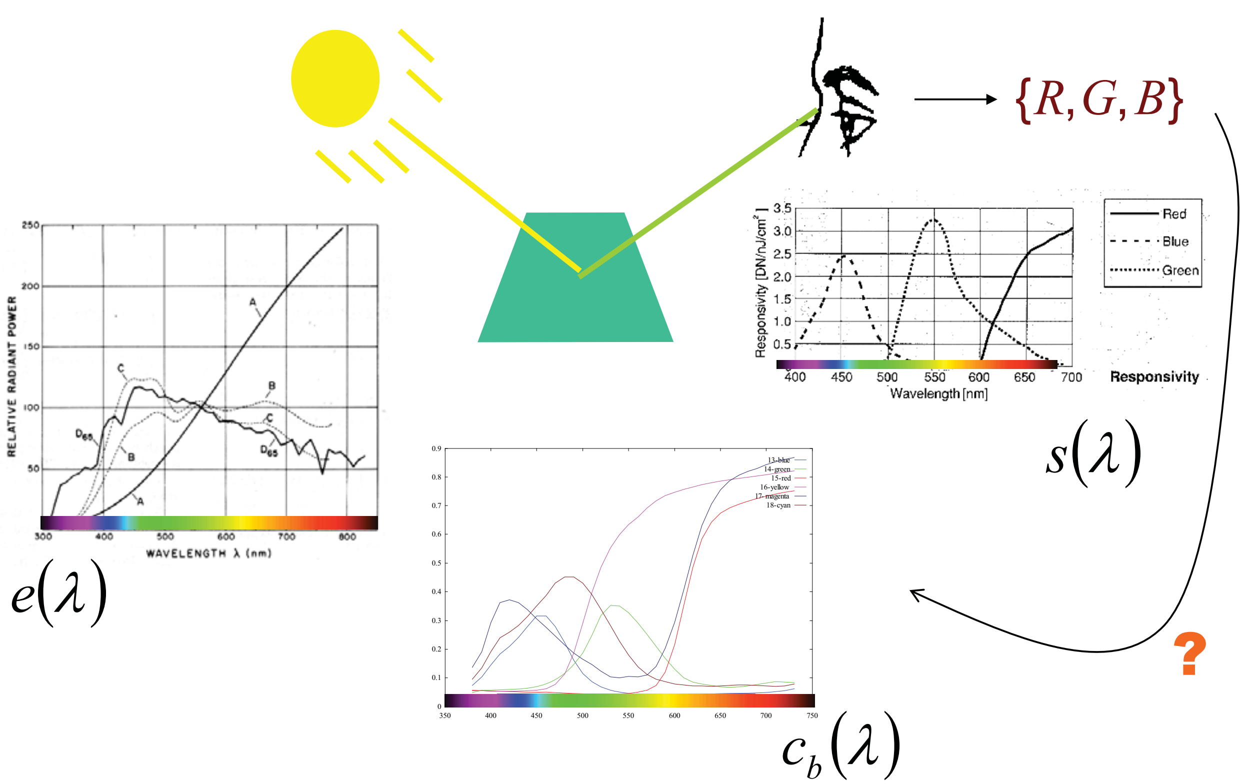

Reflecting materials

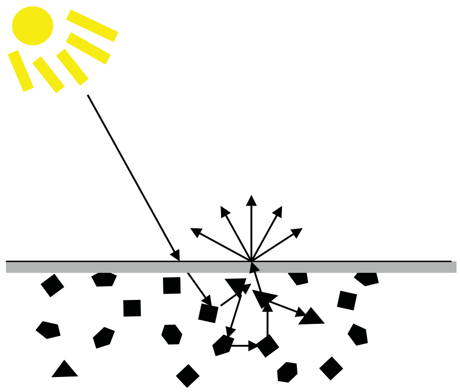

Body Reflectance
Diffuse reflection, isotropic reflection. The spectral distribution depends on colorants.

$$f_b(\lambda,\Theta) = m_b(\Theta) c_b(\lambda)$$

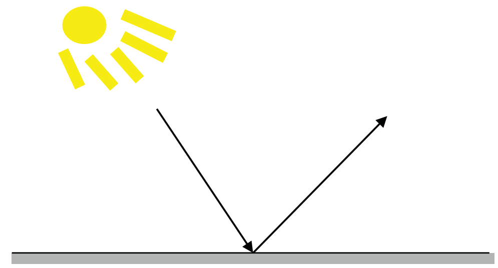

Surface Reflectance
Specular reflection. The reflection angle is similar to the incident angle. Its spectral distribution depends on the illuminant.

$$f_s(\lambda,\Theta) = m_s(\Theta) c_s(\lambda) \approx m_s(\Theta) h$$

slide credit: R. Baldrich

Dichromatic reflection model:

$$\boxed{f(\lambda,\Theta) = f_b(\lambda,\Theta) + f_s(\lambda,\Theta)}$$

$f_b(\lambda,\Theta)$: Reflected light by the object body. It depends on the pigments used to colour the object and it's the one that makes the object look coloured. (Diffuse reflectance)

$f_s(\lambda,\Theta)$: Reflected light from the surface. It has a SPD nearly the same as the incident light. (Specular or regular reflectance)

Θ: Angles that depend on light source position, observer and surface

The spectral and geometrical terms can be separated:

$$f(\lambda,\Theta) = m_b(\Theta) c_b(\lambda) + m_s(\Theta) c_s(\lambda)$$

$$\mathbf{f} = m_b \mathbf{c_b} + m_s \mathbf{c_s}$$

slide credit: R. Baldrich

Dichromatic Reflection Model

dichromatic model for matte surfaces:

$$\mathbf{f} = m_b \mathbf{c_b}$$

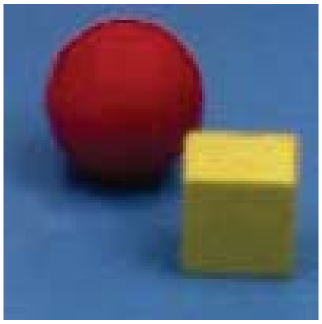

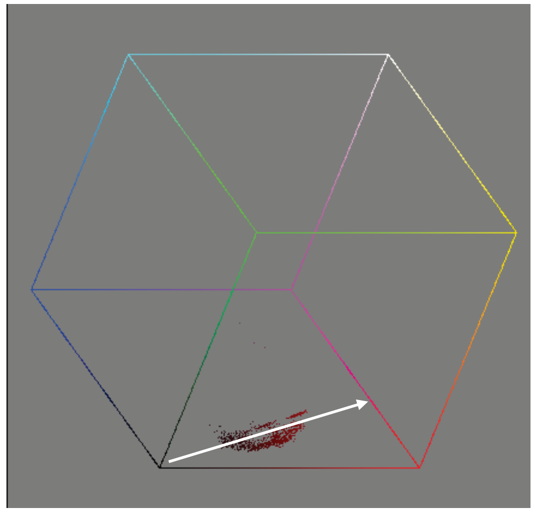

RGB-histogram

Dichromatic Reflection Model

dichromatic model for specular surfaces:

$$\mathbf{f} = m_b \mathbf{c_b} + m_s \mathbf{c_s}$$

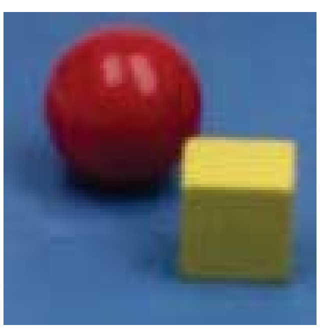

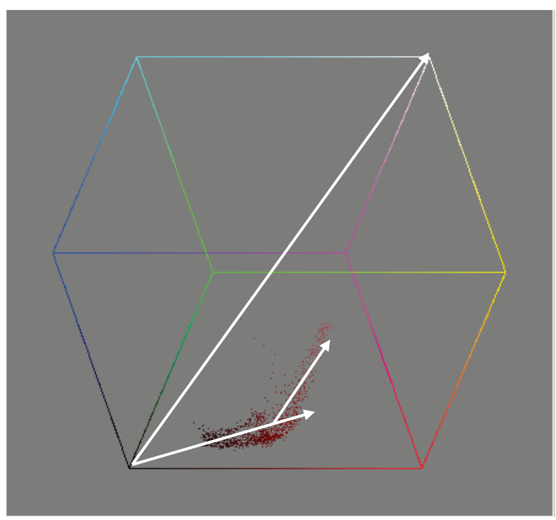

RGB-histogram

Dichromatic Reflection Model

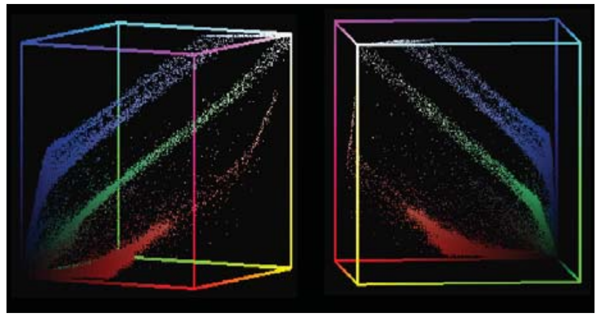

$$f(\lambda,\Theta)=m_b(\Theta)c_b(\lambda)+m_s(\Theta)c_s(\lambda)$$

- we want to describe the object independent of scene accidental events:
  - shadow- a change of $m_b(\Theta)$
  - shading- a change of $m_b(\Theta)$
  - viewpoint/orientation object - a change of $m_b(\Theta)$ and $m_s(\Theta)$
  - specularities - a change of $m_s(\Theta)c_s(\lambda)$
- the description should only be dependent on $c_b(\lambda)$

slide credit: R. Baldrich

color spaces: normalized RGB

- normalized RGB is given by:

$$\{r,g,b\}=\{\frac{R}{R+G+B},\frac{G}{R+G+B},\frac{B}{R+G+B}\}$$

- invariant for shadow and shading variations (matte surfaces):

$$r=\frac{R}{R+G+B}=\frac{\cancel{m_b}c_R^b}{\cancel{m_b}c_R^b+\cancel{m_b}c_G^b+\cancel{m_b}c_B^b}=\frac{c_r^b}{c_r^b+c_g^b+c_b^b}$$

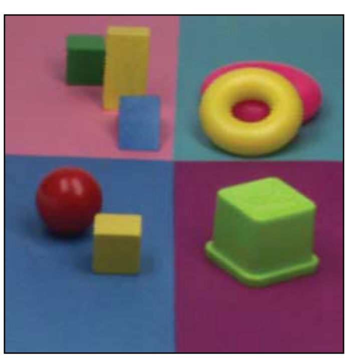

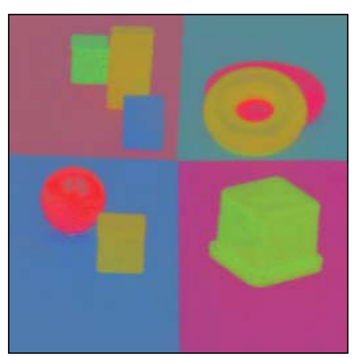

normalized RGB

color spaces: hue-saturation-intensity

• defined as: $hue = \arctan\left(\frac{\sqrt{3}(R-G)}{(R+G-2B)}\right)$

$$sat = \sqrt{\tfrac{2}{3}\left(R^2+G^2+B^2-RG-RB-GB\right)}$$

$$i \;\; = \frac{R+G+B}{\sqrt{3}}$$

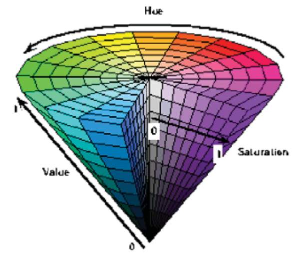

• hue is invariant for shading variations and specularities under white light:

$$hue = \arctan\left(\frac{\sqrt{3}\cancel{m^b}\left(c_R^b+\cancel{c^s}-c_G^b-\cancel{c^s}\right)}{\cancel{m^b}\left(c_R^b+\cancel{c^s}+c_G^b+\cancel{c^s}-2c_B^b-2\cancel{c^s}\right)}\right)$$

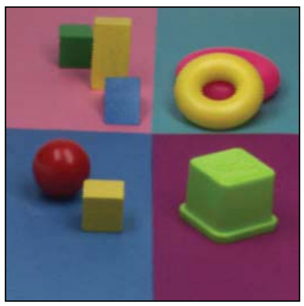

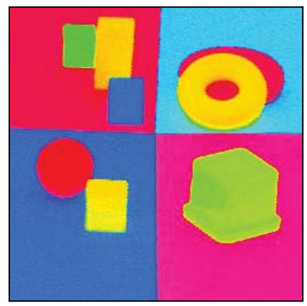

hue

Take care of instabilities

• when working in different color spaces always take instabilities into account !

• Error propagation is a convenient tool for instability evaluation:

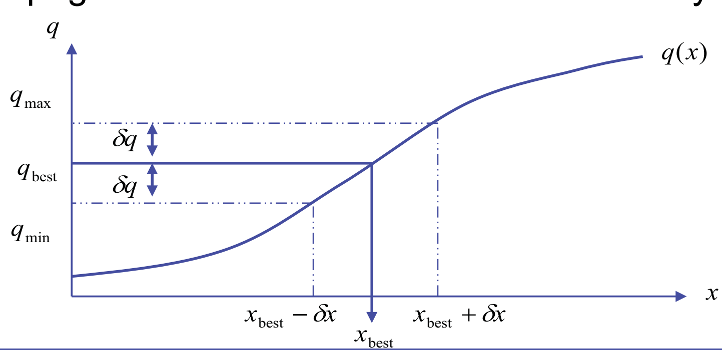

Suppose that $u,\ldots,w$ are measured with corresponding uncertainties $\sigma_u,\ldots,\sigma_w$ to compute function $q(u,\ldots,w)$.

The predicted uncertainty is defined by :

$$\sigma_{\mathrm{q}} = \sqrt{\left(\frac{\partial q}{\partial u}\sigma_u\right)^2 + \ldots + \left(\frac{\partial q}{\partial w}\sigma_w\right)^2}$$

Take care of instabilities

• when working in different color spaces always take instabilities into account !

• Error propagation is a convenient tool for instability evaluation:

Ex. 1

$$hue = \arctan\left(\frac{\sqrt{3}(R-G)}{(R+G-2B)}\right) \longrightarrow (\partial hue)^2 = \left(\partial\left(\arctan\left(\frac{\sqrt{3}(R-G)}{(R+G-2B)}\right)\right)\right)^2$$

$$(\partial hue)^2 = \left(\frac{\partial hue}{\partial R}\right)^2 \partial^2 R + \left(\frac{\partial hue}{\partial G}\right)^2 \partial^2 G + \left(\frac{\partial hue}{\partial B}\right)^2 \partial^2 B$$

$$= \frac{1}{sat^2}\partial^2 R \text{ (assuming } \partial^2 R = \partial^2 G = \partial^2 B)$$

Take care of instabilities

• when working in different color spaces always take instabilities into account !

• Error analysis is a convenient tool for instability evaluation:

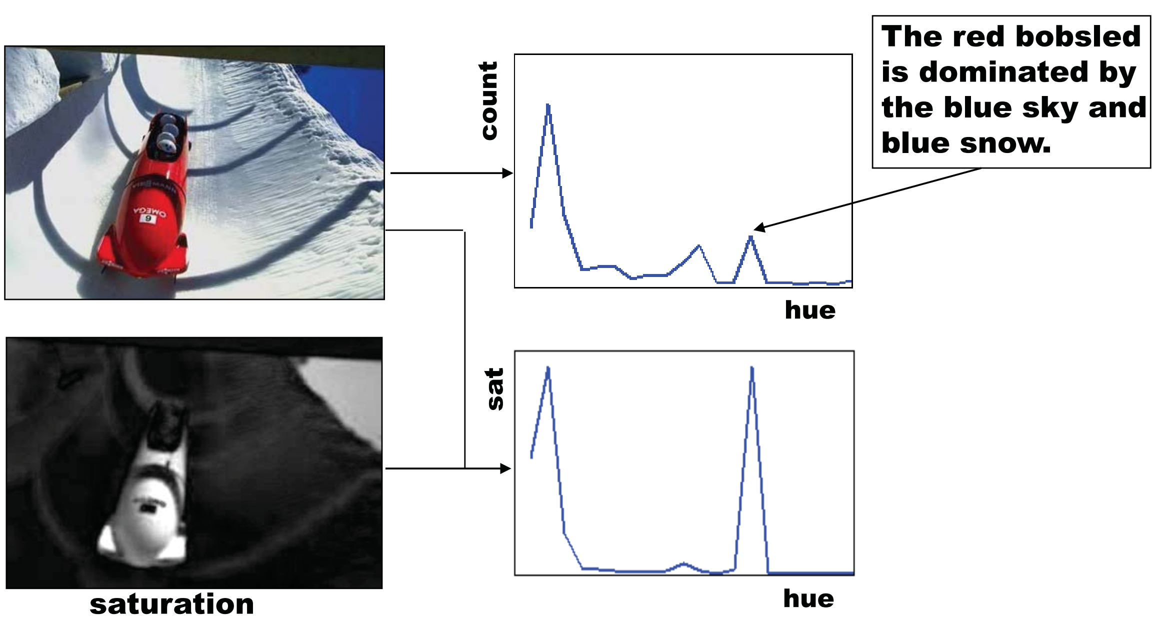

references: photometric invariants

- S.A. Shafer. *Using color to separate reflection components*. Color research and applications, 1985.
- G.J. Klinker et al.. *A physical approach to color image understanding.* IJCV, 1990.
- M.J. Swain, D.H. Ballard . *Color indexing*. IJCV, 1991
- T. Gevers, A.W.M. Smeulders. *Color based object recognition*. Pattern Recognition, 1999.
- J.M. Geusebroek et al. *Color Invariance.* PAMI, 2001.
- T. Gevers, H. Stokman. *Robust histogram construction from color invariance for object recognition.* PAMI, 2004.
- J. van de Weijer, C. Schmid. Coloring local feature extraction. ECCV, 2006
- B.A. Maxwell et al. *A Bi-illuminant Dichromatic Reflection Model for Understanding Images*, CVPR, 2008.
- T. Zickler et al. *Color Subspaces as Photometric Invariants*. IJCV, 2008.

overview

PART I (low-level)
Joost van de Weijer

1. ***Reflection Models***
   - Dichromatic reflection model
   - Color Spaces
2. ***Color Differential Structure***
   - Color Edges
   - Photometric Invariant Edge Detection
3. ***Saliency and Color Boosting***
   - Itti and Koch model
   - Color boosted
4. ***Color Constancy***
   - At the pixel
   - Low-level
   - High-level

PART II (high-level)
Theo Gevers

1. ***Interest point detection***
   - Harris Laplace
   - Color boosted
2. ***Descriptors***
   - SIFT
   - Extension to color
3. ***Object recognition (VOC/TRECVid)***
   - Dense and point sampling
   - Code book generation
   - Results
4. ***Applications***
   - Tracking in video
   - Object replacement
   - Emotion recognition
   - Head pose estimation

Color Differential Structure

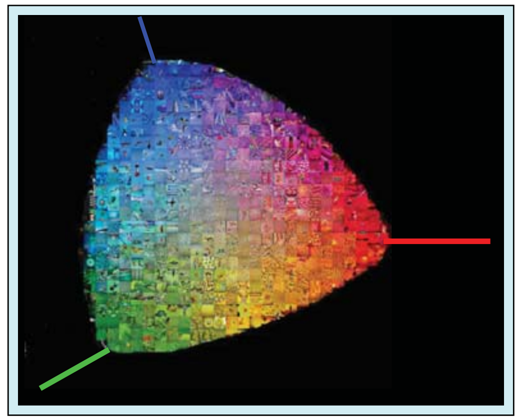

differential-based computer vision

1. How do we combine the differential structure of the various color channels ?
2. How do we incorporate color invariance theory into the measurements of the differential structure while maintaining robustness ?

isoluminance

luminance gradient: isoluminant edges are not detected.

