

Computational Vision

U. Minn. Psy 5036

Daniel Kersten

Lecture 20: Surface Material

Initialize

Spell check off

```
In[314]:= Off[General::spell1];
```

```
In[315]:= SetOptions[ArrayPlot, ColorFunction -> "GrayTones",  
  DataReversed -> True, Frame -> False, AspectRatio -> Automatic,  
  Mesh -> False, PixelConstrained -> True, ImageSize -> Small];  
SetOptions[ListPlot, ImageSize -> Small];  
SetOptions[Plot, ImageSize -> Small];  
SetOptions[DensityPlot, ImageSize -> Small, ColorFunction -> GrayLevel];  
nbinfo = NotebookInformation[EvaluationNotebook[]];  
dir = ("FileName" /. nbinfo /. FrontEnd`FileName[d_List, nam_, ___] -> ToFileName[d]);
```

Outline

Last time

Local measurements

Representing motion, Orientation in space-time

Optic flow, the gradient constraint, aperture problem

Neural systems solutions to the problem of motion measurement.

Space-time oriented receptive fields

Global integration

Sketched a Bayesian formulation--the integrating uncertain local measurements with the right priors can be used to model the perception of direction of plaid and planar motions.

Today

Later, we'll pick up on motion again--namely structure from motion in the context of determining layout and computing heading

Today, surface material:

surface properties, color, transparency, etc..

reflectance & lightness constancy

transparency, shininess,...

“Cooperative” computation or “strong fusion”.

Introduction to material perception

Material & Texture modeling

General categories of the "stuff" we see: surfaces (opaque and transparent), particle clouds (e.g. smoke, mist,..), liquids, hair, fur,...

Connection with count vs. mass nouns.

Research in computer graphics has provided major progress in the characterization of real surfaces, but realism is still a challenge.









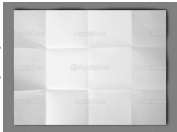
Uses of material perception

We can visually sense:

Dynamic physical properties such as the viscosity of honey, the elasticity of a rubber ball, the stiffness of cloth. While the perception of these properties is most compelling with motion, the inferences are possible even in static views given sufficient context.



Optical properties such as transparency of glass, the glossiness of granite, the wetness of pavement, the shininess of a spoon, the refractivity and reflectivity of glass, and degree of pigmentation of a piece of paper

```
In[330]:= GraphicsRow[ { Image[  ], Image[  ], Image[  ],  
  
Image[  ], Image[  ], Image[  ],  
  
Image[  ], Image[  ], Image[  ] }
```











The last image above illustrates the ambiguity associated with determining degree of pigmentation. The paper is uniform white, but it isn't obvious. Further, humans can make a judgment of brightness, while at the same time inferring the degree of pigmentation (i.e. reflectance or albedo). Sometimes the word "lightness" is used when talking about correlates with reflectance. The ability to perceive or accurately estimate surface reflectance depends on sufficient scene context. The colors of surfaces in the world look different when viewed through a small aperture.

Knowledge of material attributes are useful for a broad range of functions:

Material perception is important for action consequences (or affordances). For example:

How should I grasp it? Can I eat it? Can I shape it? How should I walk or drive on it? Is she sick? Is it slippery?

```
GraphicsRow[  
  
{ Image[  ], Image[  ], Image[  ], Image[  ],  
  
Image[  ], Image[  ], Image[  ], Image[  ] }
```



To use Jan Koenderink's expression, surface appearance is just the "crust" of underlying properties that an organism finds useful.

Learned associations are likely important, but we know relatively little about learning associations between material appearance and context such as shape. Or the role of reward. See context demos at end of lecture.

Uniform smooth materials

We'll now focus on solid uniform surface materials (rather than particle materials, like smoke or fire), and their optical properties.

Surfaces with material properties or attributes:

reflectance ("paint" or pigment or albedo)
matte and shiny
mirrors

transparency
multiplicative, additive

Physics-based generative modeling: Bidirectional reflectance distribution functions

Earlier we learned about the Lambertian shading equation, and early related models such as the "Phong model". But in general surface reflectances are more complicated.