Color Feature Detection

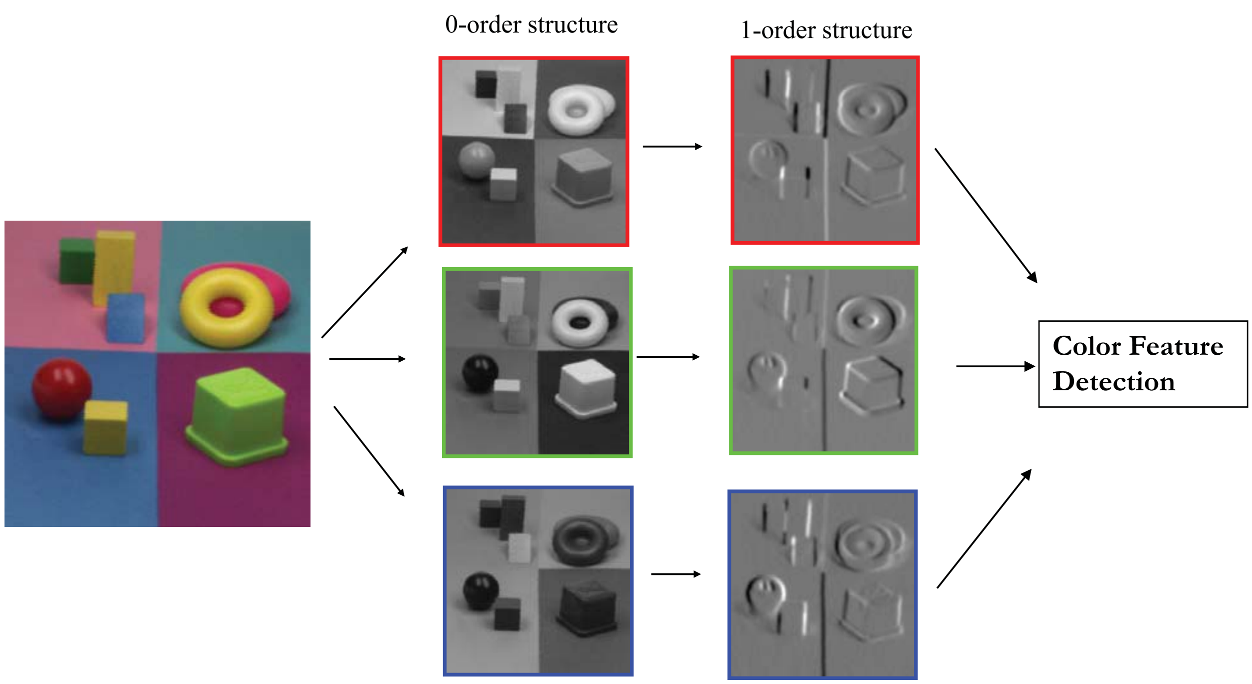

from luminance to color

vector: $R_x \quad + \quad G_x \quad = \quad 0$

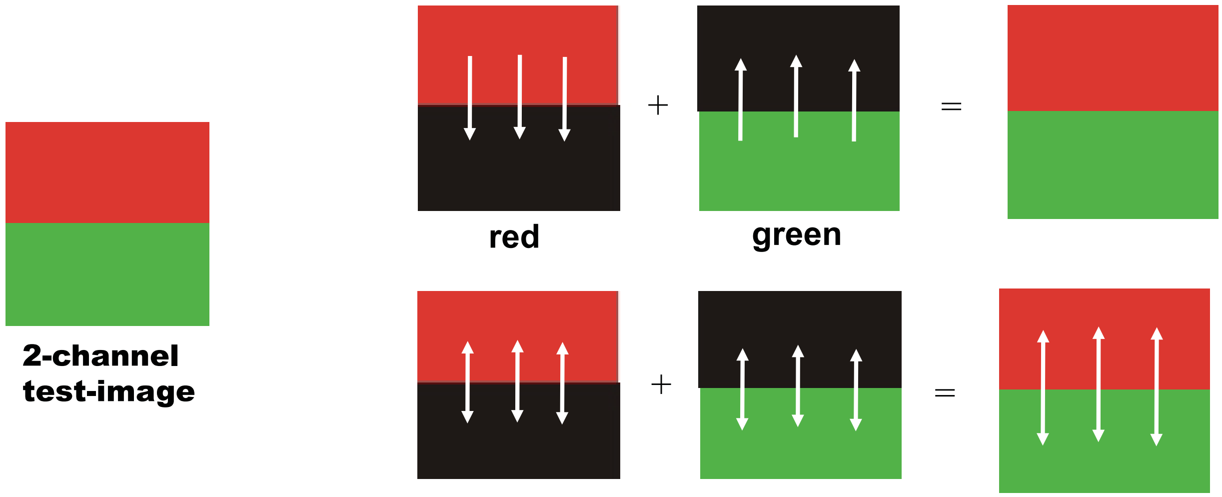

tensor: $$\begin{pmatrix} R_x^2 & R_xR_y \\ R_xR_y & R_y^2 \end{pmatrix} + \begin{pmatrix} G_x^2 & G_xG_y \\ G_xG_y & G_y^2 \end{pmatrix} = \begin{pmatrix} R_x^2+G_x^2 & R_xR_y+G_xG_y \\ R_xR_y+G_xG_y & R_y^2+G_y^2 \end{pmatrix}$$

DiZenzo. “Note on the Gradient of a Multi-Image”, *Computer Vision, Graphics, and Image Processing*, 1986.

feature detection in oriented patterns

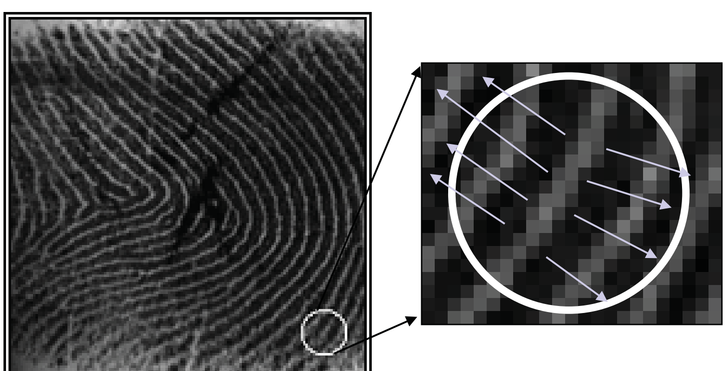

oriented texture

more tensor-based features:

- Harris corner points
- symmetry points
  (star and circle structures)
- optical flow
- orientation estimation
- curvature estimation
- ...

traditional orientation estimation:

$$\theta = \arctan\left(\frac{f_y}{f_x}\right) \rightarrow \bar{\theta} = \arctan\left(\frac{\overline{f_y}}{\overline{f_x}}\right)$$

tensor-based orientation estimation:

$$\theta = \arctan\left(\frac{2f_xf_y}{f_x^2 - f_y^2}\right) \rightarrow \bar{\theta} = \arctan\left(\frac{2\overline{f_xf_y}}{\overline{f_x^2 - f_y^2}}\right)$$

differential-based computer vision

1. How do we combine the differential structure of the various color channels ?
2. How do we incorporate color invariance theory into the measurements of the differential structure while maintaining robustness ?

Photometric Invariant Edge Detection

- we differ between three types of edges
  1. material edge
  2. shadow/shading edge
  3. specular edge
- assumptions:
  1. white illumination
  2. neutral interface reflection
  3. shadows are not colored.

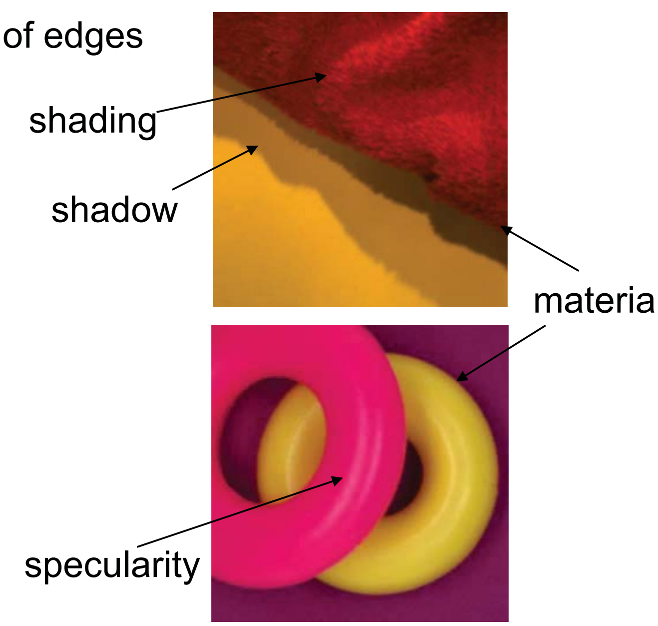

Computation of quasi-invariance

nonlinear transformation

IMAGE → FULL INVARIANT → FULL INVARIANT DERIVATIVE

linear operation

IMAGE → DERIVATIVES → QUASI- INVARIANT DERIVATIVE

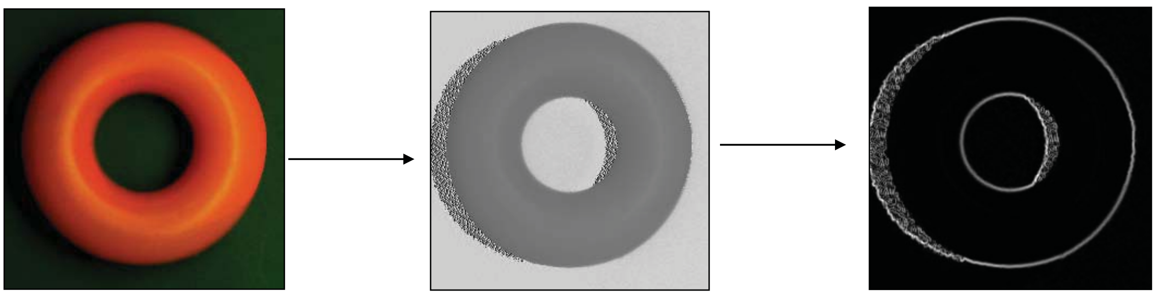

hue

hue_x

Shadow-Shading-Specular Quasi-Invariant

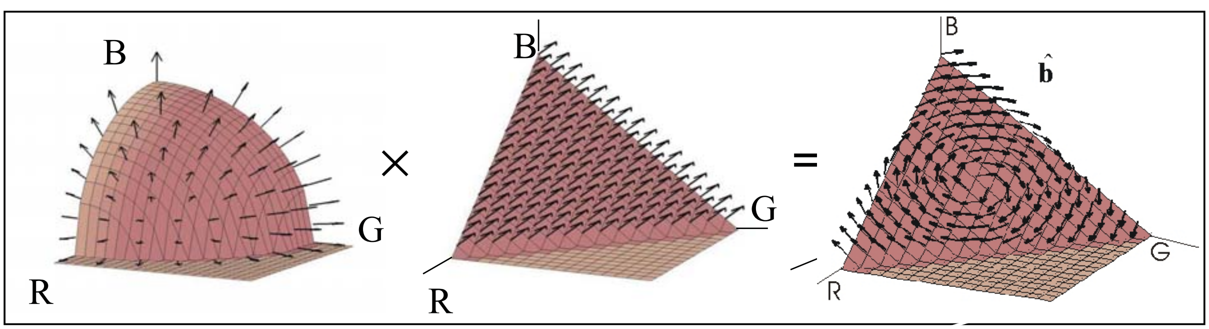

spherical coordinates	*opponent colors*	*hue-saturation-intensity*
shading variant	specular variant	shading-specular variant
shading invariant	specular invariant	shading-specular invariant

spherical coordinates

- For matte surfaces : $\mathbf{f} = m^b \mathbf{c}^b$
- all shadow-shading variation is in the radial direction

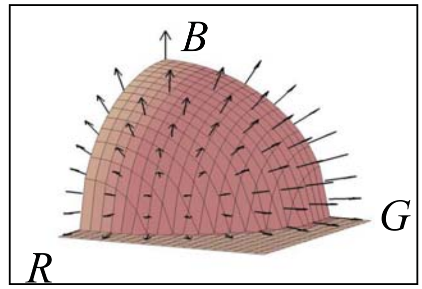

shadow/shading direction

uncertainty of c_x

$$\mathbf{f}_x = \begin{pmatrix} R_x \\ G_x \\ B_x \end{pmatrix} \xrightarrow{spherical} \begin{pmatrix} r_x \\ r\varphi_x \\ \sin\varphi\,\theta_x \end{pmatrix} = \begin{pmatrix} r_x \\ 0 \\ 0 \end{pmatrix} + r \begin{pmatrix} 0 \\ \varphi_x \\ \sin\varphi\,\theta_x \end{pmatrix} \longrightarrow \mathbf{c}_x = \begin{pmatrix} 0 \\ \varphi_x \\ \sin\varphi\,\theta_x \end{pmatrix}$$

hue-saturation-intensity

- For specular surfaces : $\mathbf{f} = m^b \mathbf{c}^b + m^s \mathbf{c}^s$
- there is no specular-shadow-shading variation in the hue-direction.

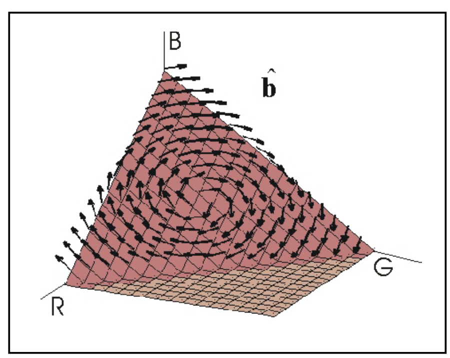

the hue direction

uncertainty of h_x

$$\mathbf{f}_x = \begin{pmatrix} R_x \\ G_x \\ B_x \end{pmatrix} \xrightarrow{hsi} \begin{pmatrix} sh_x \\ s_x \\ i_x \end{pmatrix} = \begin{pmatrix} 0 \\ s_x \\ i_x \end{pmatrix} + s \begin{pmatrix} h_x \\ 0 \\ 0 \end{pmatrix} \longrightarrow \mathbf{h}_x = \begin{pmatrix} h_x \\ 0 \\ 0 \end{pmatrix}$$

invariant edge detection applications

~~Colo Feature Extraction~~

~~Multi Image Applications~~ ~~• image retrieval~~

Color Feature Detection

Single Image Applications

• snakes

• feature extraction

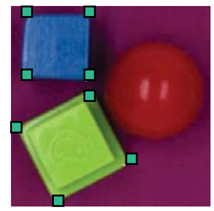

Instabilities

shadow-shading invariance:

$$\lim_{\{R,G,B\}\to 0}$$

specular-shadow-shading invariance:

$$\lim_{\{R,G,B\}\to \alpha\{1,1,1\}}$$

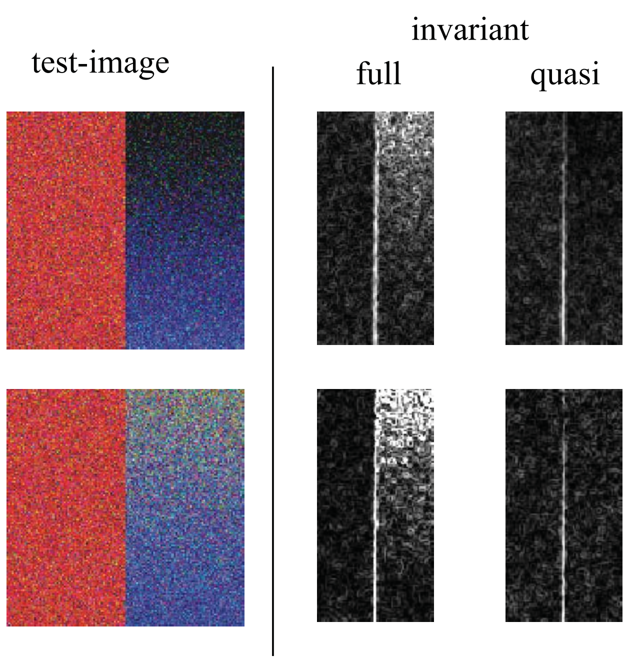

Edge Detection

• experiments conducted on pantone colorset (1012) which is used to compose 500.000 edges.

• edge detection is based on the maximum response path of the derivative energy.

• edges are tested on
  - edge displacement.
  - percentage of missed edges.

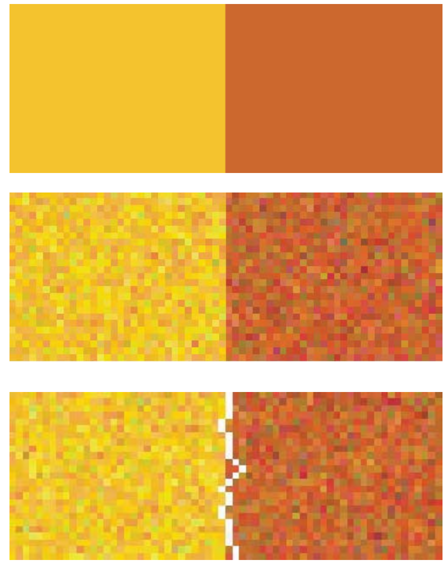

Edge Detection

• experiments conducted on pantone color set (1012) which is used to compose 500.000 edges.

• edge detection is based on the maximum response path of the derivative energy.

• edges are tested on
  - edge displacement.
  - percentage of missed edges.

shadow-shading:

	Δ	ε
full	0.21	2.0
quasi	0.043	0.99

specular-shadow-shading:

	Δ	ε
full	0.85	9.8
quasi	0.35	5.8

• Conclusion: Quasi invariants more than half the edge displacement, and have higher discriminative power.

experiments : canny edge detection

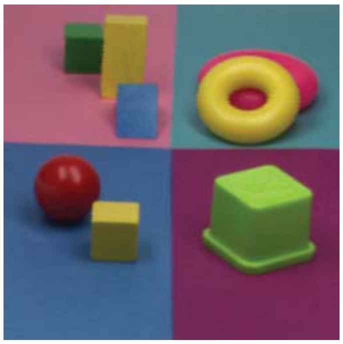

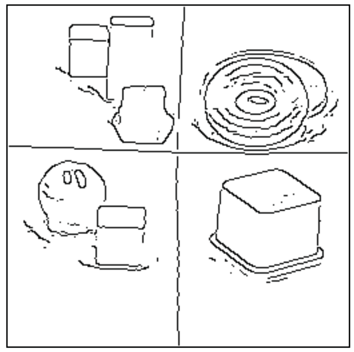

luminance-gradient

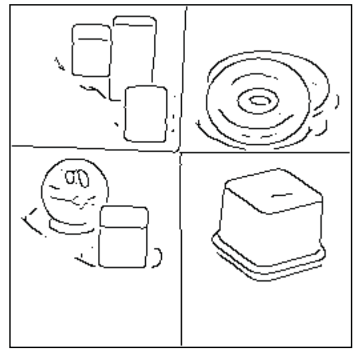

RGB-gradient

experiments : canny edge detection

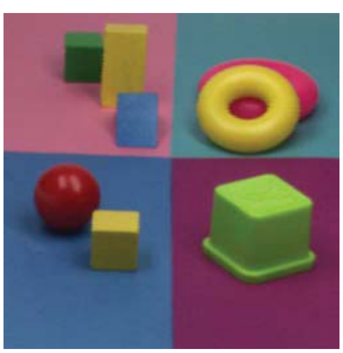

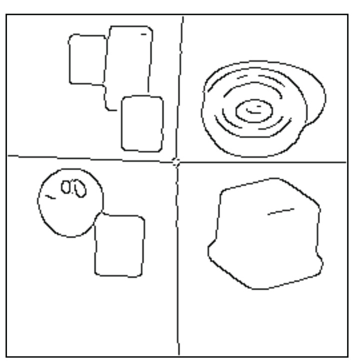

shadow-shading quasi-invariant

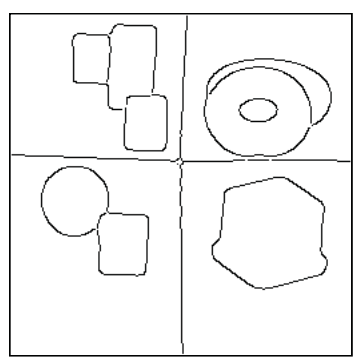

shadow-shading-specular quasi-invariant

Edge Classification

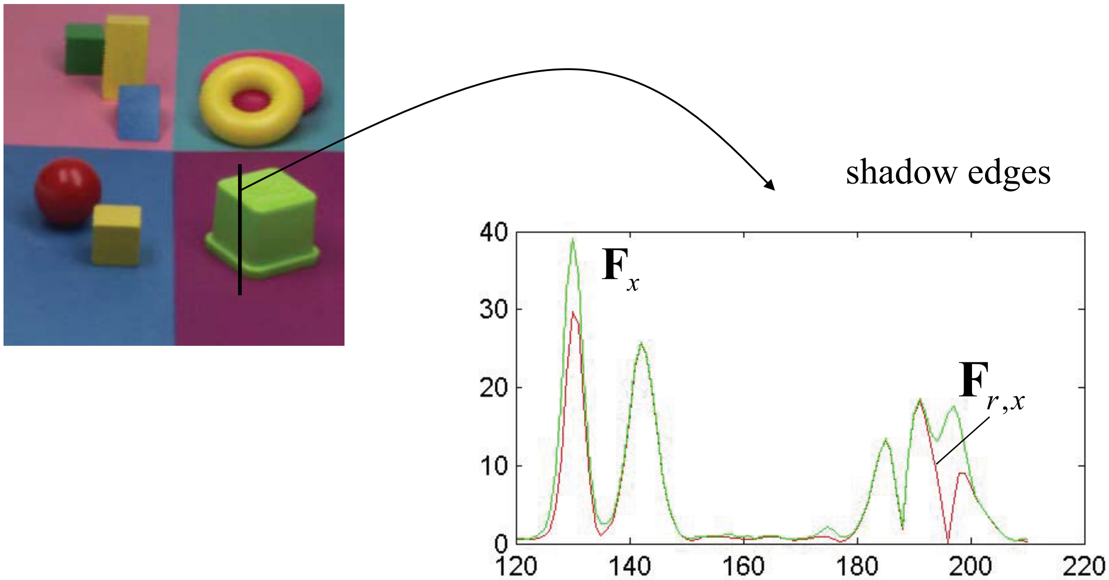

Edge Classification

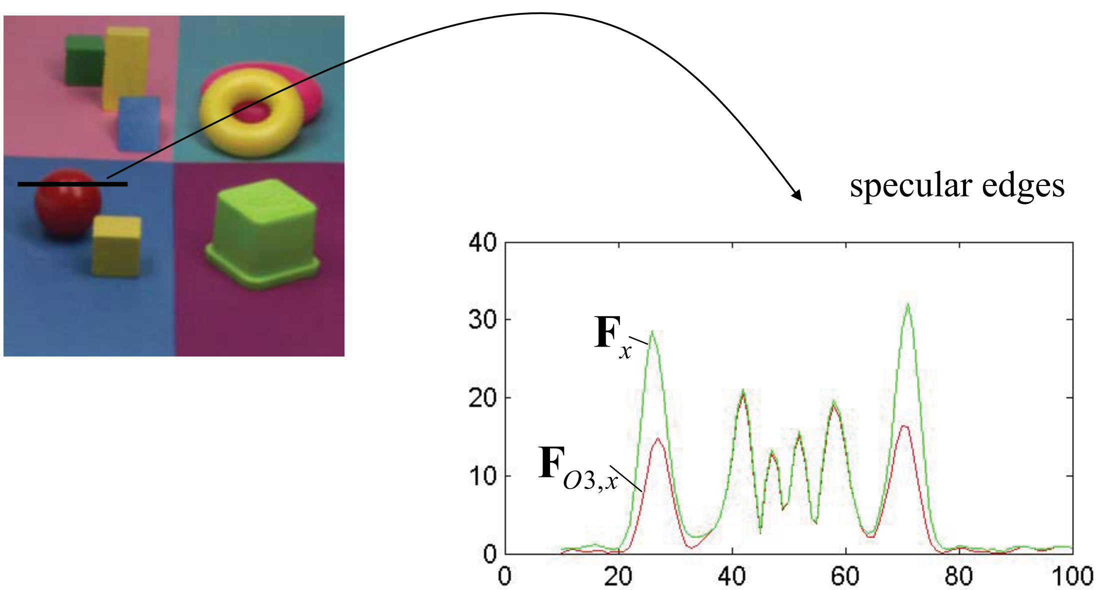

Edge Classification

red - object edge
green-shading/shadow edge
Blue – specular edge

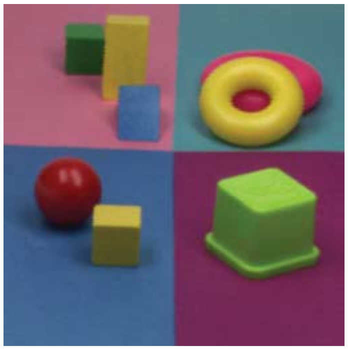

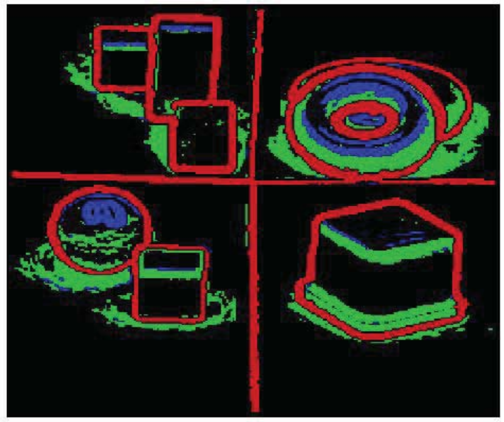

Photometric Invariant Corner Detection

• Harris corner detector combined with the quasi-invariants allows for photometric invariant corner detection