Figure from: Image-Based BRDF Measurement Including Human Skin Stephen R. Marschner* Stephen H. Westin Eric P. F. Lafortune, Kenneth E. Torrance Donald P. Greenberg

The Bidirectional Reflectance Distribution Function (BRDF) describes directional dependence of the reflected light energy. The BRDF represents, for each incoming angle, the amount of light that is scattered in each outgoing angle.

For a given wavelength, it is the ratio of the reflected radiance in a particular direction to the incident irradiance:

$$\rho(\theta_i, \phi_i, \theta_e, \phi_e) = \frac{dL_e(\theta_e, \phi_e)}{dE_i(\theta_i, \phi_i)}$$

where E is the *irradiance*, that is the incident flux per unit area ($w\text{-}m\text{-}2$), and L is the reflected *radiance*, the reflected flux per unit area per unit solid angle ($w\text{-}m\text{-}2\text{-}sr\text{-}1$). The units of BRDF are inverse steradi-

ans. Respects the physics: Reciprocity, energy conservation. See: https://en.wikipedia.org/wiki/Bidirectional_reflectance_distribution_function.

Luminance (candelas per meter squared) is the radiance weighted by the human photopic sensitivity function (sensitivity as a function of wavelength). Illuminance (in lux) is irradiance also weighted by the photopic sensitivity function.

For isotropic materials, the reflections don't change with rotations about the surface normal.

This isn't true for anisotropic materials, such as "brushed metal".

For a Lambertian (perfectly diffuse) surface, for example, the BRDF is constant. The surface brightness is the same from all directions.

The Phong model described earlier in the context of shape-from-shading can approximate only a subset of surfaces characterized by BRDFs.





Figure from: <http://graphics.stanford.EDU/~smr/brdf/bv/>

- What is the BRDF for a perfectly mirror surface? Hint: it involves delta functions.

Ward reflection model: For calculating an image from a description of the shape, the illumination, and the BRDF

The Ward model is a physically realizable cousin of the Phong model. It respects energy conservation.




Subscripts i and e below indicate incoming and outgoing rays, respectively.

$$L_e(\theta_e, \phi_e) = \iint L_i(\theta_i, \phi_i) \rho(\theta_i, \phi_i, \theta_e, \phi_e) \cos \theta_i \sin \phi_i d\theta_i d\phi_i$$

Here is an example with diffuse (Lambertian) and specular components:

$$\rho(\theta_i, \phi_i, \theta_e, \phi_e) = \frac{\rho_d}{\pi} + \rho_s \frac{e^{-\tan^2(\delta)/\alpha^2}}{4\pi\alpha^2 \sqrt{\cos \theta_i \cos \theta_e}}$$

δ is the angle between the viewer and the vector defining the mirror reflection of the incident ray (i.e. where the angle of reflection equals the angle of incidence). α can be thought of as a measure of "roughness", and ρ_d and ρ_s give the amounts of diffuse and reflected contributions. Here's a figure from Fleming et al."



Other links

<http://www.cs.princeton.edu/~smr/cs348c-97/surveyspaper.html>

<http://www.graphics.cornell.edu/research/measure/>

For examples using BRDF measurements of human skin see:

<http://www.graphics.cornell.edu/online/measurements/>

Textured materials

Note that "texture" sometimes refers to low-level cues or image statistics useful for inferring properties like slant and shape, but it is also used to refer to surface material properties that are useful to estimate or label. Estimating physical texture properties is useful as a view-invariant object property, as a correlate of surface roughness related to friction. So keep in mind that sometimes "texture" refers to an image features, and other times to 3D surface or pattern properties. One can have a uniform surface texture, with very non-uniform image appearance and statistics (e.g. folded cloth, crumpled paper).

In this lecture, we focus on texture as a material property.