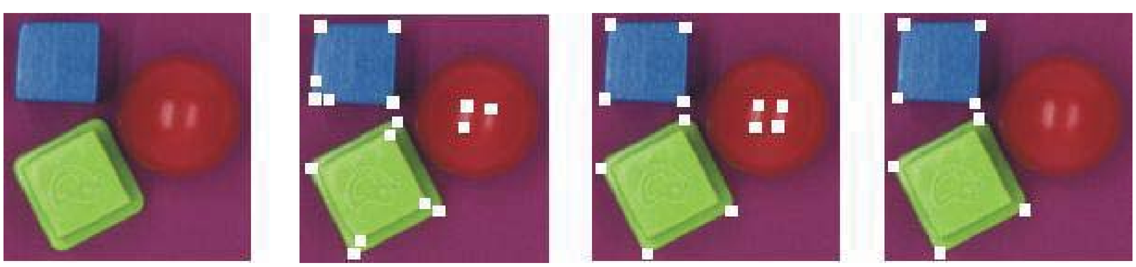

RGB | shadow-shading | specular-shadow-shading

experiments : Hough transform

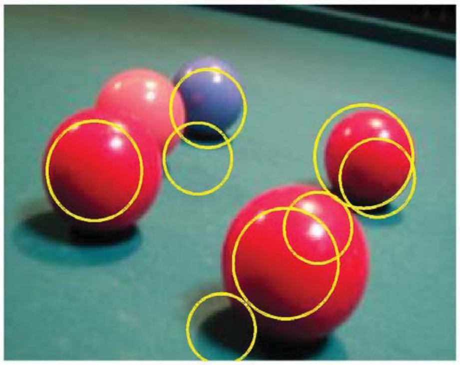

RGB-gradient

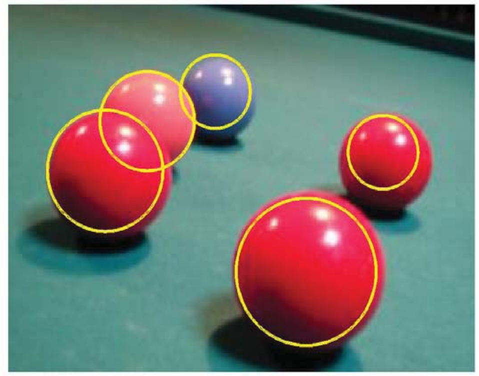

shadow-shading-specular quasi-invariant

references: color differential structure

- S. DiZenzo. *A note on the gradient of a multi-image.* Computer Vision, Graphics, and Image Processing, 1986.
- G. Sapiro and D. Ringach. *Anisotropic diffusion of multivalued images with applications to color filtering.* IEEE Image Processing, 1996.
- J.M. Geusebroek et al. *Color Invariance.* IEEE Trans. Pattern Analysis and Machine Intelligence, 2001.
- J. van de Weijer, Th. Gevers, J-M Geusebroek. *Quasi-invariant edge and corner detection*, IEEE Trans. Pattern Analysis and Machine Intelligence, 2006.
- J. van de Weijer, Th. Gevers, A.W.M. Smeulders, *Robust Photometrical Invariant Features from the Color Tensor*, IEEE T. Image Processing, 2006.

overview

PART I (low-level)
Joost van de Weijer

1. ***Reflection Models***
   - Dichromatic reflection model
   - Color Spaces
2. ***Color Differential Structure***
   - Color Edges
   - Photometric Invariant Edge Detection
3. ***Saliency and Color Boosting***
   - Itti and Koch model
   - Color boosted
4. ***Color Constancy***
   - At the pixel
   - Low-level
   - High-level

PART II (high-level)
Theo Gevers

1. ***Interest point detection***
   - Harris Laplace
   - Color boosted
2. ***Descriptors***
   - SIFT
   - Extension to color
3. ***Object recognition (VOC/TRECVid)***
   - Dense and point sampling
   - Code book generation
   - Results
4. ***Applications***
   - Tracking in video
   - Object replacement
   - Emotion recognition
   - Head pose estimation

Color Salient Features

Saliency Detection

• Goal: direct our gaze rapidly towards objects of interest in our environment.

• Visual attention is know to be driven by both *bottom up* (image based) and *top-down* (task based) cues.

• Bottom-up saliency uses simple visual attributes such as *intensity*, *contrast*, *color opponency*, *orientation* , *direction* and *velocity* of *motion*.

• What matters is *feature contrast* rather than absolute feature strength (as in center surround systems).

overview approach

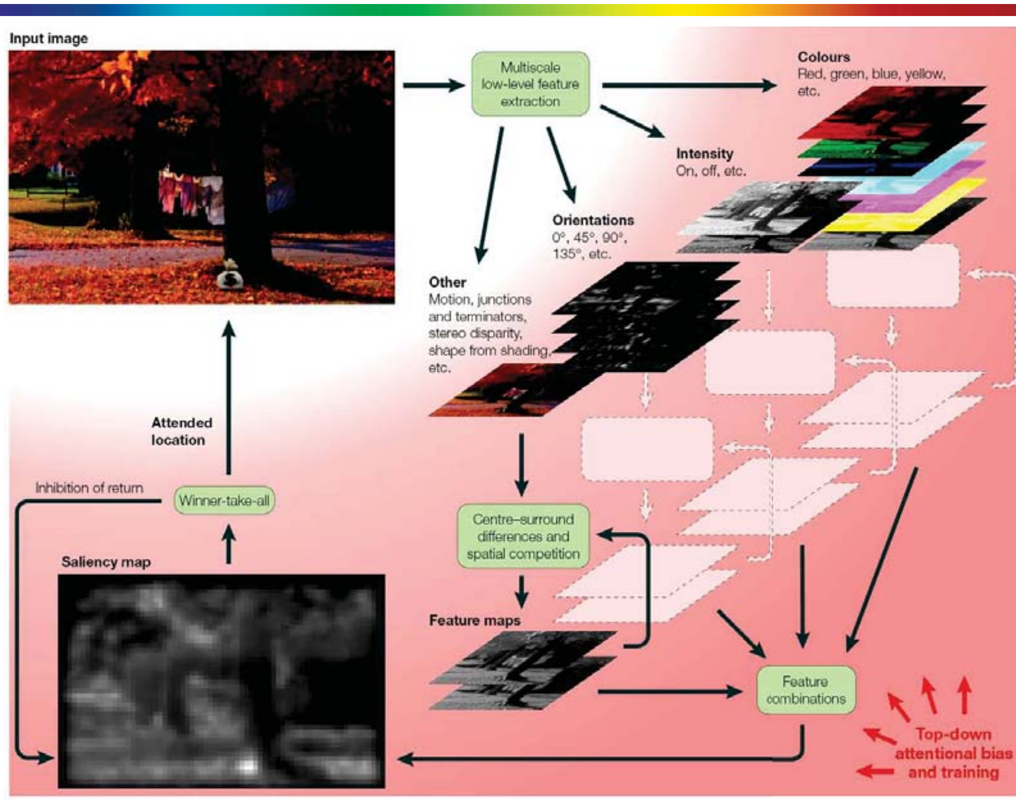

L.Itty, C. Koch "Computational Modelling of Visual Attention", Nature Reviews Neurosciende, 2001.

Computational Modeling of Visual Attention

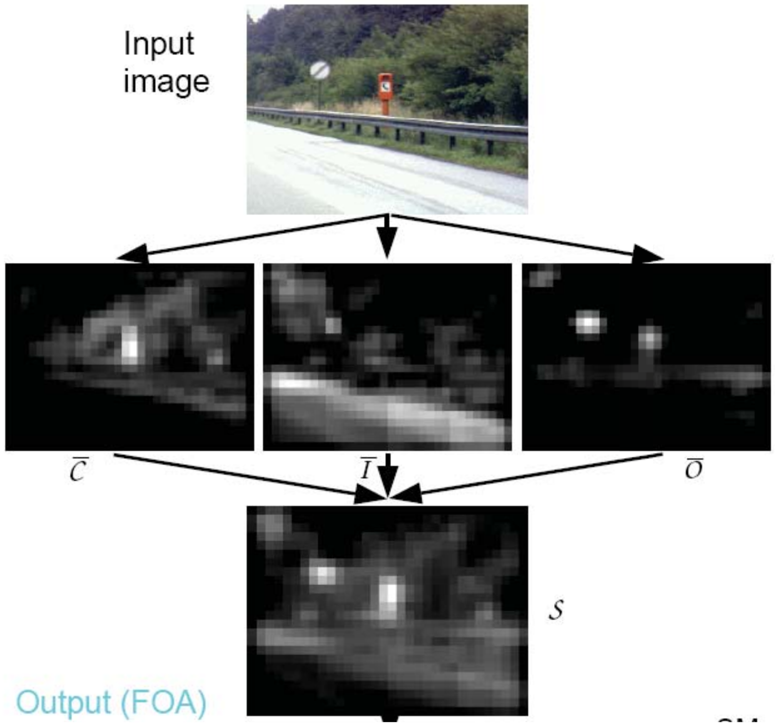

L.Itty, C. Koch "Computational Modelling of Visual Attention", Nature Reviews Neurosciende, 2001.

black-white focus of detectors

luminance-based points

color-based points

color distinctiveness

- the information content of an event, v, is equal to :

$$I(v) = -\log(p(v)) = -\log(p(\mathbf{f})p(\mathbf{f}_x)p(\mathbf{f}_y))$$

$$v = (R \quad G \quad B \quad R_x \quad G_x \quad B_x \quad R_y \quad G_y \quad B_y)$$

- equation differential-based salient point detectors : $H(\mathbf{f}_x, \mathbf{f}_y)$

$$\boxed{\textit{Color Boosting Saliency:} \quad p(\mathbf{f}_x) = p(\mathbf{f'}_x) \leftrightarrow |g(\mathbf{f}_x)| = |g(\mathbf{f'}_x)|}$$

statistics of color images:

- The statistics of $\mathbf{f}_x$ is computed by looking of the 40.000 images of the Corel database.

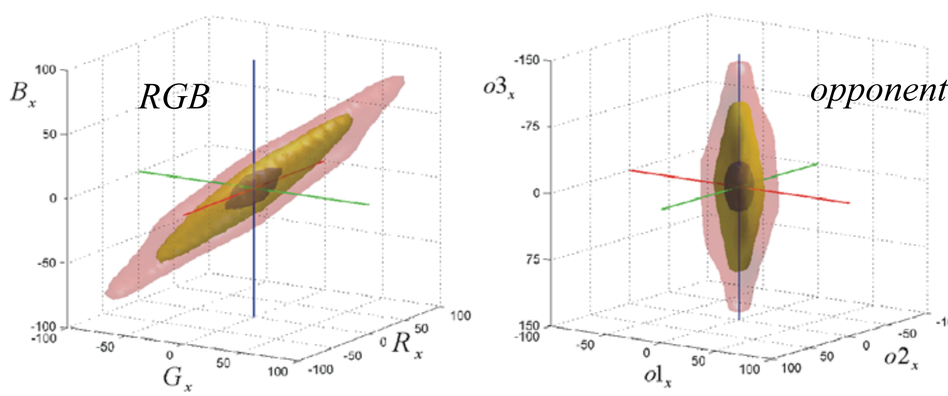

- Isosalient surfaces can be approximated by aligned ellipsoids in decorrelated color spaces.

statistics of color images:

Color Boosting Saliency: $p(\mathbf{f}_x) = p(\mathbf{f}'_x) \leftrightarrow |g(\mathbf{f}_x)| = |g(\mathbf{f}'_x)|$

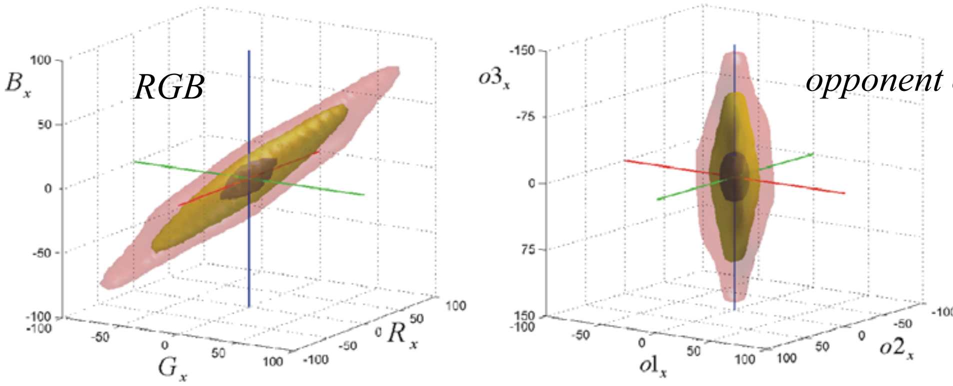

opponent colorspace

Color Boosting function: $$g(\mathbf{f}_x) = \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix} h(\mathbf{f}_x)$$

saliency points

input car-image

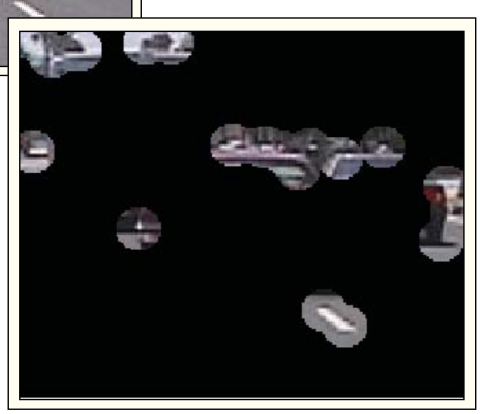

RGB-based (first 20 points)

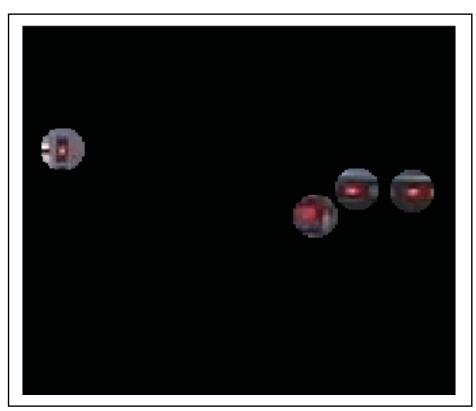

saliency boosting (first 4 points)

generality approach: global optimal regions

RGB gradient

color boosting

experiment: quantitative analysis

Quantitative evaluation of color boosting on a retrieval experiment.

- Nister database: around 10.000 images
- detector: DoG (color boosted)
- descriptor: SIFT+hue

experiment: quantitative analysis

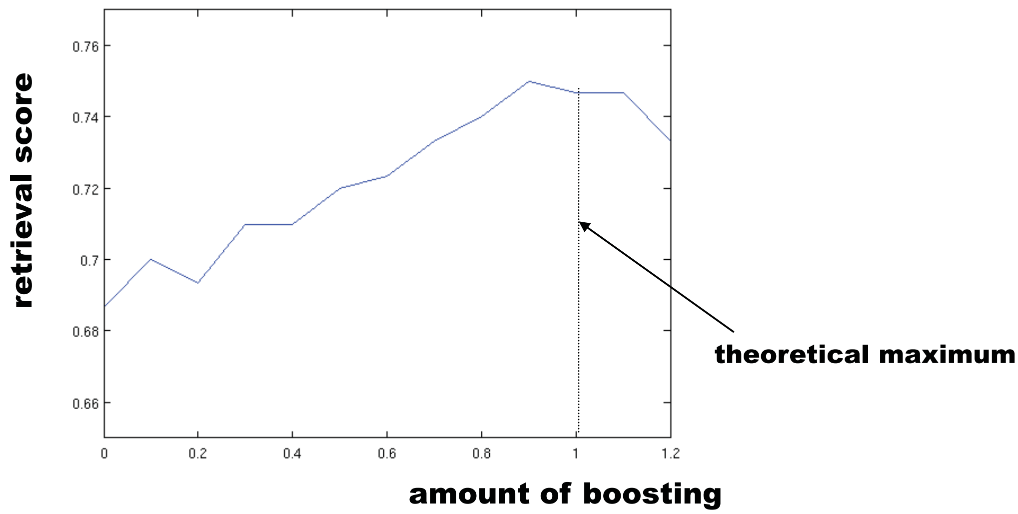

- color boosting improves results between 5-10 percent
- the obtained maximum score is ‘equal’ to the theoretical maximum.

The do’s and dont’s of Color Features

1. Take care in combining different channels:
   Tensor-based features solve the opposing vector problem.
2. Look at what kind of photometric invariance your problem needs:

   Do not take derivatives of circular color spaces.

   Compute first derivatives, then color space transform.

   Quasi-invariants are more stable for feature detection.
3. When working with invariance take instabilities into account.
   Use error analysis to find certainty measures for your invariants.
4. When considering photometric invariance always also take discriminative power into account.
5. From information theory an optimal color space for salient feature detection can be derived.
6. Color information is highly corrupted in compressed data. In compression (jpeg, mpeg) chrominance is subsampled.

overview

PART I (low-level)
Joost van de Weijer

1. ***Reflection Models***
   - Dichromatic reflection model
   - Color Spaces
2. ***Color Differential Structure***
   - Color Edges
   - Photometric Invariant Edge Detection
3. ***Saliency and Color Boosting***
   - Itti and Koch model
   - Color boosted
4. ***Color Constancy***
   - At the pixel
   - Low-level
   - High-level

PART II (high-level)
Theo Gevers

1. ***Interest point detection***
   - Harris Laplace
   - Color boosted
2. ***Descriptors***
   - SIFT
   - Extension to color
3. ***Object recognition (VOC/TRECVid)***
   - Dense and point sampling
   - Code book generation
   - Results
4. ***Applications***
   - Tracking in video
   - Object replacement
   - Emotion recognition
   - Head pose estimation

Color Constancy at a Pixel

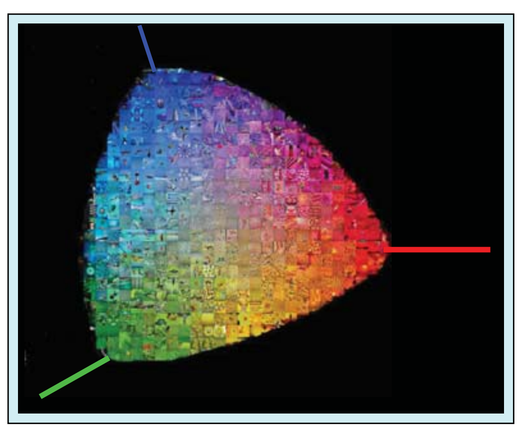

problem statement

How do we recognize colors to be the same under varying light sources ?

color constancy : the ability to recognize colors of objects invariant of the color of the light source.

Colour constancy algorithms

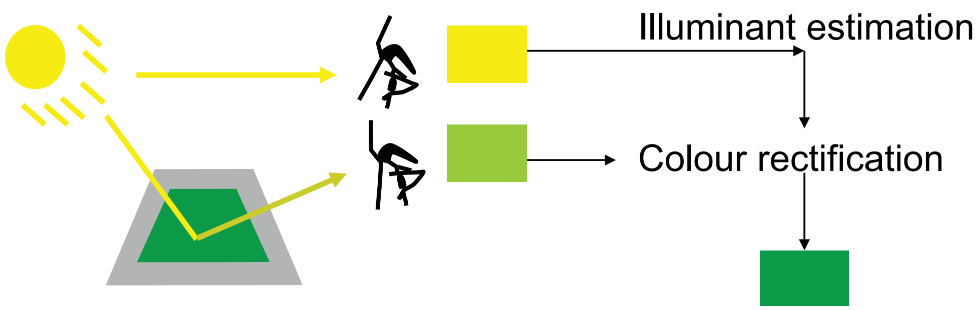

Invariant Normalizations

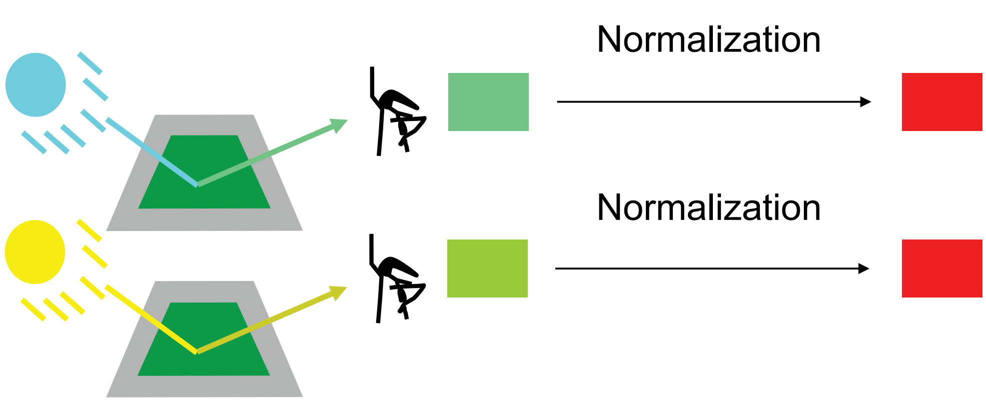

color constancy at a pixel

Assumptions :

1. Lambertian model:
   - linear relation pixel values and intensity light.
   - no specularities and interreflections.
2. perfectly narrow-band sensors (Dirac delta functions).
3. the illuminants are Planckian.

However, the final algorithm is shown to be robust to deviations from the assumptions.

Surface reflectance

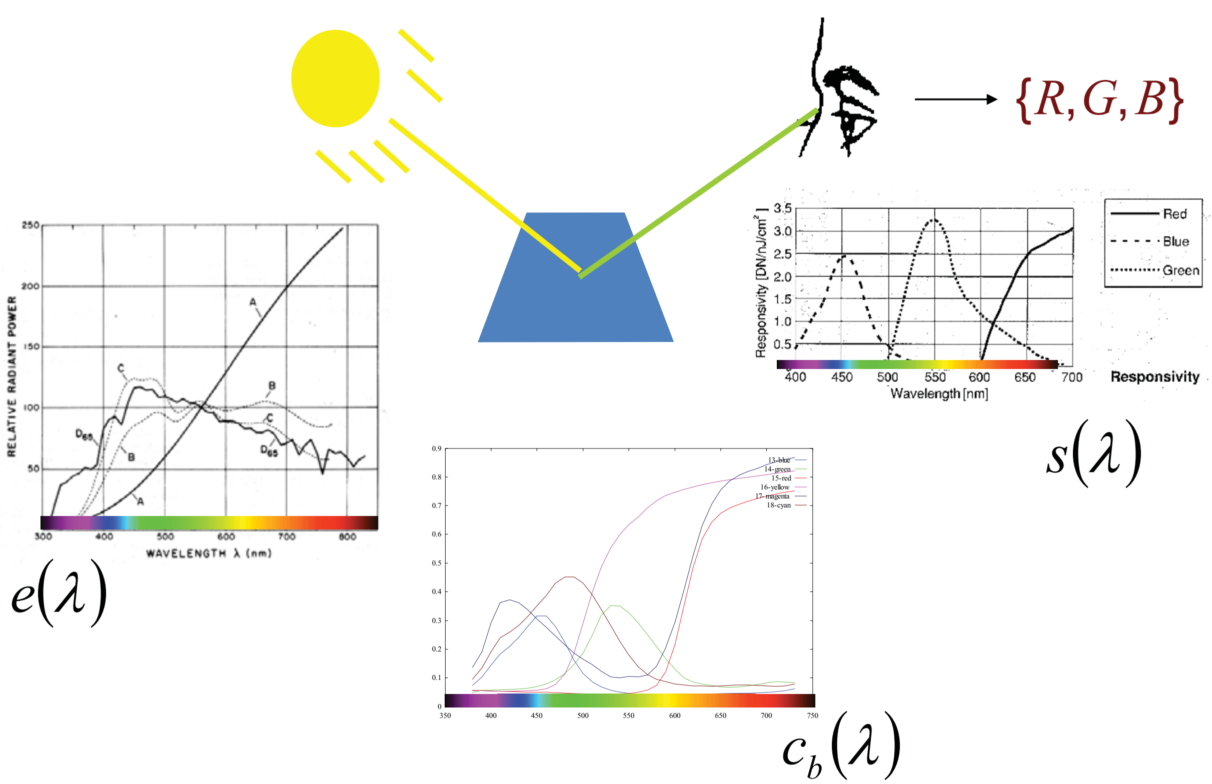

Dirac delta functions

$$p_k = \int_{\omega} e(\lambda) c_b(\lambda) s_k(\lambda) d\lambda$$

assumption: Dirac sensors

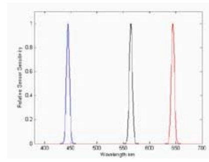

$$p_k = \int_{\omega} e(\lambda) c_b(\lambda) q_k \partial(\lambda - \lambda_k) d\lambda$$

$$\boxed{p_k = e(\lambda_k) c_b(\lambda_k) q_k}$$

Planckian illuminants

Planck's law of black body radiation states the spectral intensity of electromagnetic radiation from a black body at temperature T as a function of wavelength:

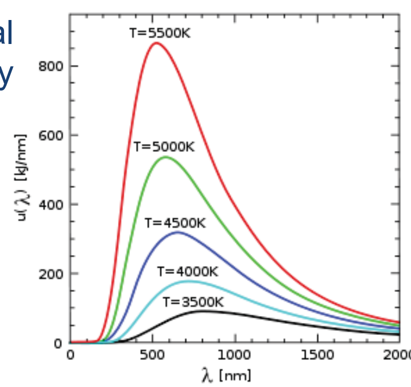

Wien's approx: $E(\lambda,T) = \frac{c_1}{\lambda^5} e^{-\frac{c_2}{T\lambda}}$

The **Planckian locus** is the path that the color of a black body as the blackbody temperature changes.

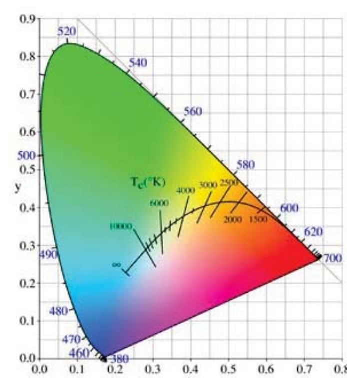

Planckian illuminants

Planck's law of black body radiation states the spectral intensity of electromagnetic radiation from a black body at temperature T as a function of wavelength:

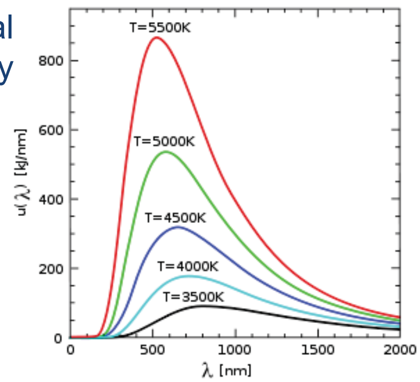

Wien's approx: $E(\lambda, T) = \dfrac{c_1}{\lambda^5} e^{-\frac{c_2}{T\lambda}}$

The **Planckian locus** is the path that the color of a black body as the blackbody temperature changes.

Daylight illuminants can be approximated by Planckian illuminants.

(indoor illuminants to some extend

2500K Household light bulbs

3000K Studio lights, photo floods

4000K Clear flashbulbs

5000K Typical daylight; electronic flash)

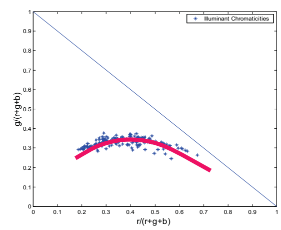

Color constancy at a pixel

$$p_k = e(\lambda_k) c_b(\lambda_k) q_k \xrightarrow{\text{Planckian light}} p_k = \frac{c_1}{\lambda_k^5} e^{-\frac{c_2}{T\lambda}} c_b(\lambda_k) q_k$$

Consider the logarithm of the chromaticity coordinates:

$$\chi_j = \log\left(\frac{p_k}{p_p}\right) = \log\left(\frac{\lambda^{-5} e^{-\frac{c_2}{T\lambda}} c_b(\lambda_k) q_k}{\lambda^{-5} e^{-\frac{c_2}{T\lambda}} c_b(\lambda_p) q_p}\right)$$

$$\chi = \mathbf{s} + \frac{1}{T}\mathbf{e} \Leftarrow \chi_j = \log\left(\frac{s_k}{s_p}\right) + \frac{1}{T}(e_k - e_p)$$

depends on surface color

depends on illuminant color

$e_k \equiv -c_2 / \lambda_k$

$s_k = \lambda_k^{-5} c_b(\lambda_k) q_k$

color constancy at a pixel - examples

examples log chromaticity plots:

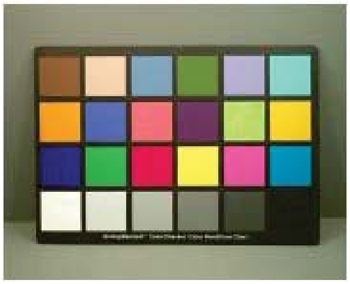

Macbeth Color Checker

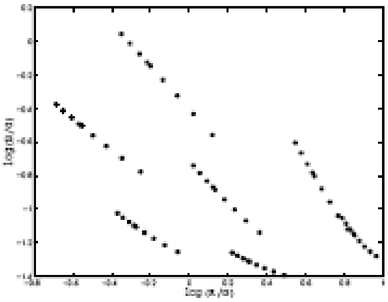

HP912 Digital Still Camera

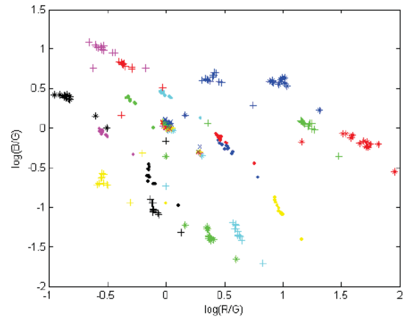

Nikon D-100

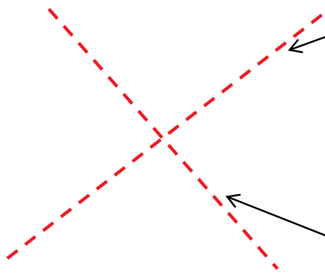

illuminant invariant direction axis

Every pixel can be represented in a illuminant invariant representation !

illuminant variant axis (camera dependent)

images source: Eli Arbel

examples illuminant invariant

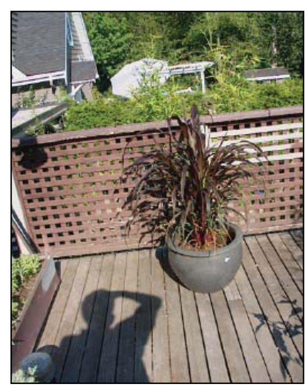

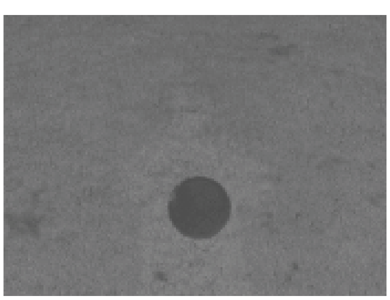

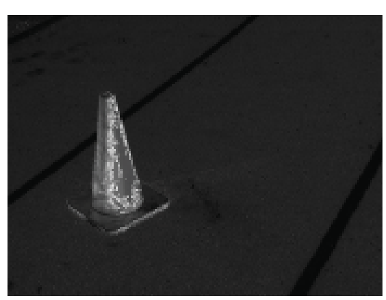

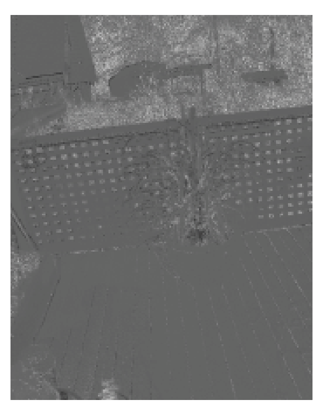

Since shadows are a change in illuminant these representation are shadow free.

shadow detection

Comparison of the edge maps of the original and the shadow invariant image allows for shadow detection.

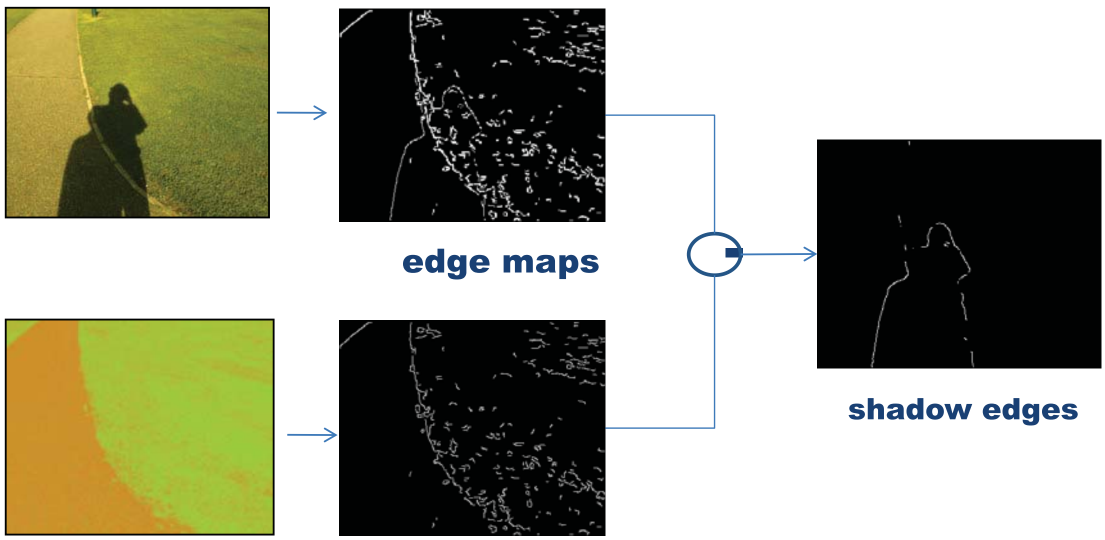

examples:

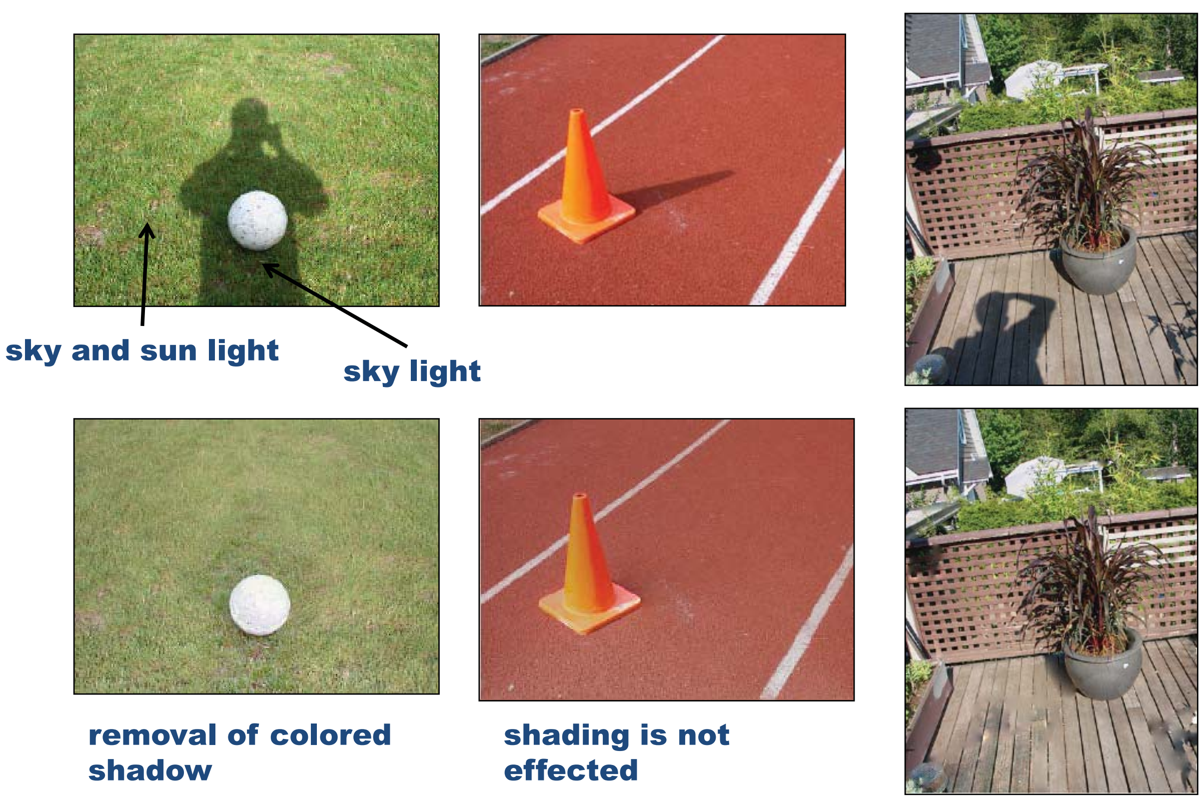

references:

1. B. H. Tenenbau. *Recovering intrinsic scene characteristics from images.* **Computer Vision Systems, 1978**.
2. Y. Weiss. *Deriving intrinsic images from image sequences*. **ICCV 2001**.
3. G. D. Finlayson, S.D. Hordley. *Color Constancy at a Pixel*. **JOSA 2001**.
4. G.D. Finlayson, S.D. Hordley, C. Lu, M.S. Drew, *On the reomoval of shadows from images*. **PAMI 2006**.
5. E. Arbel, H Hel-Or*, Texture-Preserving Shadow Removal in color Images Containing Curved Surfaces*. **CVPR 2007**.
6. F. Liu, M. Gleicher. *Texture-Consistent Shadow Removal.* **ECCV 2008**.

Gamut Mapping

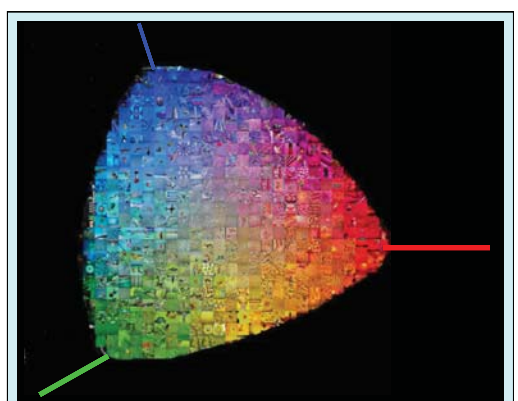

regular gamut mapping

"In real-world images, for a given illuminant, one observes only a limited number of different colors."

Solux 4700K

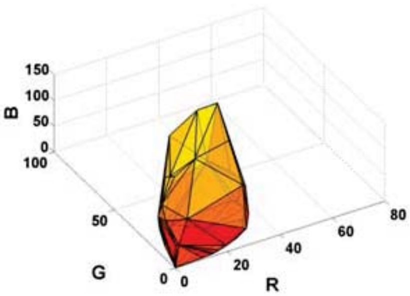

Solux 4700K + Roscolux filter

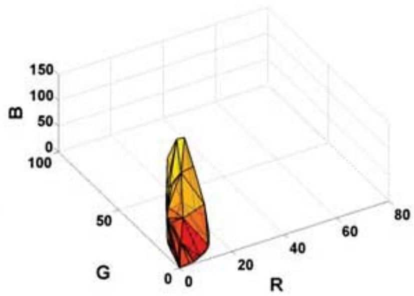

Sylvania Warm White Fluorescent

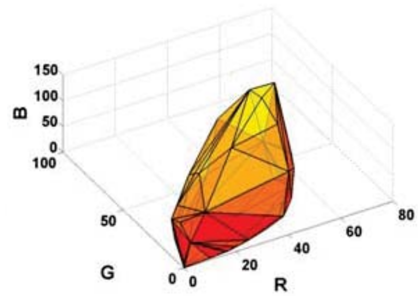

regular gamut mapping

Gamut mapping algorithm:

- Obtain input image.

regular gamut mapping

Gamut mapping algorithm:

- Obtain input image.
- Compute gamut from image.

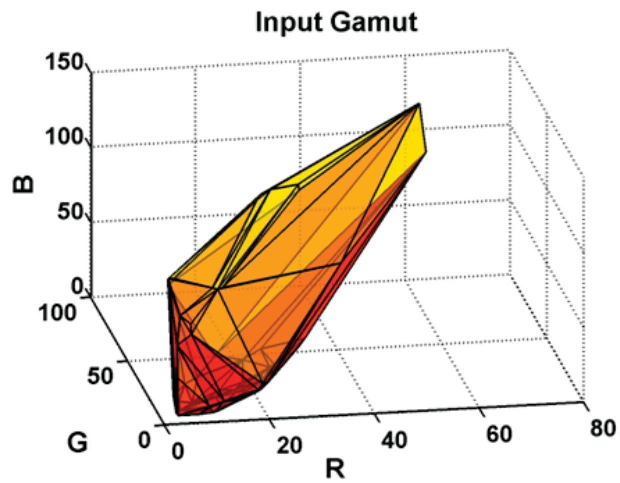

regular gamut mapping

Gamut mapping algorithm:

- Obtain input image.
- Compute gamut from image.
- Determine feasible set of mappings from input gamut to canonical gamut.

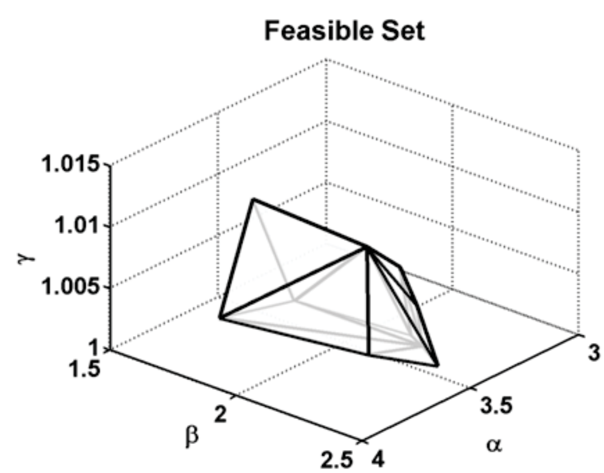

regular gamut mapping

Gamut mapping algorithm:

- Obtain input image.
- Compute gamut from image.
- Determine feasible set of mappings from input gamut to canonical gamut.
- Apply some estimator, to select one mapping from this set.