Textures can be:

regular ("herringbone pattern") or stochastic ("fur")

This applies to both image and surface properties. See: Liu, Y. (2010). Computational Symmetry in Computer Vision and Computer Graphics. Foundations and Trends® in Computer Graphics and Vision, 5(1-2), 1–195. <http://doi.org/10.1561/06000000008>

Textures can result from particles that:

"cohere" or not e.g. asphalt vs. sand & gravel

Textures can be dynamic, i.e. "flows" of liquid or particles. Optic flows can depend on a wide range of underlying physical causes:

liquid, particles, rigid shiny surface, non-rigid surface, eye or camera motion.
flows can be dominated by interactions between like particles, or be “buffeted” by invisible particles.

There can be similar patterns despite very different causes, e.g. “social forces” in crowd movement.

Textures can be caused by:

reflectance/pigment variations or bump (small geometric) variations

perceptually it isn't always easy to tell the difference, and may not matter depending on visual function.

For example, consider the visual and tactile differences between real wood and synthetic wood finishes.

(Note that image texture can also result from a completely uniform (shiny) material reflecting a textured environment)

In summary, the key characteristics of texture:

spatial variations are small with respect to the global scale of the surface structure
spatial homogeneity of repeating elements or tokens, or repeating statistics.

The notion of repeating tokens is easy to appreciate. And there is a history in the mathematics of symmetry which has formalized such textures. (E.g. group theory has shown that there are 17 basic symmetry groups in 2D).

In contrast, repeating statistics raises the key question of what statistics are being computed and over what spatial size. This is fundamentally a question of perceptual processes.

Appearance-based measurements

How can one characterize the generative model for textures? Much more complicated because of small-scale, but not micro-scale surface non-uniformity.

One approach is to build a database from texture measurements analogous to the BRDF, called a BTF or “Bidirectional Texture Function”. Rather than a physical model, the BTF is a look-up table based on a large set of measurements. See:

<http://www1.cs.columbia.edu/CAVE/projects/btf/>

Random synthesis and learning of stochastic image textures...

Next lecture. See Heeger and Bergen in Supplementary material, Zhu et al.,

Reflectance estimation & lightness

Introduction

The trichromacy theory of human color vision (Young, Maxwell, Helmholtz in the 19th century) states that three sensor types could explain the data on color matching has been highly successful. However,

it can not explain some phenomena. Simultaneous and successive color contrast effects required an elaboration of trichromacy to opponent-color coding.

What about color constancy? Color constancy refers to the observation that the color (of an object) can remain relatively unchanged with both spatial and chromatic changes in illumination. Historically, much experiment and thought went into accounting for the phenomena of color appearance over the last century and a half. However, the problem of what color vision is for, received less attention.

Although people have made various conjectures about the function of color vision, the advent of computational vision helped to clarify and motivate research into color's function. One idea is that it could aid in segmentation--the localization of boundaries in the absence of luminance contours. Another idea is that color could provide a surface attribute, relatively invariant over illumination variations useful for recognition. So just like shape is informative for object recognition, so might surface color. This second explanation gives us a different slant on color constancy--because in order to have a reliable surface attribute, we must be able to compute it given variations in secondary variables. Surface color is not given to the eye directly because the color and spatial distribution of the illumination typically varies. If we computed the color of an object by simply registering its wavelength composition, the object would rarely appear the same as it was moved about the room, or from indoors to outdoors. To obtain some measure of color constancy, vision discounts illumination as a secondary variable. The neural mechanisms of color vision support the estimation of invariant intrinsic surface properties that are in some sense closer to the parameters of the BRDF than to the physical light input to the eye.

If color constancy is a result of the neural system's attempt to estimate a surface attribute, then what is that attribute, and how do we estimate it? Let's first look at simplified versions of the problem--in particular, lightness constancy.


Functional vs. mechanistic explanations

Overview of lightness effects

Recall the Land & McCann "Two squares and a happening", and the two-cylinders version of it.

<http://web.mit.edu/persci/demos/Lightness/gaz-teaching/index.html>

Simultaneous contrast




<http://web.mit.edu/persci/demos/Lightness/gaz-teaching/flash/contrast-movie.swf>

Lightness & spatially uniform illumination

Simple constancy: simultaneous contrast mechanisms vs. functional algorithms

Local contrast



Spatially uniform illumination. $L = RE$.