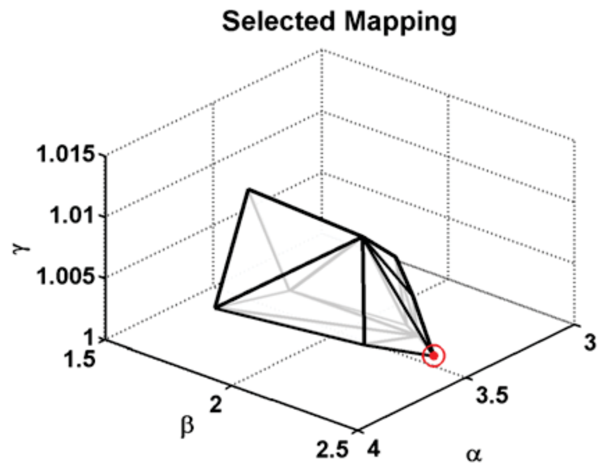

regular gamut mapping

Gamut mapping algorithm:

- Obtain input image.
- Compute gamut from image.
- Determine feasible set of mappings from input gamut to canonical gamut.
- Apply some estimator, to select one mapping from this set.
- Use mapping on input image to recover the corrected image, or on canonical illuminant to estimate the color of the unknown illuminant.

Color Constancy from Color Derivatives

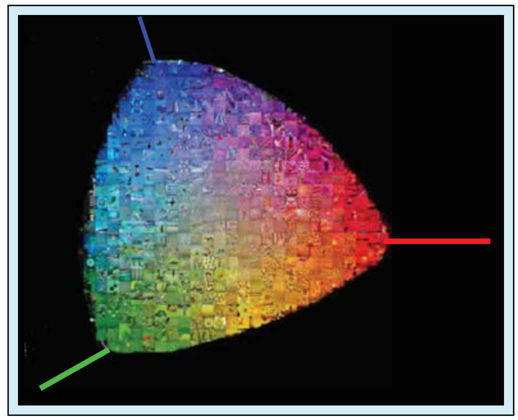

Color Constancy

color constancy : the ability to recognize colors of objects invariant of the color of the light source.

Grey world hypothesis : the average reflectance in a scene is grey.

White patch hypothesis : the highest value in the image is white.

Grey-world: $\sum_{m=1}^{M} \mathbf{f}_i(\mathbf{x}) \propto \mathbf{c}$

white-patch: $\left(\sum_{m=1}^{M} (\mathbf{f}_i(\mathbf{x}))^{\infty} \right)^{\frac{1}{\infty}} \propto \mathbf{c}$

Shades of Grey hypothesis : the n-Minkowsky norm based average of a scene is achromatic.

- unifies Grey-World and White Patch : $e^p \approx \sqrt[p]{\int |\mathbf{f}(\mathbf{x})|^p d\mathbf{x}}$

Color Constancy

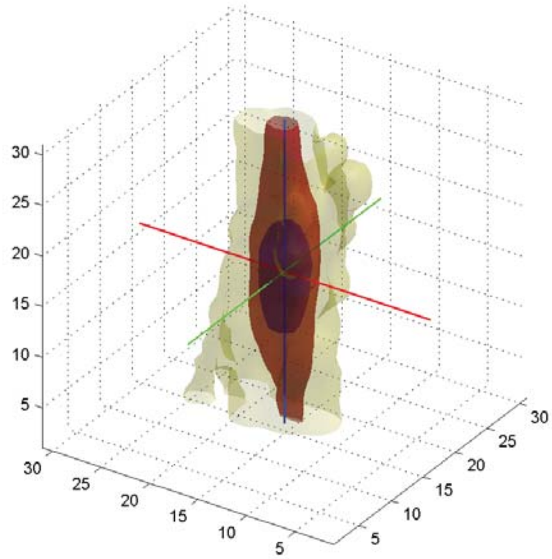

Color Constancy

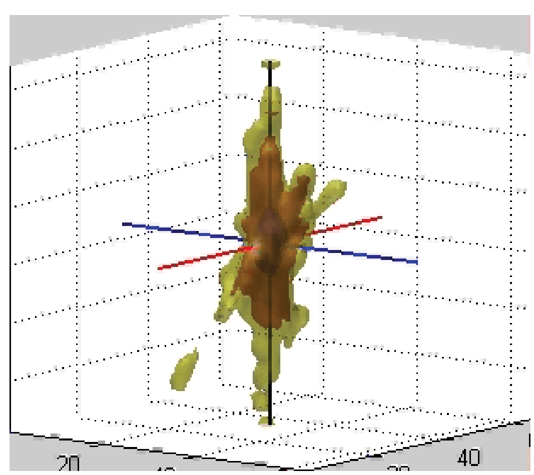

Color Constancy

color constancy : the ability to recognize colors of objects invariant of the color of the light source.

Grey world hypothesis : the average reflectance in a scene is grey.

White patch hypothesis: the highest value in the image is white.

generalization I: the L-norm: $\left(\sum_{m=1}^{M} \left(\mathbf{f}_i(\mathbf{x}) \right)^k \right)^{\frac{1}{k}} \propto \mathbf{c}$

Grey edge hypothesis : the average edge in a scene is grey.

generalization II: L-norm + differentiation order: $\left(\sum_{i=1}^{M} \left| \frac{\partial^n \mathbf{f}_i(\mathbf{x})}{\partial \mathbf{x}^n} \right|^p \right)^{\frac{1}{p}} \propto \mathbf{c}$

Color Constancy in 4 lines of matlab code !

```
function Illuminant=GreyEdgeCC(im,mink,sigma,dif)

im  =  gauss_derivative(im,sigma,dif);
im  = reshape(im,size(im,1)*size(im,2),3);
Illuminant= 1./power( sum ( power( im, mink) ), 1/mink );
Illuminant = Illuminant./norm(Illuminant) ;
```

general color constancy framework

Low-level color constancy:

$$\left(\sum_{i=1}^{M}\left|\frac{\partial^n \mathbf{f}_i(\mathbf{x})}{\partial \mathbf{x}^n}\right|^p\right)^{\frac{1}{p}} \propto \mathbf{c}$$

$n=0, p=1$

$$\sum_{m=1}^{M}\mathbf{f}_i(\mathbf{x}) \propto \mathbf{c}$$

grey-world

$n=0, p=\infty$

$$\left(\sum_{m=1}^{M}\left(\mathbf{f}_i(\mathbf{x})\right)^{\infty}\right)^{\frac{1}{\infty}} \propto \mathbf{c}$$

white-patch

$n=0, p=k$

$$\left(\sum_{m=1}^{M}\left(\mathbf{f}_i(\mathbf{x})\right)^{k}\right)^{\frac{1}{k}} \propto \mathbf{c}$$

shades-of-gray

$n=1, p=1$

$$\left(\sum_{m=1}^{M}\left|\frac{\partial \mathbf{f}_i(\mathbf{x})}{\partial \mathbf{x}}\right|^p\right)^{\frac{1}{p}} \propto \mathbf{c}$$

grey-edge

G. Finlayson, E. Trezzi, “Shades of gray and colour constancy”, *CIC 2004*
J. van de Weijer, T. Gevers “Edge-Based Color Constancy”, *IEEE IP 2007*

Color Constancy: experiment

- test set: 23 objects under 11 illuminants (Computational Vision Lab: Simon Fraser)

- $\text{angular error} = \cos(\hat{e} \cdot e)$

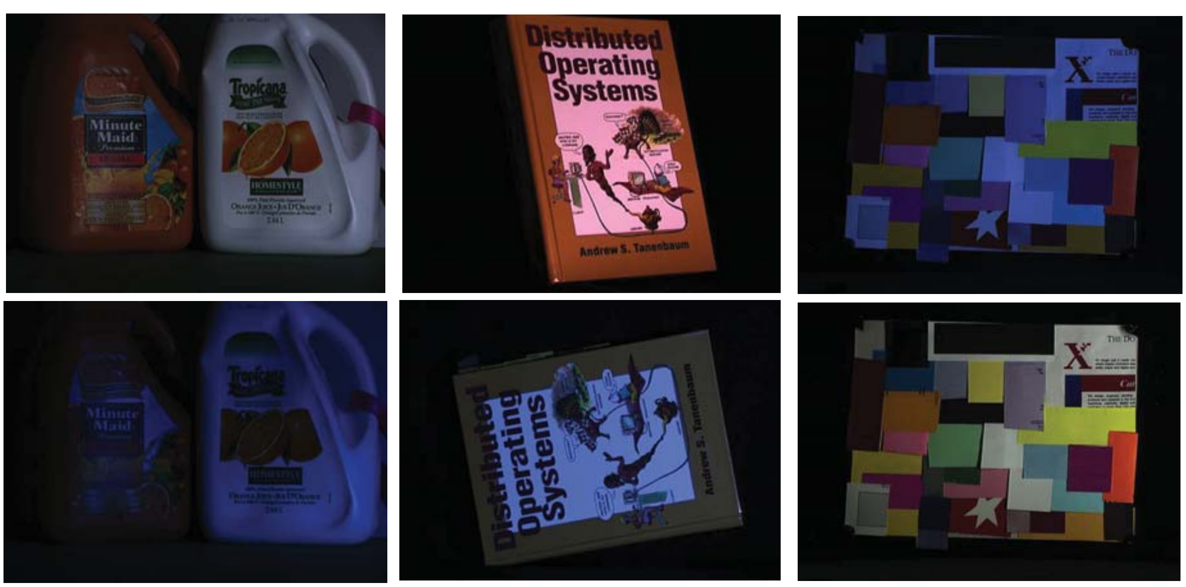

Color Constancy: experiment

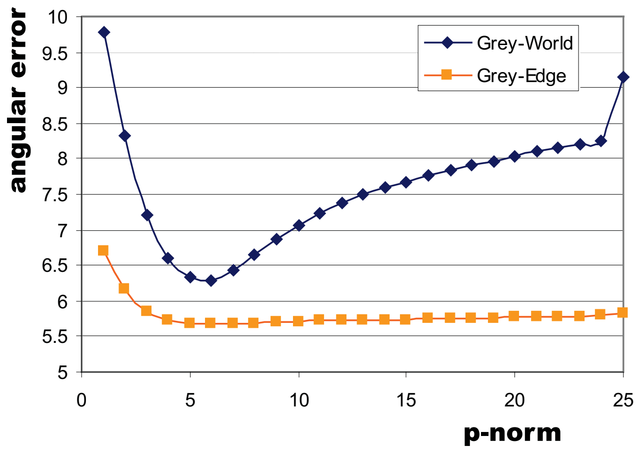

	error
Grey-World	9.8
White-Patch	9.2
General Grey-World	5.4
Grey-Edge	5.6
2nd order Grey-Edge	5,2
Color by Correlation	9,9
Gamut Mapping	5,6
GCIE, 11 Lights	4,9
GCIE, 87 Lights	5,3

Color Constancy: experiment

- real-world data set (F. Ciurea and B. Funt : Vision Lab - Simon Fraser)

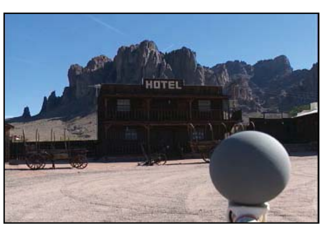

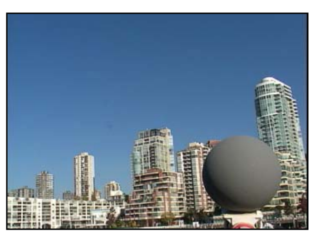

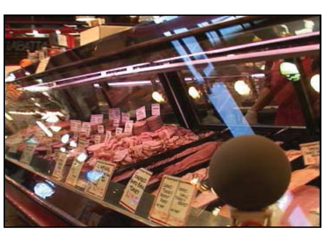

Color Constancy: experiment

• real-world data set (F. Ciurea and B. Funt : Vision Lab - Simon Fraser)

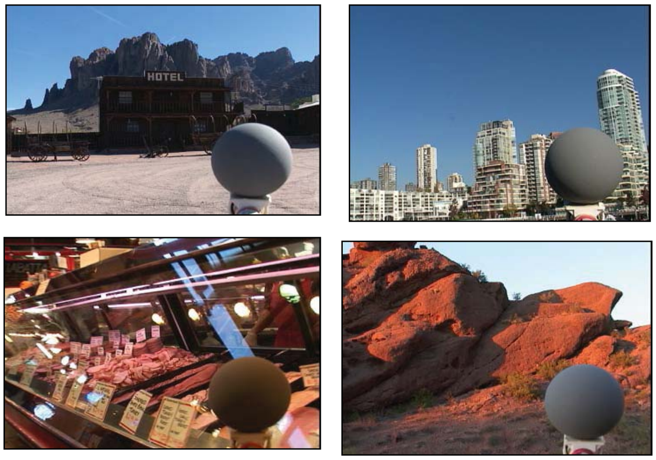

	median
Grey-World	7.3
White-Patch	6.7
General Grey-World	4.7
Grey-Edge	4.1
2nd order Grey-Edge	4.3

derivative-based gamut mapping

"In real-world images, for a given illuminant, one observes only a limited number of different colored edges."

Solux 4700K

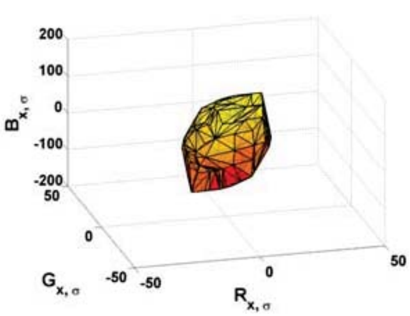

Solux 4700K + Roscolux filter

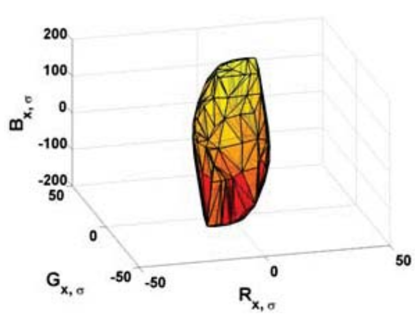

Sylvania Warm White Fluorescent

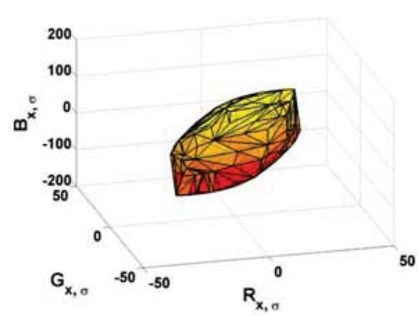

A. Gijsenij, T. Gevers, J. van de Weijer, "Edge-Based Color Constancy", *IJCV 2010*

Experiments (real-world images)

Some examples:

Original	Ideal	Derivative-based	Regular Gamut

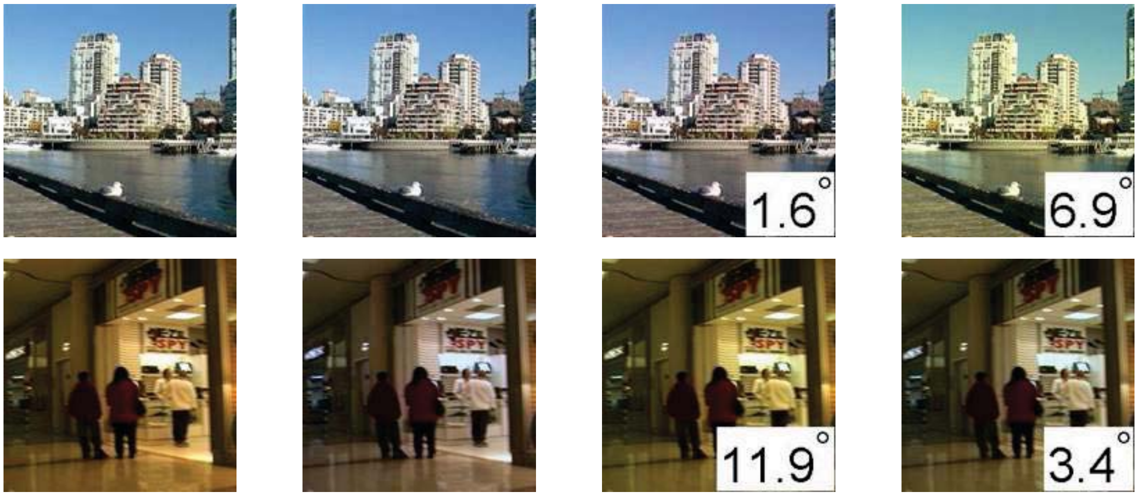

How do you choose the best cc-algorithm ?

High-Level Color Constancy

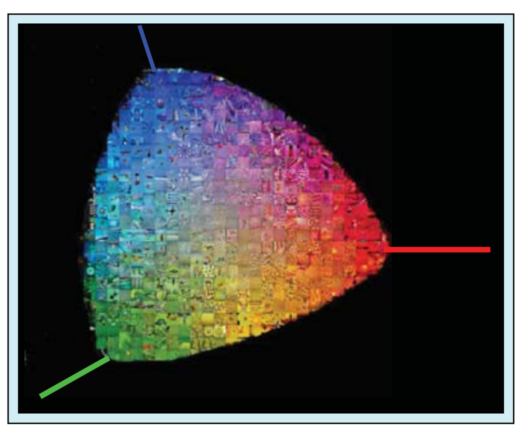

Natural Image Statistics

- Could it be that different scenes prefer different color constancy methods ?

Geusebroek and Smeulders (2005) – Weibulls

Examples:

slide credit: Arjan Gijsenij

Natural Image Statistics

Distribution of edge responses follows Weibull distribution.

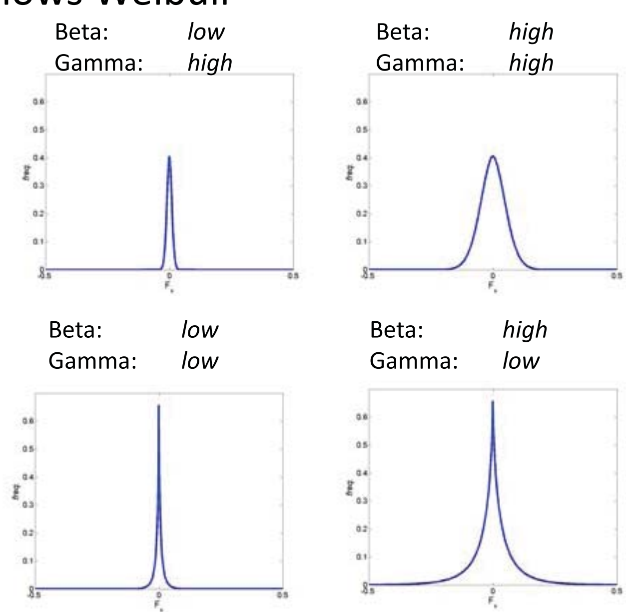

Two parameters:

- β – Contrast of the image. A higher value indicates more contrast.
- γ - Grain size. A higher value indicates more fine textures.

slide credit: Arjan Gijsenij

Color Constancy – Selection

Postsupervised Prototype Classification:

Compute Weibull-parameters for all images

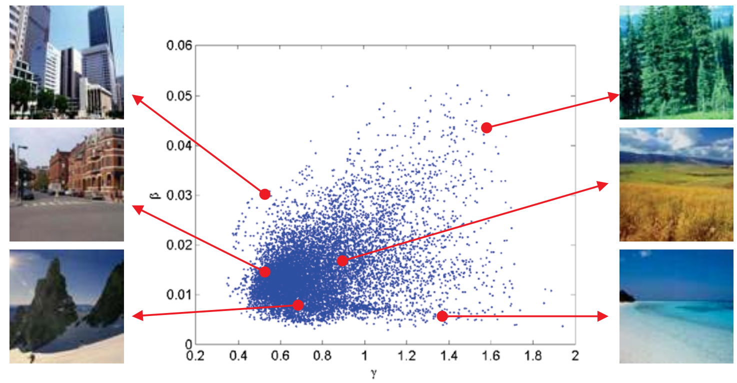

slide credit: Arjan Gijsenij

Color Constancy – Selection

Postsupervised Prototype Classification:

Compute Weibull-parameters for all images

Partition weibull-parameters using *k*-means

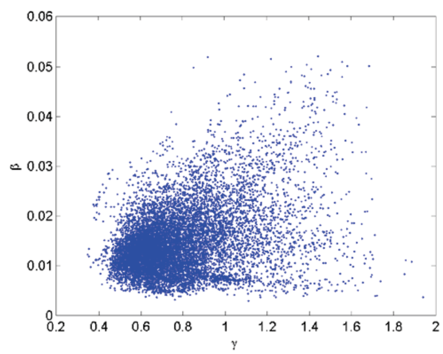

slide credit: Arjan Gijsenij

Color Constancy – Selection

Postsupervised Prototype

Classification :

Compute Weibull-parameters for all images

Partition weibull-parameters using *k*-means

Label cluster centers according to the minimum mean angular error

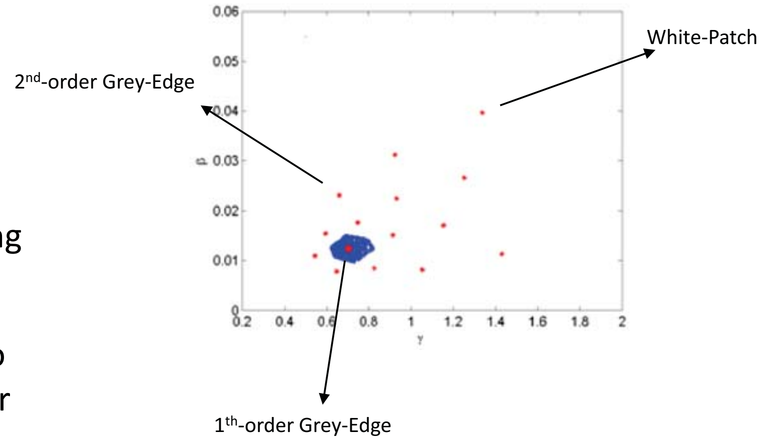

slide credit: Arjan Gijsenij

Color Constancy – Selection

Postsupervised Prototype

Classification :