The simple generative model, together with a task assumption that assumes vision values an invariant object property,

suggests that R should be a better predictor of lightness than L .

Lightness $\rightarrow R$, where R is between 0 and 1.

Let L_1 and L_2 be the luminance of the small center disks, and R_1 and R_2 be their reflectances.

Relative reflectance and local contrast: $L_1/L_2 = (R_1E)/(R_2E) = R_1/R_2$

So perceived ratios of lightness should match luminance ratios. What about the sense that

lightness as a particular absolute value, e.g. "white", "black", "medium gray"?

Normalization or anchoring problem (see Gilchrist).

Estimate of $R_1 \sim L_1/L_{avg}$

or $R_1 \sim L_1/L_{max}$?

There is a physical constraint on reflectance. In the natural world a typical range is about 10 to 1, or with really white and really black surfaces up to 30 to 1.


Role of contrast and adaptation mechanisms--See Kraft and Brainard (1999).

See Appendix for a simple Bayesian model of reflectance estimation.


Spatially varying illumination

Recall Land & McCann's "Two squares and a happening" in the discussion of edge detection (Lecture 13)

The left half looks lighter than the right half.




But the intensity across a horizontal line tells a different story:



The two ramps are identical. One can apply the above rule that the luminance ratio at edges provides a good estimate of the reflectance ratio at edges to build a lightness algorithm (See Appendix). Algorithms can be applied to other related stimuli such as the O'Brien Cornsweet effect.

Craik-O'Brien-Cornsweet effects

```
size = 256; y[x_] :=  $\frac{1.}{\text{Abs}[\frac{x}{2}] + 6}$ ; y1 = Table[y[x], {x, 0,  $\frac{\text{size}}{2} - 1$ }]
y12 = Join[-Reverse[y1], y1 - 0.03]
ListPlot[y12 + 0.5, PlotRange → {-1, 1}, Joined → True]
```



```
picture = Table[Table[y12[[i]], {i, 1, size}], {i, 1, size}];
picture2 = Table[Table[-(y12[[i]]), {i, 1, size}], {i, 1, size}];
gp1 = ArrayPlot[picture, Frame → False,
  Mesh → False, PlotRange → {-1, 1}, AspectRatio → Automatic]
gp2 = ArrayPlot[picture2, Frame → False, Mesh → False,
  PlotRange → {-1, 1}, AspectRatio → Automatic];
```



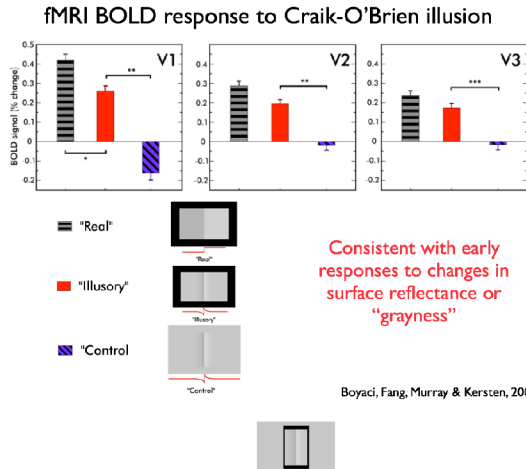
Lightness algorithms in "flat land"

The Appendix provides some examples of historical approaches to computing lightness.

Visual cortex: Respond to luminance or lightness change?

If we could measure the activity in the first cortical area, V1, would responses there correspond to intensity change or lightness change?

Using functional magnetic resonance imaging, we can image V1 activity in subjects viewing a contrast-reversing Craik-O'Brien illusion. The result is that V1 responds to lightness contrast with almost the same strength as to a perceptually equivalent luminance change.



See supplementary pdf notes on course web page. Boyaci, H., Fang, F., Murray, S. O., & Kersten, D. (2007). Responses to lightness variations in early human visual cortex. *Curr Biol*, 17(11), 989-993. (pdf) and

Boyaci, H., Fang, F., Murray, S. O., & Kersten, D. (2010). Perceptual grouping-dependent lightness processing in human early visual cortex. *Journal of Vision*, 10(9). <http://jwww.journalofvision.org/content/10/9/4.full>

The animation below shows what contrast-reversing lightness illusion looks like:


```
ListAnimate[{gp1, gp2, gp1, gp2}, 1, AnimationRunning -> False]
```



Reflectance estimation & indirect lighting


A key missing component of our generative models so far is failing to take into account how light bounces around between surfaces.

A good first question to ask is: Does the visual system take into account the geometry of reflected light? Consider indirect lighting such as would arise in a corner. Below is a photograph of a corner consisting of a white paper on the left and a red paper on the right. The white paper looks pinkish (consistent with the physical spectrum induced by the light landing on it from the red paper). Normally we don't notice the color of reflected light, perhaps because given sufficient cues to shape, it can discount the color of indirect lighting. How can we test this?




http://gandalf.psych.umn.edu/~kersten/kersten-lab/Mutual_illumination/BlojKerstenHurlbertDemo99.pdf
(See Bloj MG, Kersten D, Hurlbert AC (1999))


Imagine making the following card using a color printer:



Fold it, arrange the lighting as show below, and then look at it from above. The physics of the situation is illustrated below.



Now if you look at the paper steadily, you would experience a spontaneous reversal in the shape going from a concave corner to a convex "roof". When the paper looks like a roof, the white side appears more pinkish than when the paper appears to be a concave. Why is this so?



Two generative models

Blouin et al. measured the perceptual effect and produced a quantitative explanation showing that the visual system seemed to have

"built-in" knowledge about the effects of indirect vs. direct lighting. They used the following generative models to produce a Bayesian estimate of surface color matches:

Indirect plus direct lighting (1 bounce, Funt & Drew, 1991)

$$I(\lambda) = E(\lambda)\rho_1(\lambda)\left[\cos(\alpha_1) + f_{21}(x)\rho_2(\lambda)\cos(\alpha_2)\right] + n$$

$$f_{21}(x) = \frac{1}{2}\left[1 + \frac{\cos(\beta) - (x/w)}{\sqrt{(x/w)^2 - 2(x/w)\cos(\beta)}}\right]$$

Direct lighting only (0 bounce)

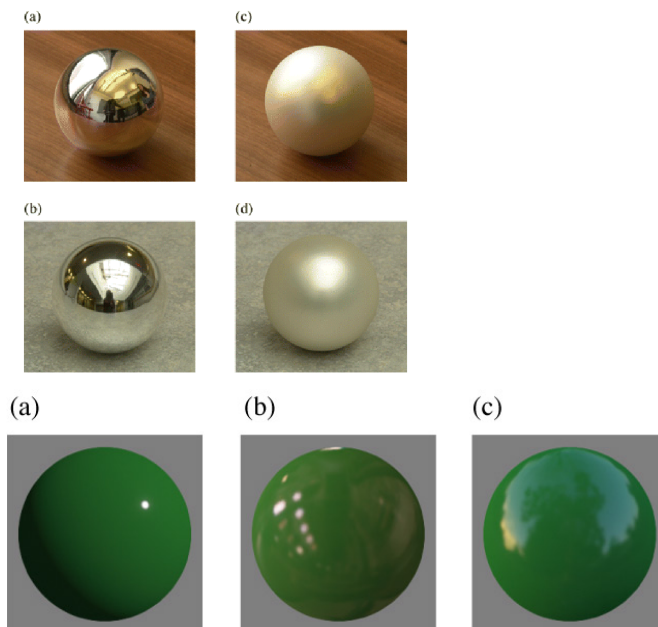
$$I(\lambda) = E(\lambda)\rho_1(\lambda)\left[\cos(\alpha)\right] + n$$


Bottom line: Human color matches consistent with built-in knowledge of the generative laws of reflection.

Perception of shiny materials

Static patterns

The perception of shiny materials is a striking illustration of invariant material perception. The top left and bottom left panels of the figure below show a ball made of the same material--in effect the same ball in every respect. Yet, because the material is nearly perfectly reflective, the appearances depend entirely on the illumination of the surrounding environment.