Compute Weibull-parameters for all images

Partition weibull-parameters using *k*-means

Label cluster centers according to the minimum mean angular error

Build 1-NN Classifier on these cluster centers

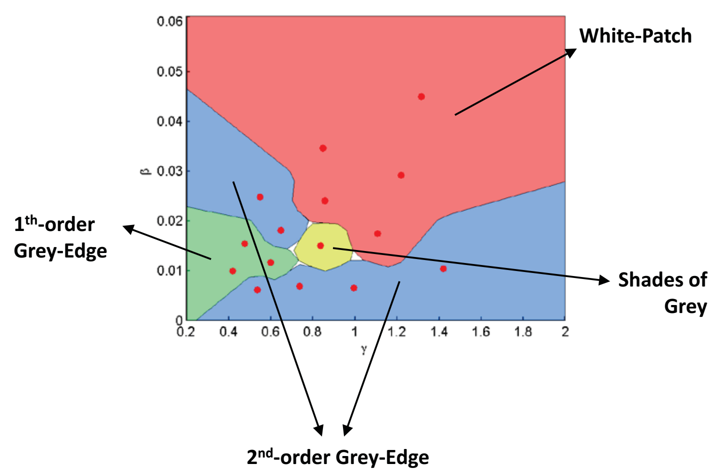

slide credit: Arjan Gijsenij

Experiments

Data set consisting of 11000+ images

The *true* illuminants are known (ground truth)

Grey sphere is *masked during experiments*

Performance measure → **angular error:**

$$\cos^{-1}(\hat{\mathbf{e}}_l \cdot \hat{\mathbf{e}}_e)$$

slide credit: Arjan Gijsenij

Experiments – Results

Original	Ideal	Selection	White-Patch	Grey-World

slide credit: Arjan Gijsenij

Experiments – Performance

Method	Mean	Median
Grey-World	7.9°	7.0°
White-Patch	6.8°	5.3°
General Grey-World	6.2°	5.3°
1th-Order Grey-Edge	6.2°	5.2°
2nd-Order Grey-Edge	**6.1°**	**5.2°**
Gamut mapping	8.5°	6.8°
Color-by-Correlation	6.4°	5.2°

slide credit: Arjan Gijsenij

Experiments – Performance

Method	Mean	Median
2nd-Order Grey-Edge (baseline)	6.1°	5.2°
Selection – 5 methods	5.7° (-7%)	4.7° (-10%)
Combining – 5 methods	5.6° (-8%)	4.6° (-12%)
Combining – 75 methods	**5.0°(-18%)**	**3.7° (-29%)**

slide credit: Arjan Gijsenij

Color Constancy from High-Level Visual Information

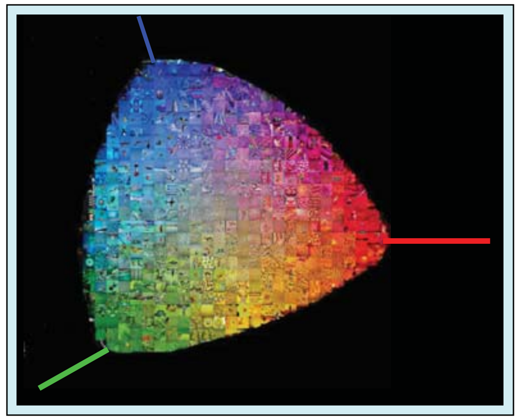

problem statement

How do we recognize colors to be the same under varying light sources ?

color constancy : the ability to recognize colors of objects invariant of the color of the light source.

computational color constancy

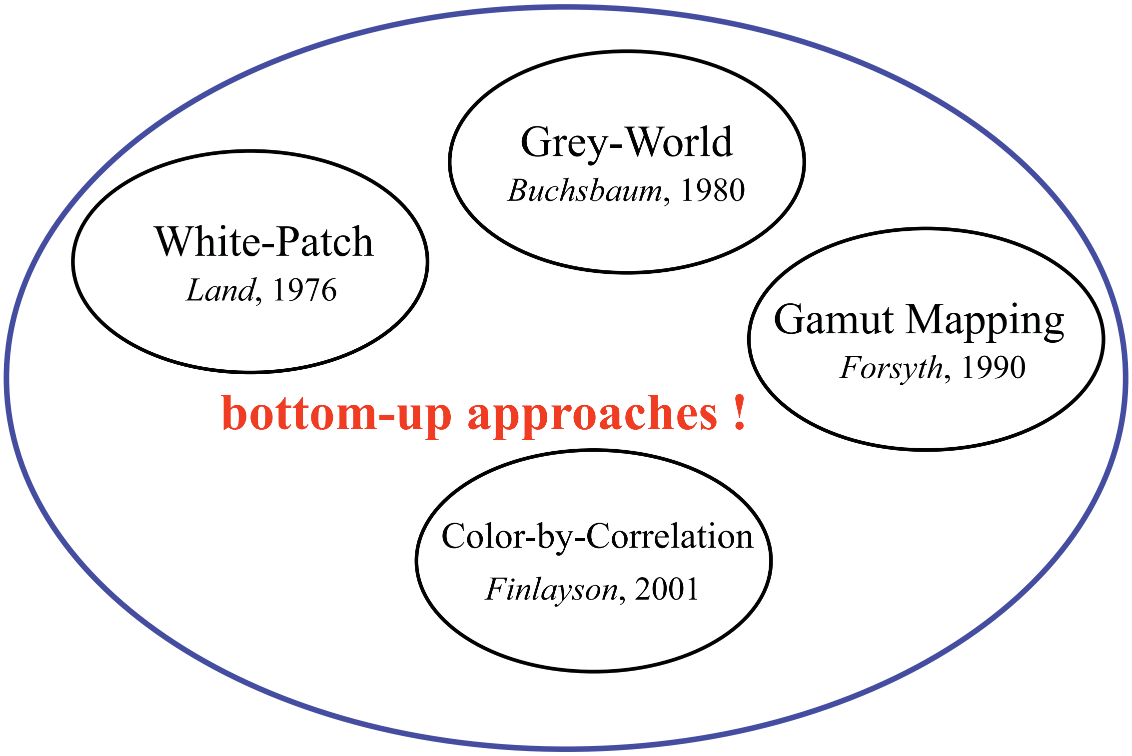

top-down color constancy

psychophysical motivation:

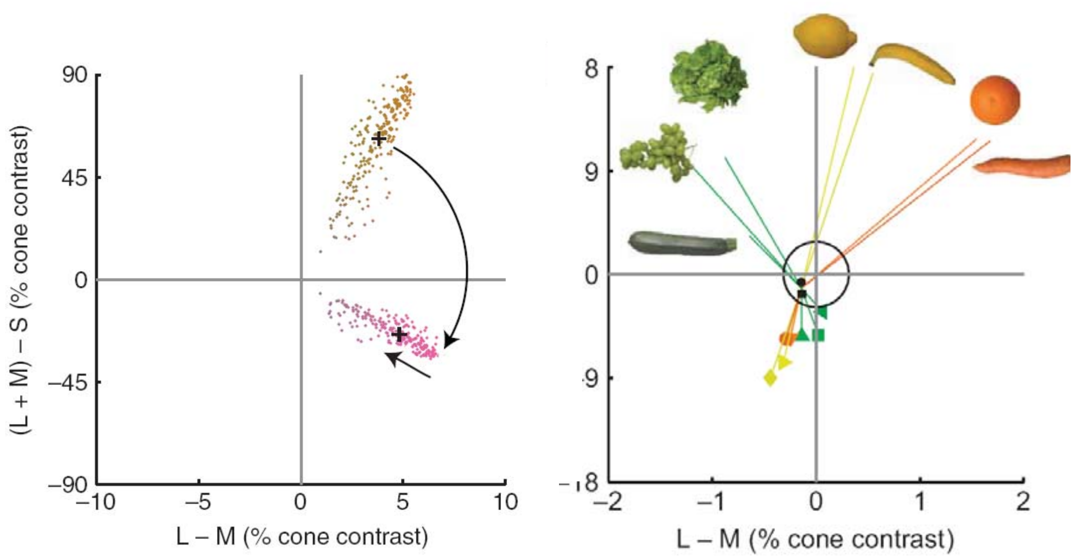

Hansen et al. "Memory modulates color appearance", *nature neuroscience*, 2006.

problem statement

How do we recognize colors to be the same under varying light sources ?

color constancy : the ability to recognize colors of objects invariant of the color of the light source.

How can we apply high-level visual information for computational color constancy ?

overview our approach

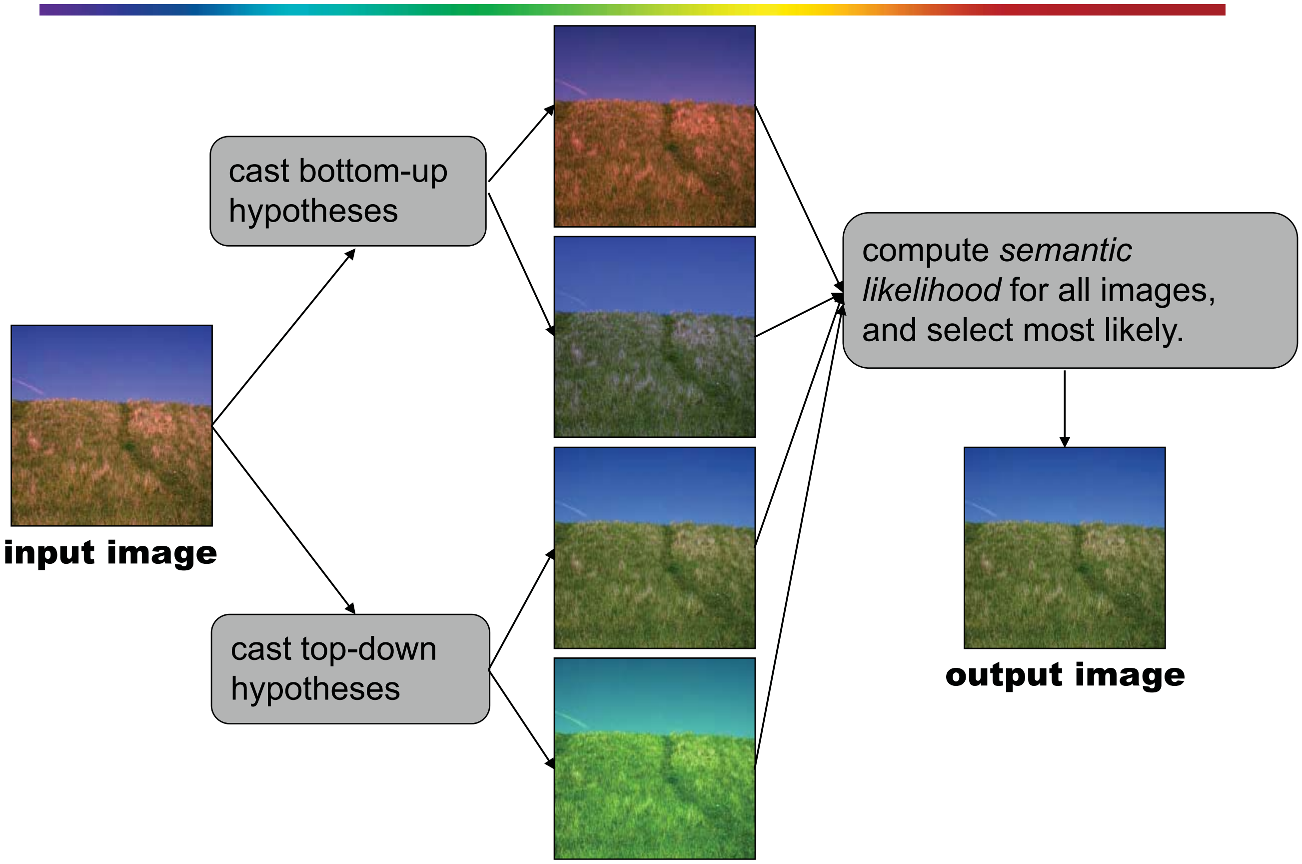

plsa-based image segmentation

- We use Probabilistic Latent Semantic Analysis (pLSA) to compute the semantic likelihood of an image.

Image representation

- dense extraction of 20x20 pixel patches on 10x10 pixel grid
- each patch described by discretized features, the words .
  - texture: SIFT (750 visual words, k-means)
  - color: hue (100 visual words, k-means)
  - position: patch location indicated by cell in a 8x8 grid

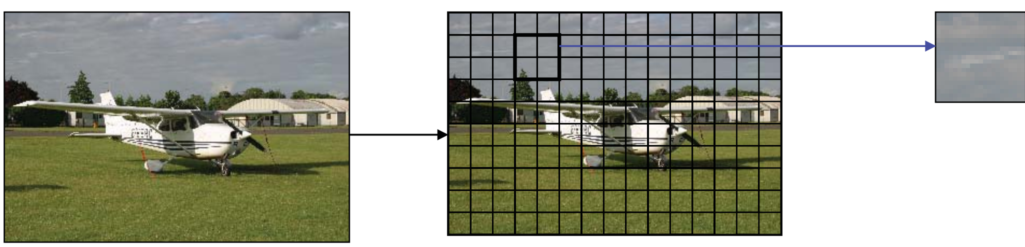

plsa-based image segmentation

- We use Probabilistic Latent Semantic Analysis (pLSA) to compute the semantic likelihood of an image.

An image is modeled as a mixture of semantic topics:

$$p(w|d) = \sum_z p(w|z)\, p(z|d)$$

visual word **image** **semantic topics** **image-specific mixture proportions**

$$p(w|z) = \prod_{m=1}^{M} p(w^m|z)$$

The $p(w^m|z)$ can either be learned supervised or unsupervised.
We assume them to be learned from images taken under a white illuminant.

likelihood image $p(d) = \prod_w p(w|d)$

plsa-based image segmentation

supervised learning

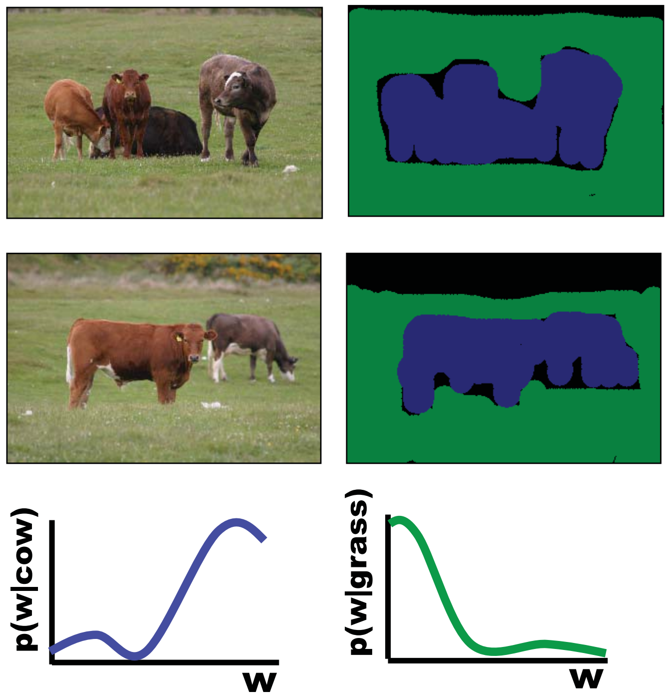

p(w|z)

test image

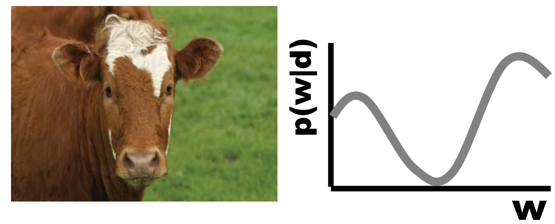

$$p(w|d)=\sum_z p(w|z)\,p(z|d)$$

unknown

using EM: p(z|d)={0.6,0.4}

semantic image segmentation

plsa-based image segmentation

unsupervised learning

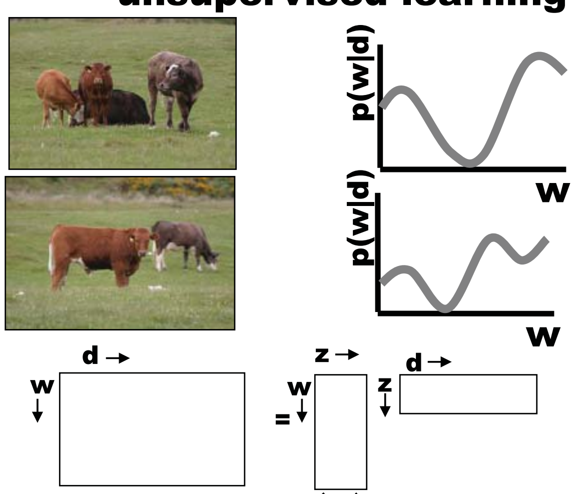

$$p(w|d)=\sum_z p(w|z)\,p(z|d)$$

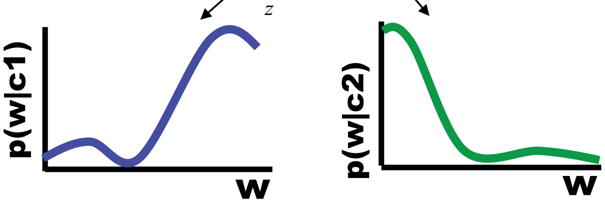

p(w|z)

test image

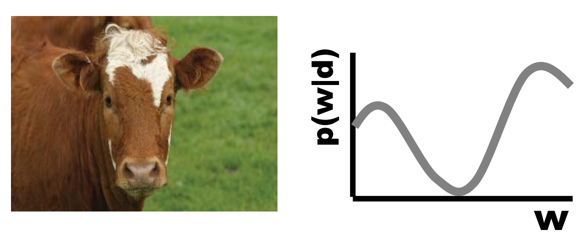

$$p(w|d)=\sum_z p(w|z)\,p(z|d)$$

unknown

using EM: p(z|d)={0.6.0.4}

semantic image segmentation

semantic likelihood image

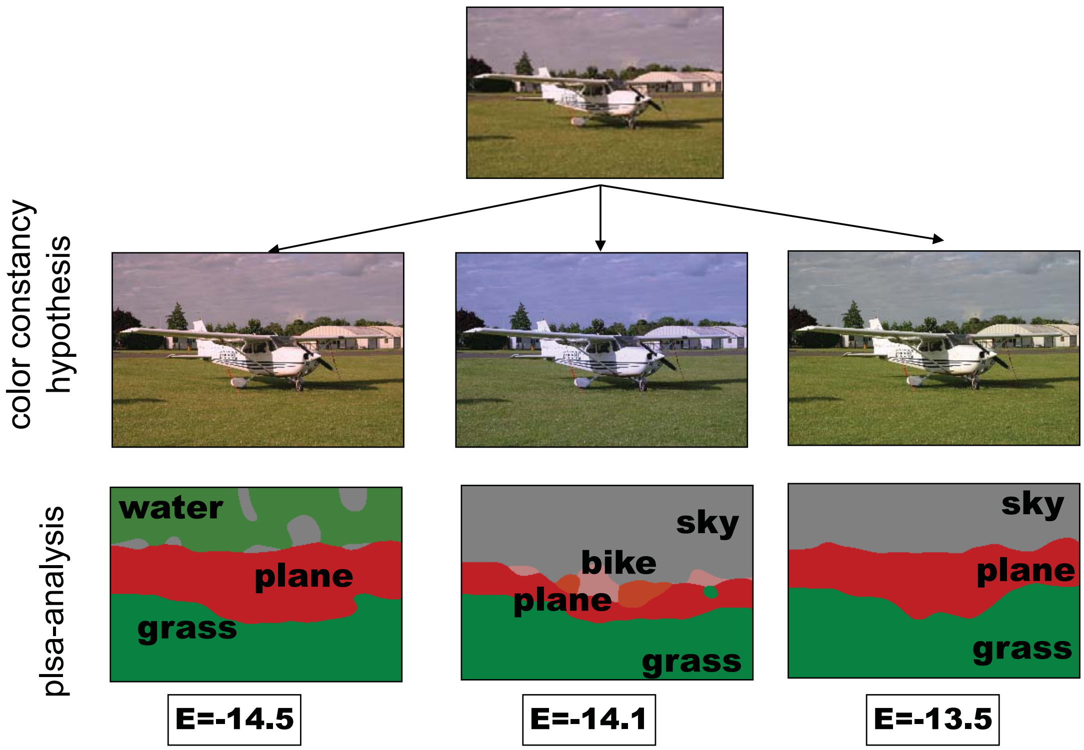

casting hypotheses: bottom-up

Low-level color constancy:

$$\left(\sum_{i=1}^{M} \left| \frac{\partial^n \mathbf{f}_i(\mathbf{x})}{\partial \mathbf{x}^n} \right|^p \right)^{\frac{1}{p}} \propto \mathbf{c}$$

$n=0, p=1$

$$\sum_{m=1}^{M} \mathbf{f}_i(\mathbf{x}) \propto \mathbf{c}$$

grey-world

white-patch

shades-of-gray

$n=1, p=1$

$$\left(\sum_{m=1}^{M} \left| \frac{\partial \mathbf{f}_i(\mathbf{x})}{\partial \mathbf{x}} \right|^p \right)^{\frac{1}{p}} \propto \mathbf{c}$$

grey-edge

We will use n={0,1,2} and p={2,12} to cast a total of 6 bottom-up hypotheses.

G. Finlayson, E. Trezzi, "Shades of gray and colour constancy", ***CIC 2004***
J. van de Weijer, T. Gevers "Edge-Based Color Constancy", ***IEEE TIP 2007***

casting hypotheses: top-down

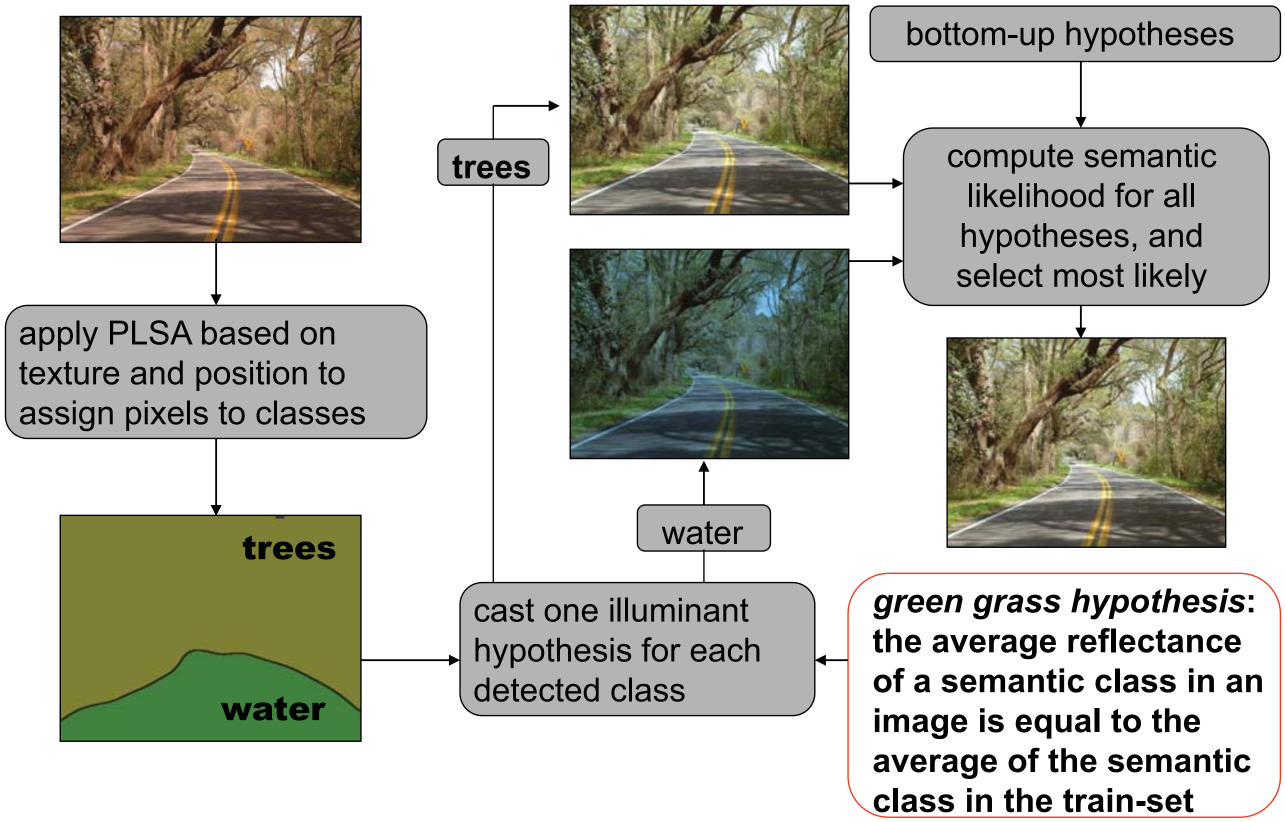

experiment: illuminant estimation

Data Set contains both indoor and outdoor scenes from a wide variety of locations (150 training, 150 testing)

Topic-word distributions are learned unsupervised on the texture and position cue (color is ignored in training).

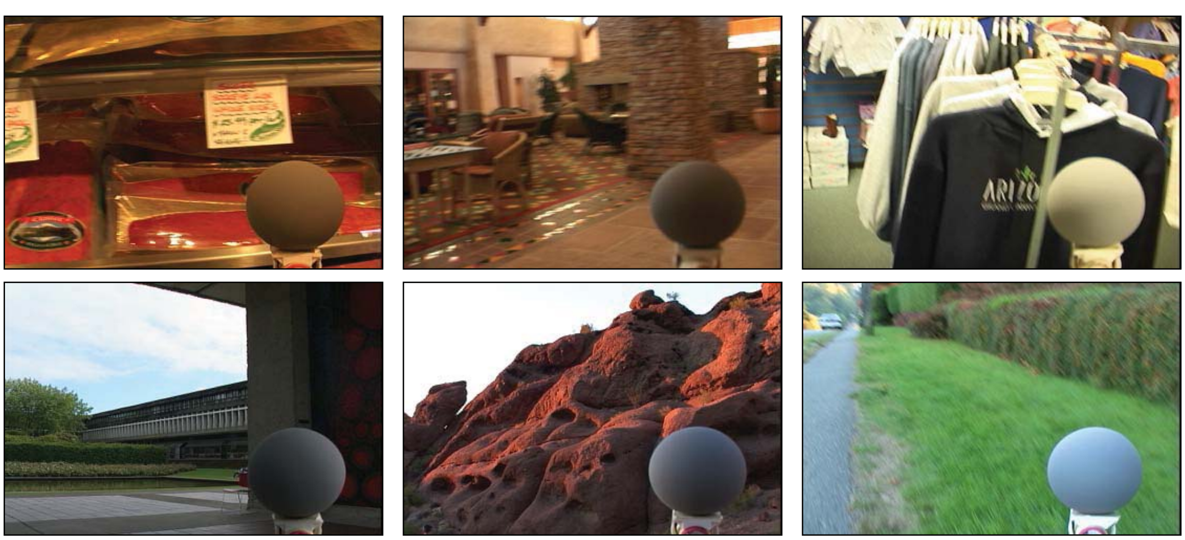

F. Ciurea and B. Funt "A large database for color constancy research", ***CIC 2004.***

experiment: illuminant estimation

results in angular error:

		standard color constancy		high-level selection		
	no cc	worst BU	best BU	BU	TD	BU & TD
indoor	12.8	12.3	6.1	5.3	5.6	5.3
outdoor	5.5	7.4	4.9	4.7	4.7	4.5

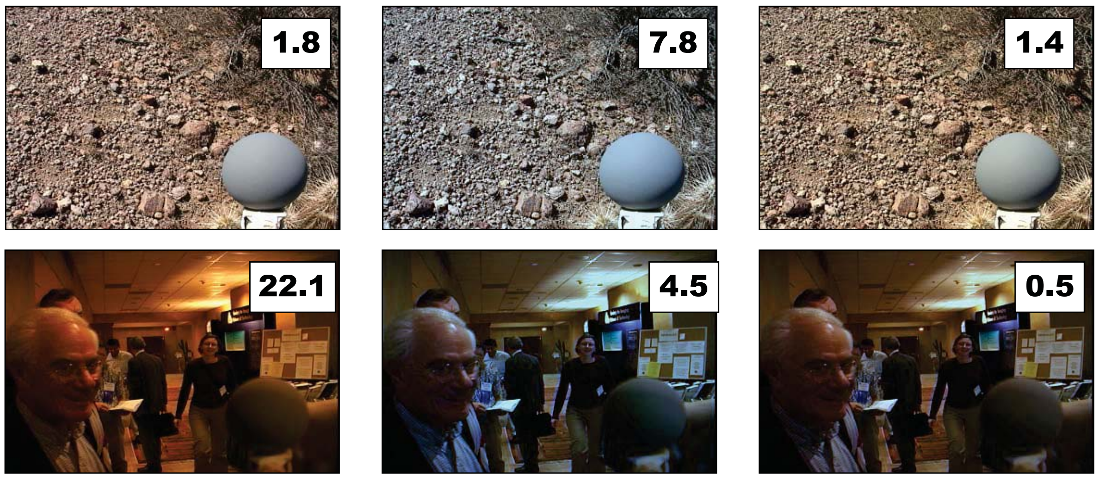

input image bottom-up top-down

experiment: semantic segmentation

Data Set training: labelled images of Microsoft Research Cambridge (MSRC) set, together with ten images collected from Google Image for each class. Traning: 350 images. Test : 36 images.

Topic-word distributions are learned supervised.

Classes: building, grass, tree, cow, sheep, sky, water, face and road.

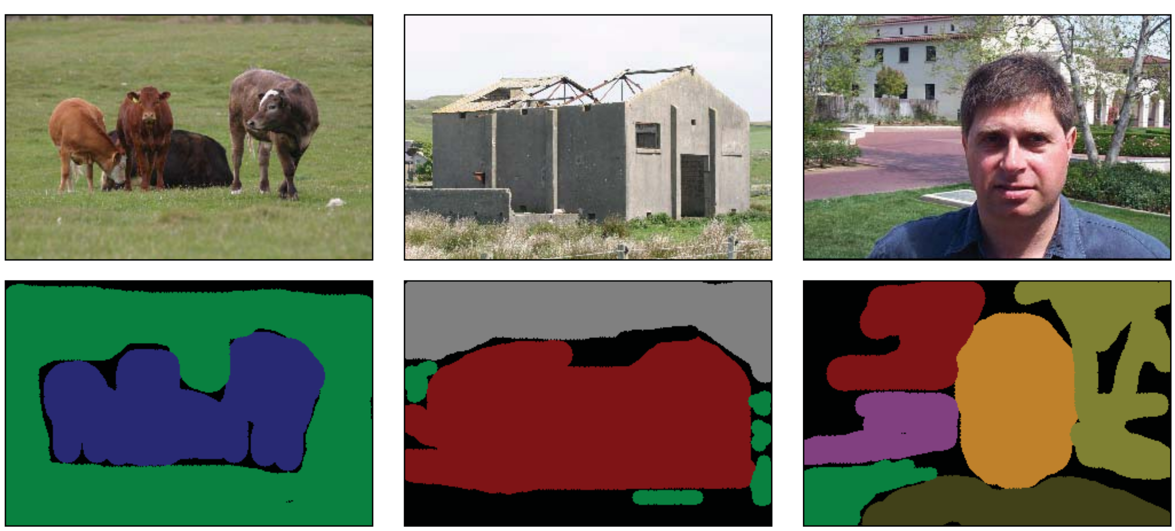

J. Shotton et al. "Textonboost", ***ECCV 2006.***

experiment: pixel classification

results pixel classification in %:

	standard color constancy		high-level selection		
no cc	worst BU	best BU	BU	TD	BU & TD
39.6	41.4	52.2	53.4	59.5	64.2

Summary Color Constancy

• The Planckian locus describes natural light illuminants.

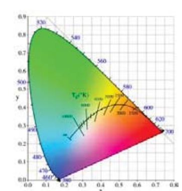

• Color constancy at the pixel allows for shadow removal.

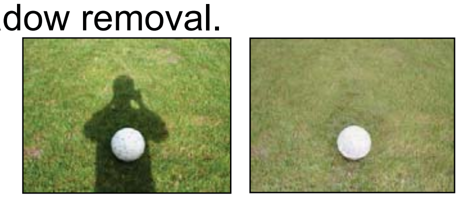

•The general grey-world algorithm generalizes a set of low-level color constancy algorithms, including white patch, grey-world, grey-edge, and shades –of-grey.

$$\left(\sum_{i=1}^{M} \left| \frac{\partial^n \mathbf{f}_i(\mathbf{x})}{\partial \mathbf{x}^n} \right|^p \right)^{\frac{1}{p}} \propto \mathbf{c}$$

• Top-down information improves both color constancy performance and semantic segmentation results.

references: color constancy

- D.A. Forsyth, "*A novel algorithm for color constancy*." IJCV, 1990.
- G.D. Finlayson, M.S. Drew, B.V. Funt, "*Color by correlation: A simple, unifying framework for color constancy*", PAMI 2001.
- K. Barnard, L. Martin, B.V. Funt, "*A comparison of computational color constancy algorithms-part II: Experiments with data*" IEEE transactions on Image Processing, 2002.
- G.D. Finlayson, S.D. Hordley, and I. Tastl. "*Gamut constrained illuminant estimation*", ICCV'03.
- G.D. Finlayson and E. Trezzi. "*Shades of gray and colour constancy*", IS&T/SID, CIC'04.
- J. van de Weijer, Th. Gevers, A. Gijsenij, "*Edge-Based Color Constancy*", TIP 2005.
- A. Gijsenij, T. Gevers, "*Color Constancy using Natural Image Statistics*",CVPR 2006.
- A. Chakrabarti, K. Hirakawa, T. Zickler, "*Color Constancy Beyond Bags of Pixels"*, CVPR 2008.
- A. Gijsenij, T. Gevers, J. van de Weijer, "Edge-Based Color Constancy", *IJCV 2010.*

acknowledgements:

Amsterdam ISLA	: Theo Gevers, Jan-Mark Geusebroek, Arnold Smeulders, Arjan Gijsenij.
INRIA Rhone-Alpes	: Cordelia Schmid, Jakob Verbeek, Marcim Marszalek.
CVC Barcelona	: Maria Vanrell, Ramon Baldrich, Juan Toledo.

Color Naming

learning color names

task: Object colors in many images are often not explicitly labeled. Can we label these image automatically with color names ?

Ebay user: “Find me all yellow cars ?”

learning color names

From linguistic studies it is known that the development of color names follows a similar pattern for all languages.

Development color names in languages:

white | red | green | blue | brown | purple
black | | yellow | | | pink
| | | | | orange
| | | | | grey

The english language has 11 basic color terms.

learning color names

- **Use google image to assemble a set of weekly labeled images.**

black blue brown green orange purple white

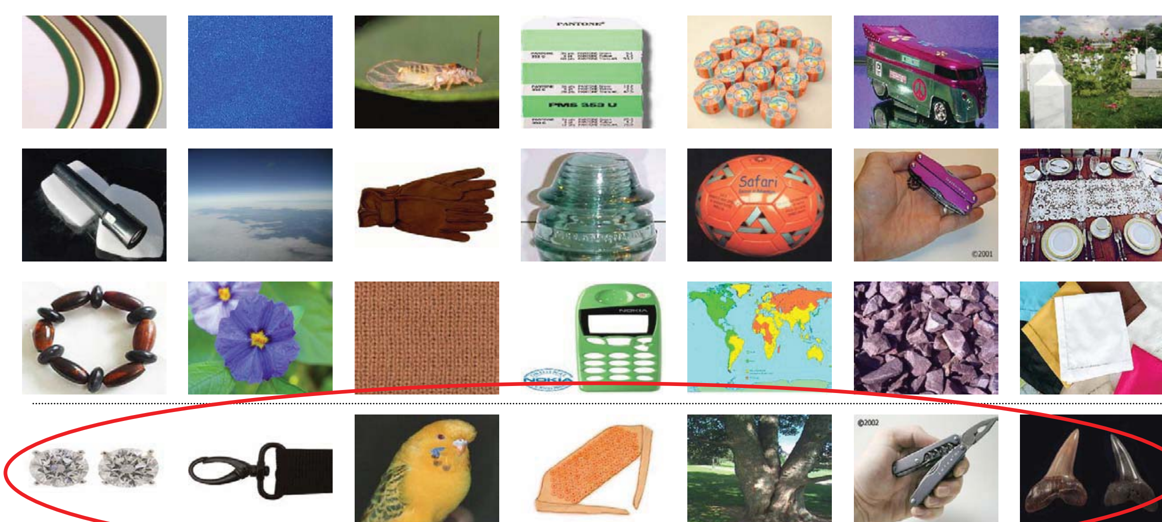

false positives **Images retrieved with Google image**

learning color names

Labeled input images:

yellow

yellow

...

red

LAB-histogram representation:

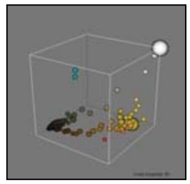

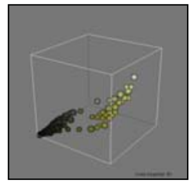

...

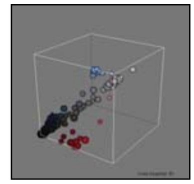

PLSA-bg

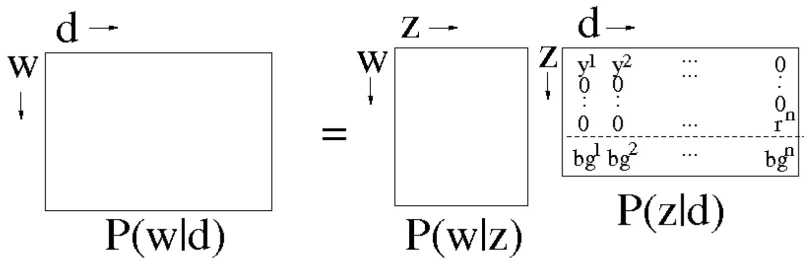

Color name distribution:

yellow

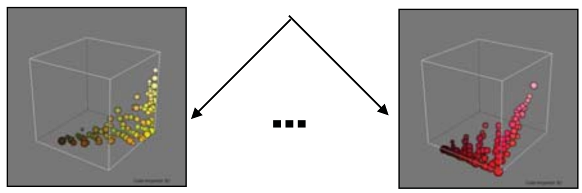

... red

learning color names

task: Object colors in many images are not explicitly labeled.
Can we label these image automatically with color names ?

Ebay user: "Find me all yellow cars ?"

Result automatic labeling pixels:

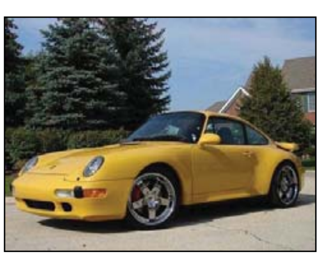

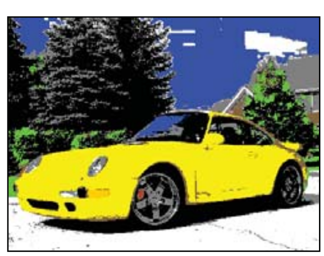

Example: classification soccer data

- Achromatic colors are very abundant in the world, about 45 % (67 % with brown) .

black	blue	brown	grey	green	orange	pink	purple	red	white	yellow
19	12	23	19	10	2	2	2	4	6	1

statistics based 40.000 corel images.

- when using photometric invariance always consider discriminative power.

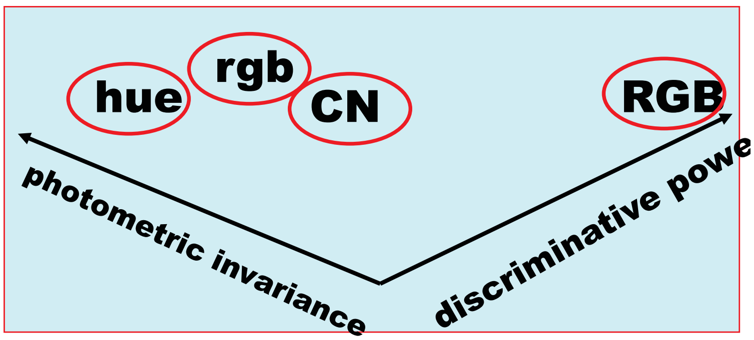

Example: classification soccer data

- Achromatic colors are very abundant in the world, about 45 % (more than 60 % with brown) .

black	blue	brown	grey	green	orange	pink	purple	red	white	yellow
19	12	23	19	10	2	2	2	4	6	1

statistics based 40.000 corel images.

- when using photometric invariance always consider discriminative power.

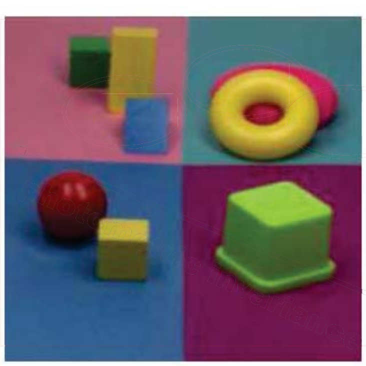
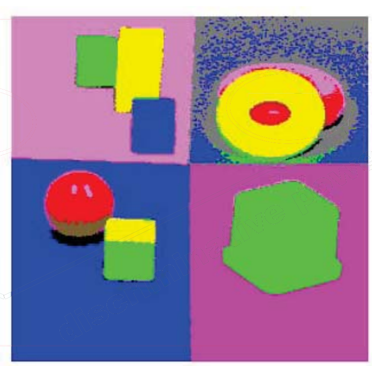

Results flower data set:

- **test color names for image classification on a flower data set of 1360 images over 17 classes.**




dataset		**flower**	
method	color	shape	color & shape
HSV-SIFT	-	-	78
hue	40	65	79
opponent	39	65	79
color names	57	65	81

references: color naming

- B. Berlin, P. Kay. *Basic Color terms: their universality and evolution.* Berkeley: University of California, 1969.
- A. Mojsilovic. *A computational model for color naming and describing color composition of images.* IEEE TIP 14(5), 2005.
- K. Yanai, K. Barnard, Image region entropy: a measure of *visualness* of web images associated with on concept, ACMMM 2005.
- R. Benavente, M.Vanrell, R. Baldrich. *Parametric fuzzy sets for automatic color naming*, JOSA 25(10), 2008.
- J. van de Weijer, Cordelia Schmid, Jakob Verbeek, Diane Larlus. *Learning Color Names for Real-World Applications*. IEEE TIP 2009.

Blur Robust and Color Constant Image description

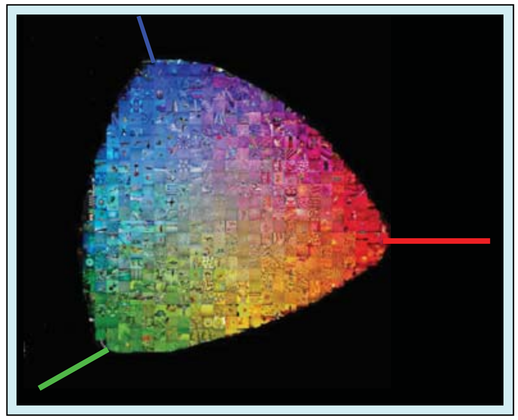

problem statement

How do we recognize colors to be the same under varying light sources ?

$$\begin{pmatrix} R' \\ G' \\ B' \end{pmatrix} = \begin{pmatrix} \alpha & 0 & 0 \\ 0 & \beta & 0 \\ 0 & 0 & \gamma \end{pmatrix} \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

color constancy : the ability to recognize colors of objects invariant of the color of the light source.

Change of illuminant can be modeled by the *diagonal model*.