From: Fleming RW, Dror RO, Adelson EH (2003) Real-world illumination and the perception of surface reflectance properties. *J Vis* 3:347-368.

One of the main results of this study was to show that human judgments of specular attributes increases in accuracy when a shiny object is placed in a complex but still naturalistic environment. What is "naturalistic"? One ingredient is that the images have the kurtotic histogram properties that we studied in an earlier lecture. This is consistent with the presence of edges being important. Bright points are important too. Recognizable reflected objects are not necessary. Fleming et al used the above Ward model and had human subjects estimate qualities related to the degree of specularity (the image of a point light source on a really shiny object is quite sharp, i.e. a point), and the relative amounts of specular to diffuse components. See α and ρ_s in the Ward model above.

<http://journalofvision.org/3/5/3/article.aspx>

Moving patterns: Shiny or matte?

Optic flow can disambiguate the perception of whether a surface is shiny or matte.



<http://gandalf.psych.umn.edu/~kersten/kersten-lab/demos/MatteOrShiny.html>

See:

Doerschner, K., Fleming, R. W., Yilmaz, O., Schrater, P. R., Hartung, B., & Kersten, D. (2011). Visual motion and the perception of surface material. *Current Biology*, 21(23), 2010–2016.

doi:10.1016/j.cub.2011.10.036

and

Doerschner K, Kersten D, & Schrater P. (2009) Rapid Classification of Surface Reflectance from Image Velocities. *Computer Analysis of Images and Patterns, Lecture Notes in Computer Science*, Springer Berlin / Heidelberg, volume 5702, pp. 856-864.

They analyzed the optic flow patterns produced by shiny vs. matte objects. Specularities tend to "stick"

to points of high curvature. Objects that have a range of curvatures might be expected to show peculiarities in the histograms of speed components extracted from the optic flow. Doerschner et al. (2009) showed that the histograms of curvy shiny objects tend to be more bimodal than for matte objects.

Doerschner et al. (2011) showed how a number of optic-flow based cues play a role in distinguishing shiny from matte materials.

Cooperative computation, context, and more demos

Interaction between material perception and structure-from-motion


<http://gandalf.psych.umn.edu/users/kersten/kersten-lab/demos/transparency.html>

The role of context

Context is important :

```
In[328]:= ListAnimate[{upsideup, rightsideup},  
  AnimationRunning → False, AnimationRepetitions → 1]
```

Out[328]=



```
Image[];
```

Missing window pranks

<https://youtu.be/mHaNd1U1950>

<https://youtu.be/1DWpgMHprlU>

Dynamic properties

Some demos from Eitan Grinspun's lab. <http://www.cs.columbia.edu/~eitan/>

Appendices

Simple constancy: Bayesian reflectance estimation of one isolated patch in flatland

This simple example shows how marginalization or integrating out secondary variables can constrain an otherwise unconstrained problem...even without hypothesizing explicit priors on reflectance or illumination distributions. (Freeman, 1994)

$L = R \cdot EI + \text{noise}$. Given L , what is R ? EI is the secondary variable that we want to discount by marginalization.

Let the illumination range be $[0, 10]$, and the reflectance range $[0, 1]$. Then the luminance range is also $[0, 10]$.

Let the luminance noise be Gaussian with a standard deviation less than 10% of L , say 1 for simplicity. Then the probability of an observation L given R and EI is proportional to:


```
likeli[L_, R_, EI_] := Exp[-(.5)*(L-R*EI)^2]/Sqrt[2. Pi];
```

```
prior[R_] := PDF[UniformDistribution[{0, 1}], R]
```


where we assume that the noise has a Gaussian distribution.