Colour constancy algorithms

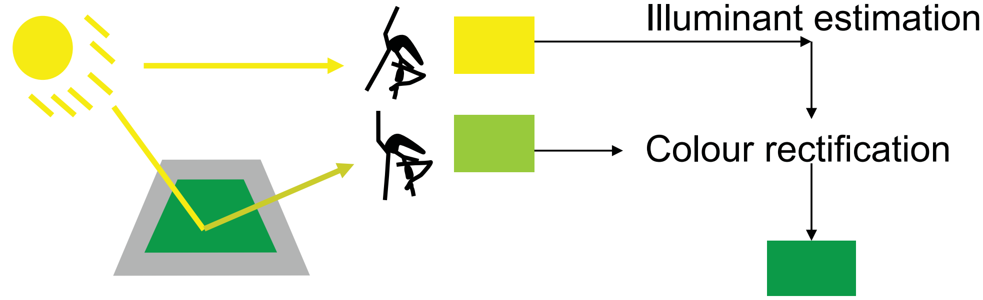

Invariant Normalizations

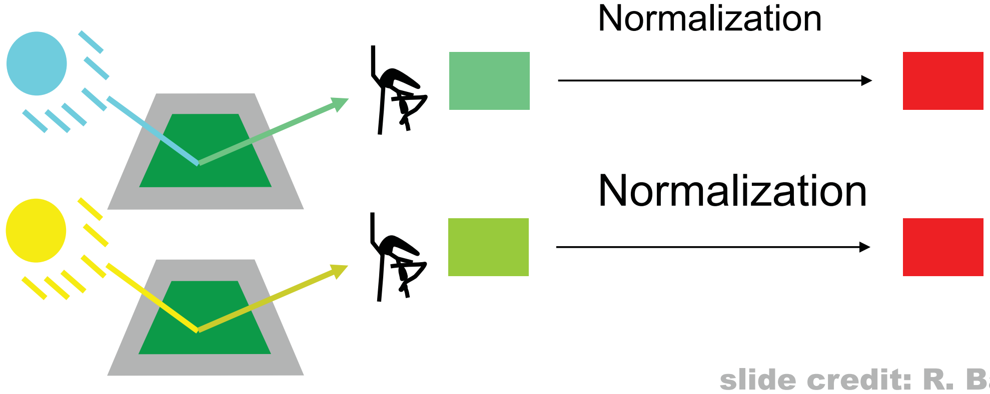

slide credit: R. Baldrich

Color Constant Derivatives

- A color constant representation of a single color patch is impossible.
- The difference between two color patches can be represented invariant to the color illuminant.

Funt and Finlayson:

Mondrian-world: $\mathbf{f}(\mathbf{x}) = m^b \mathbf{c}^b(\mathbf{x})\mathbf{e}$

$$p = \frac{R^1}{R^2} = \frac{m^b {c_1}^R e^R}{m^b {c_2}^R e^R} = \frac{{c_1}^R}{{c_2}^R}$$

$$\ln p = \ln\frac{R^1}{R^2} = \ln R^1 - \ln R^2 = \tfrac{\partial}{\partial x}\ln R$$

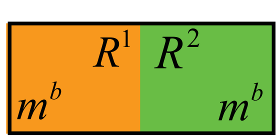

Gevers and Smeulders:

3D-world: $\mathbf{f}(\mathbf{x}) = m^b(\mathbf{x})\mathbf{c}^b(\mathbf{x})\mathbf{e}$

$$m = \frac{R^1G^2}{R^2G^1} = \frac{{m_1}^b {c_1}^R e^R}{{m_2}^b {c_2}^R e^R}\frac{{m_2}^b {c_2}^G e^G}{{m_1}^b {c_1}^G e^G} = \frac{{c_1}^R {c_2}^G}{{c_2}^R {c_1}^G}$$

$$\ln m = \ln\frac{R^1G^2}{R^2G^1} = \ln\frac{R^1}{G^1} - \ln\frac{R^2}{G^2} = \tfrac{\partial}{\partial x}\ln\frac{R}{G}$$

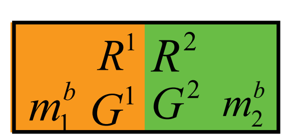

Changing Light Source

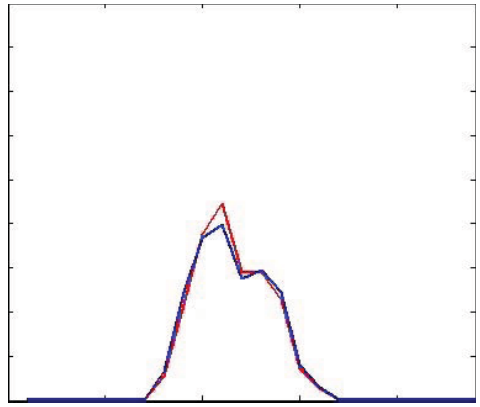

Histogram color constant edge description

Changing Blur & Light Source

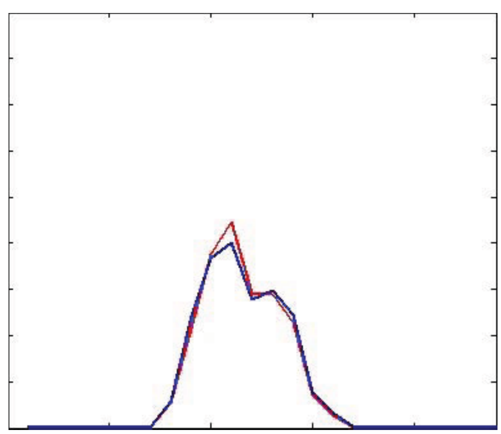

Histogram color constant edge description

Why is this a problem ?

• Image blur is frequently encountered phenomenon.
• Possible causes are : out-of-focus, relative motion between camera and object, and aberrations of the optical system.

Obtaining Invariance to Image Blur

• A color constant representation of a single color patch is impossible.
• The difference between two color patches can be represented invariant to the color illuminant.

Funt and Finlayson:

Mondrian-world: $\mathbf{f}(\mathbf{x}) = m^b \mathbf{c}^b \mathbf{e}(\mathbf{x})$

$$p = \frac{R^1}{R^2} = \frac{m^b {c_1}^R e^R}{m^b {c_2}^R e^R} = \frac{{c_1}^R}{{c_2}^R}$$

$$\ln p = \ln \frac{R^1}{R^2} = \ln R^1 - \ln R^2 = \tfrac{\partial}{\partial x} \ln R$$

Consider a blurred image: $R' = R \otimes G^{\sigma_s}$

$$\tfrac{\partial}{\partial x}^{\sigma_d} \ln R = \frac{R_x^{\sigma_d}}{R^{\sigma_d}} \qquad \tfrac{\partial}{\partial x}^{\sigma} \ln R' = \frac{R_x^{\sqrt{\sigma_d^2+\sigma_s^2}}}{R^{\sqrt{\sigma_d^2+\sigma_s^2}}}$$

On the edge the following holds:

$$R^{\sqrt{\sigma_s^2}} = R^{\sqrt{\sigma_d^2+\sigma_s^2}} \qquad R_x^{\sqrt{\sigma_d^2}} = C(\sigma_s) R_x^{\sqrt{\sigma_d^2+\sigma_s^2}}$$

Robustness with respect to blur is obtained by:

$$\varphi_p^1 = \arctan\left(\frac{R_x G}{G_x R}\right) \qquad \varphi_p^1 = \arctan\left(\frac{G_x B}{B_x G}\right)$$

Color Constancy & Blur Robust

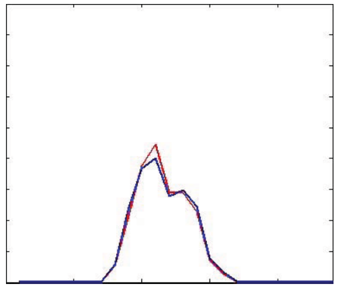

color constant edge description

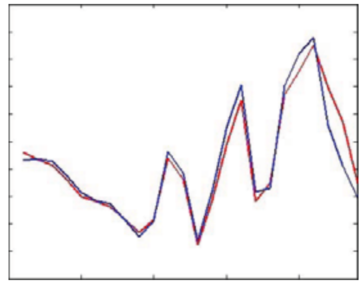

color constant & blur robust
edge description

Retrieval Experiment I

- Twenty different objects where captured under 11 different object orientations and 11 different light sources (Simon Fraser).
- We compare the retrieval results of the color constant description with the color constant and blur robust description.
- Error given in Normalized Average Rank (NAR).

rank	1	2	>2	ANAR
p	180	5	15	0.010
φ_p	169	17	14	0.012
m	155	22	23	0.024
φ_m	115	23	65	0.049

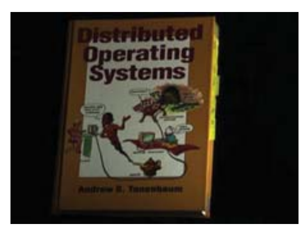

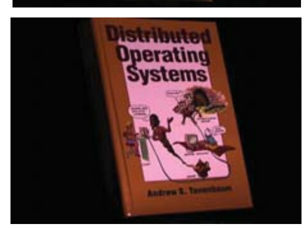

Retrieval Experiment I

• Twenty pairs of images with varying image blur.
• We compare the retrieval results of the color constant description with the color constant and blur robust description.

rank	1	2	>2	ANAR
p	7	2	11	0.365
φ_p	16	3	1	0.018
m	6	2	12	0.303
φ_m	13	1	6	0.053

Extra slides

Dichromatic Reflection Model in Chromaticity Representation

Chromaticities:

$$\{r,g,b\} = \{\frac{R}{R+G+B}, \frac{G}{R+G+B}, \frac{B}{R+G+B}\}$$

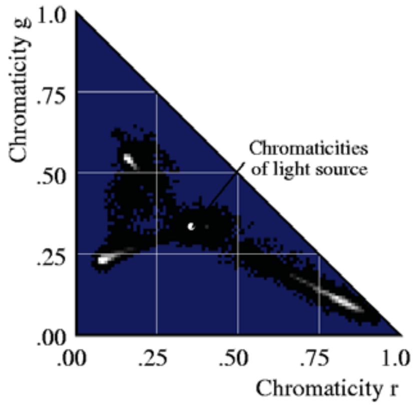

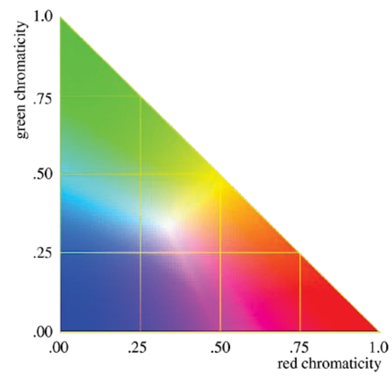

experiments

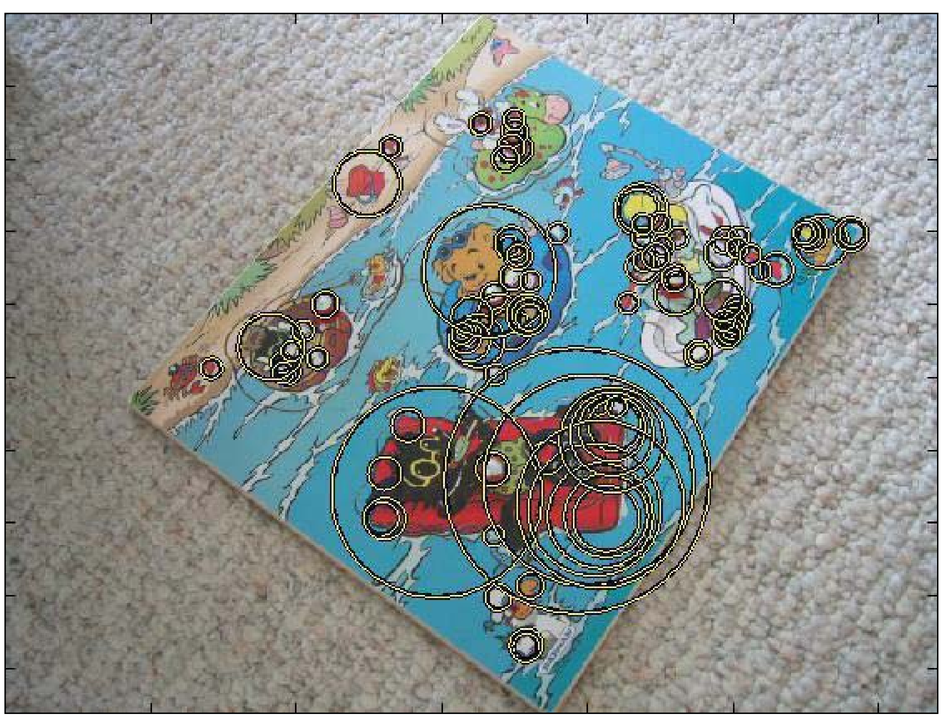

no boosting

experiments

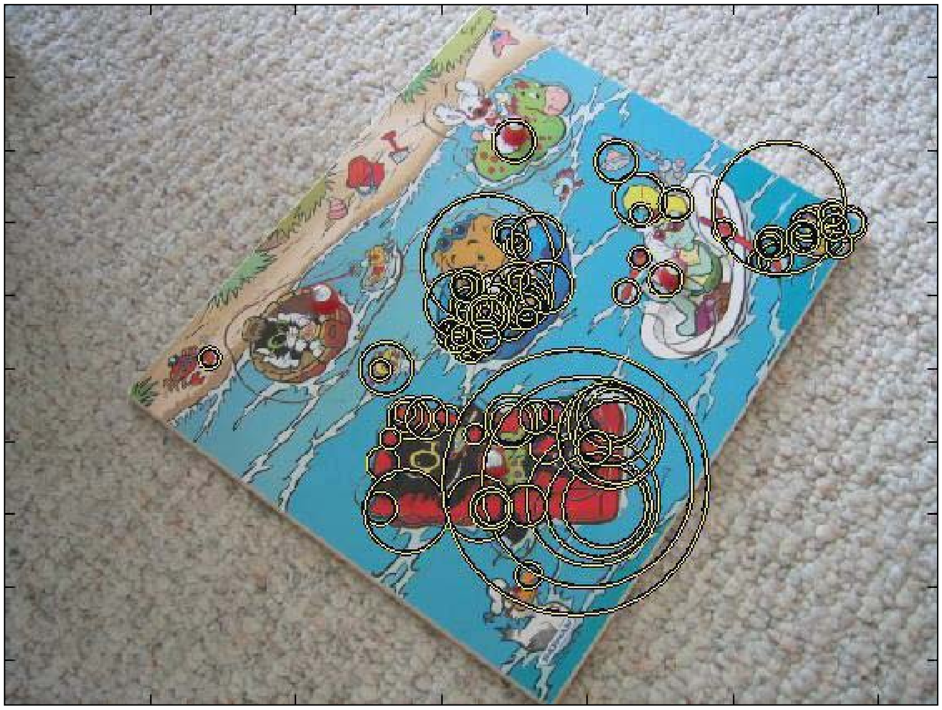

boosting

repeatability: photometric robustness

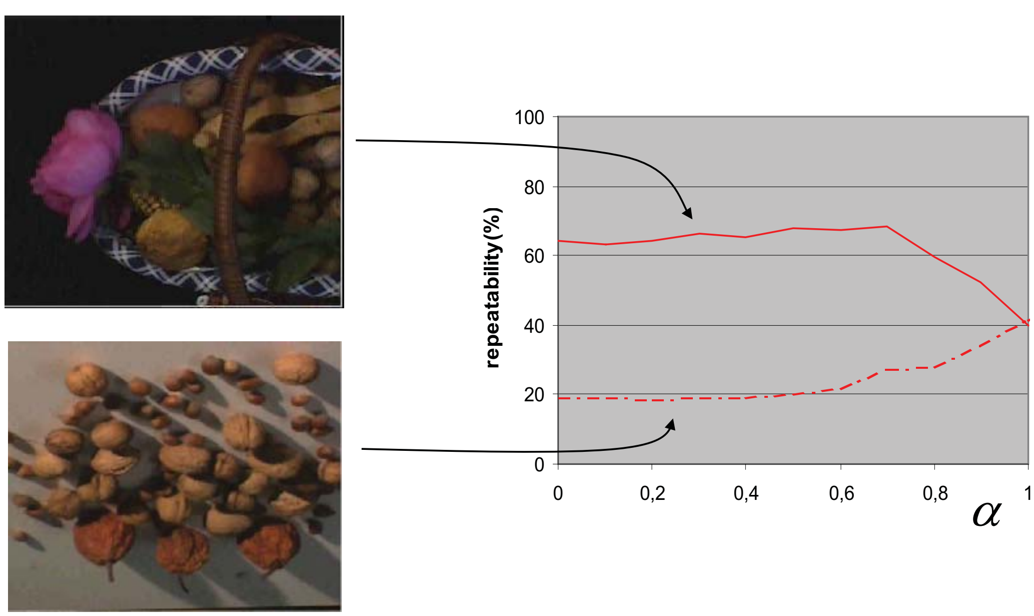

The salient space is highly photometrically invariant.

color distinctiveness

- color distinctiveness is measured by the information content:

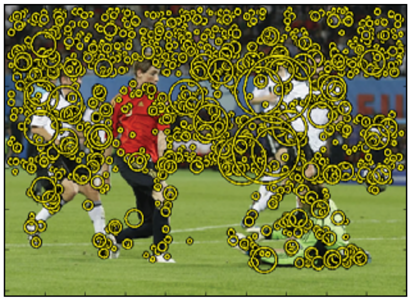

$$I(v) = -\sum_{\text{all salient points}} \log(p(v))$$

	20 points		100 points	
	incr(%)	decr(%)	incr	decr
opponent	63	0.0	22	1.0
spherical	49	1.0	13	1.3
HSI	57	0.3	17	1.1

Dichromatic Model

- dichromatic model:

$$\mathbf{F} = e(m^b\mathbf{C}^b + m^s\mathbf{C}^s)$$

body + specular

intensity illuminant

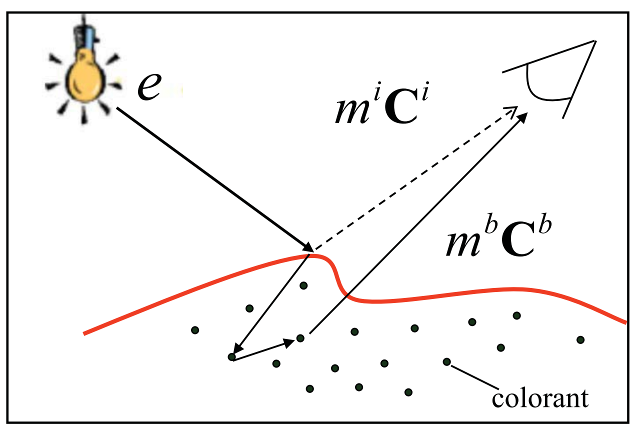

- first order photometric structure:

$$\mathbf{F}_x = \{R_x, G_x, B_x\} = m^b\mathbf{C}^b_x + \left(e_x m^b + e m^b_x\right)\mathbf{C}^b + e m^i_x \mathbf{C}^i$$

material + (shadow + shading) + specular

experiments

- visual inspection

- evaluation criteria:
  1. repeatability : salient point detection should be stable under varying viewing conditions.
  2. distinctiveness : salient point should focus on events with a low probability of occurrence.
  3. optimal : is the transformation optimal ?

The do’s and dont’s of Color Features

1. Take care in combining different channels:
   Tensor-based features solve the opposing vector problem.
2. Look at what kind of photometric invariance your problem needs:
   Quasi-invariants are more stable for feature detection
   Do not take derivatives of circular color spaces.
   Compute first derivatives, then color space transform.
3. When working with invariance take instabilities into account:
   Use error analysis to find certainty measures for your invariants.
4. When considering photometric invariance always also take discriminative power into account.
5. In case of image derivative-based descriptors take be aware of blur.
   Divisions of derivative based descriptors are often robust to blur.
6. From information theory an optimal color space for salient feature detection can be derived.
7. Color information is highly corrupted in compressed data. In compression (jpeg, mpeg) chrominance is subsampled.

Colour constancy algorithms

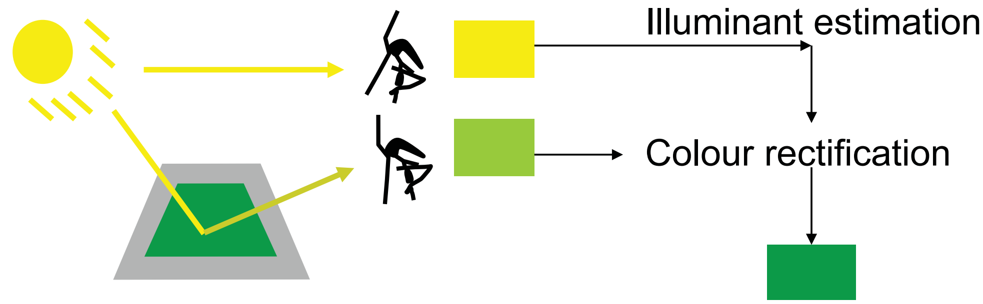

Invariant Normalizations

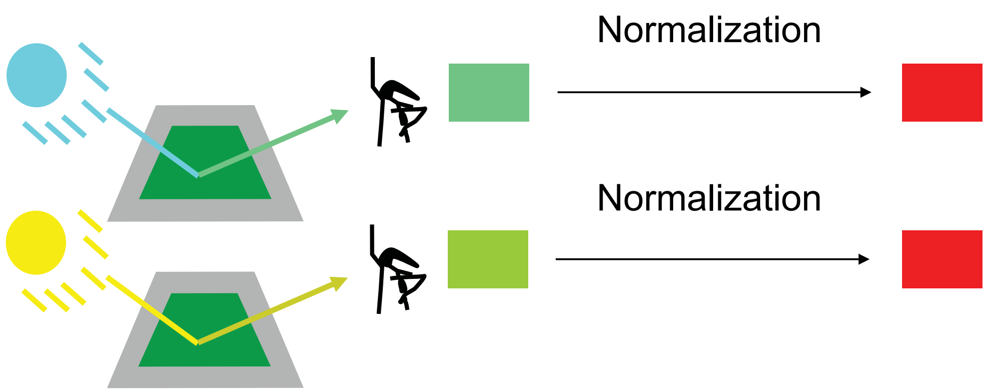

experiments : robust optical flow estimation

input sequence

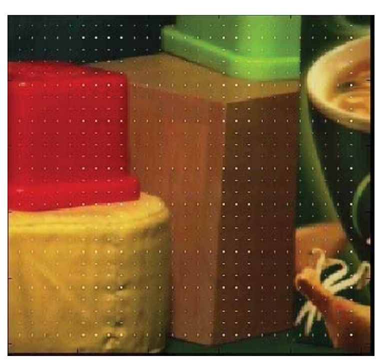

photometric optical flow

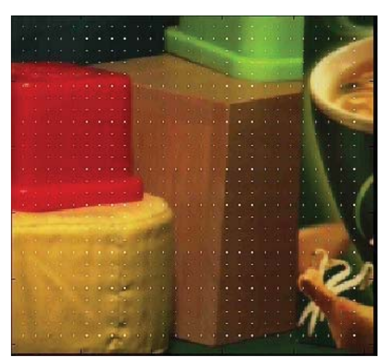

robust photometric optical flow