Here is a plot of the likelihood for a luminance value of 1:

```
Plot3D[prior[R] * likeli[1, R, EI],
  {R, 0.1, 0.9}, {EI, 0.1, 10}, AxesLabel -> {"R", "E", "p"}]
```




```
DensityPlot[likeli[1, R, EI], {R, 0.1, 0.9}, {EI, 0.1, 10}]
```



Marginalize over illumination, EI to get the likelihood of R for a given value of L :


```
pr[L_, R_] := Evaluate[Integrate[prior[R] * Exp[-(L-R*EI)^2], {EI, 0, 10}]]
```

```
Plot3D[pr[L, R], {L, 0, 10}, {R, 0.0, 1},
  AxesLabel -> {"R", "L", "p"}, PlotRange -> {0, 10}]
```




Here is the relative probability of R given L = 0.75

```
pt = {{r0 = R /. NMinimize[-pr[.75, R], R][[2]], 0}, {r0, 10}};
Plot[pr[0.75, R], {R, 0, 1}, PlotRange -> {0, 10}, Epilog -> {Red, Thick, Line[pt]}]
```



```
LvsR = Table[{L, R /. NMinimize[-pr[L, R], R][[2]]}, {L, .01, 10, .1}];
ListPlot[LvsR]
```




Some notes on color illusions


Munker-White illusion

<http://web.mit.edu/persci/people/bart/DemoLinks.html>

White M (1981) The effect of the nature of the surround on the perceived lightness of graybars within square-wave test gratings. Perception 10:215-230.



Neon color spreading, etc.



From: <http://neuro.caltech.edu/~carol/VanTuijl.html>


<http://www.cogsci.uci.edu/%7Eeddhoff/illusions.html>

Hoffman DD (1998) Visual Intelligence. New York: W. W. Norton & Company.


http://www.michaelbach.de/ot/col_neon/

Land, Horn and others

Given $L(x,y)=S(x,y)E(x,y)$, where L is the intensity/luminance data of the image (using an achromatic world), we attempt to estimate $S(x,y)$.




Rather than seeking a spatial filter explanation (e.g. do edge detection, and then fill in the region up to the edges with the color of the edges), consider the following functional explanation of this illusion:



The lightness value we assign is correlated with S, not with L. So how can we estimate S?

The idea is to assume that the image intensity changes (or changes in r, b, or g) are due to slowly varying illumination together with piece-wise constant reflectances. The slowly varying illumination needs to be filtered out. Land's scheme was to use ratios:



We want to estimate

$$\frac{L_1}{L_5} = \frac{L_1}{L_2} \cdot \frac{L_2}{L_3} \cdot \frac{L_3}{L_4} \cdot \frac{L_4}{L_5}$$

$$\frac{L_1}{L_5} = \frac{S_1 E_1}{S_2 E_2} \cdot \frac{S_2 E_2}{S_3 E_3} \cdot \frac{S_3 E_3}{S_4 E_4} \cdot \frac{S_4 E_4}{S_5 E_5}$$

But we would like to discount small changes in L, so we can use the rule:

$$\text{if } \left| 1 - \frac{L_i}{L_{i+1}} \right| < t$$

$$\text{then set } \frac{L_i}{L_{i+1}} = \frac{S_i E_i}{S_{i+1} E_{i+1}} = 1$$

$$1 \cdot 1 \cdot \frac{L_3}{L_4} \cdot 1 = 1 \cdot 1 \cdot \frac{S_3 E_3}{S_4 E_4} \cdot 1 \cong \frac{S_3}{S_4} = \frac{S_1}{S_5}$$

and thus by using luminance ratios that are sufficiently large in the product, we obtain an estimate of the relative reflectance

$$S_5 \cong S_1 \left(\frac{L_4}{L_1} \right)$$

where the luminance ratio is measurable. We can obtain estimates for $i = 1, 2, 3, 4, 6$. (See Appendix to see how Land extended this lightness algorithm model to color).

Horn's algorithm

Luminance (L) = reflectance (S) x illumination (E)

1) Take logs to turn the multiplication into addition:

$$1) C(x, y) = \log L(x, y) = \log(SE) = \log S(x, y) + \log E(x, y) = S' + E'$$

2) High-pass filter to amplify the edges

$$\nabla^2 C(x, y) = \nabla^2 S' + \nabla^2 E' \quad (\text{or } \nabla^2 G * C)$$

3) Threshold all values below some finite threshold

$$t(x, y) = T[\nabla^2 C] = T[\nabla^2 S' + \nabla^2 E'] = \nabla^2 S'(x, y)$$

Using Poisson's equation, solve for $S'(x, y)$.

$$S' = t * g$$

$$F[S'] = F[t]F[g]$$

The mathematical complication is because the problem is two-dimensional. In one dimension, only beginning calculus is required to understand how to solve a simple differential equation--just integrate.

Both Horn's method and Land's have some problems:

- 1) Normalization (anchoring problem) is actually more complicated, because we have taken second derivatives, leaving an extra degree of freedom in the integration process.
- 2) Spatial scale and threshold
- 3) restricted to flatland.


In fact most of the alternatives face the same problems. One could imagine various ways of filtering out the illumination, for example, using spatial frequency representations of the image..but this does not help.

Mathematica demonstration of a 1D lightness calculation in flatland

One explanation is that the visual system takes a spatial derivative of the intensity profile. Recall from calculus that the second derivative of a linear function is zero. So a second derivative should filter out the slowly changing linear ramp in the illusory image. We approximate the second derivative with a discrete kernel (-1,2,-1). Let's apply this to a line across the Craik-O'Brien illusion above.

The steps are: 1) take the second derivative of the image;

```
filter = {1, -2, 1};
(*Take the second derivative at each location*)
fspicture = ListConvolve[filter, picture[[128]]];
ListPlot[fspicture, Axes -> False]
```




2) threshold. To handle gradients that aren't perfectly linear, we add a threshold function to set small values to zero before re-integrating:

```
threshold[x_, τ_] := If[Abs[x] > τ, x, 0];
SetAttributes[threshold, Listable];
fspicture = threshold[fspicture, 0.0025];
ListPlot[fspicture, Axes → False]
```

3) re-integrate


```
ListPlot[fspicture, Joined → True, PlotRange → {-0.1, 0.1}];
integratefspicture = FoldList[Plus, fspicture[[1]], fspicture];
integratefspicture2 = FoldList[Plus, integratefspicture[[1]], integratefspicture];
ListPlot[integratefspicture2, Joined → True, Axes → False]
```




Color constancy

Land's demonstrations

Beginning in the 1950's, Edwin Land has shown the sophistication of human color constancy in a number of striking demonstrations (Land, E.H., 1983). In one experiment, three lights (long, medium, and short wave lamps) illuminate a Mondrian consisting of a collection of patches of paper of various colors.





We consider two phases, each characterized by a different global illumination of the whole Mondrian. In the first phase, we pick out two patches, a white (W) and a yellow (Y) one, on which to focus our attention. A radiometer is used to measure the amount of each of the three components radiating off the yellow patch. Now, in the second phase, we adjust the irradiance of each of the colored lights so that we get the same readings for the white patch as we had for the yellow patch in the first phase. And as a consequence, the spectral composition of the yellow patch changes too, because it is now receiving the same changed illumination as the white patch. Based on spectral composition, we might predict that the white patch of the first phase would be made to appear yellow in the second phase. but it doesn't. color constancy is maintained, and the white patch appears white, and the yellow appears yellow. How can this be done?

Kraft and Brainard (1999) have measured color constancy under nearly natural viewing conditions. Their results rule out all three classic hypotheses: local adaptation, by adaptation to the spatial mean of the image, or by adaptation to the most intense image region. What more is needed to explain to constancy beyond these simple visual mechanisms?

References

- Adelson, E. H. (1999). Lightness Perception and Lightness Illusions. In M. S. M. Gazzaniga (Ed.), *The New Cognitive Neurosciences* (pp. 339-351). Cambridge, MA: MIT Press.
- Adelson, E. H. (1993). Perceptual organization and the judgment of brightness. *Science*, 262, 2042-2044.
- Adelson, E. H., & Pentland, A. P. (1996). The Perception of Shading and Reflectance. In D. C. Knill & W. Richards (Eds.), *Perception as Bayesian Inference* (pp. 409-423). Cambridge: Cambridge University Press.
- Anderson BL (2003) The role of occlusion in the perception of depth, lightness, and opacity. *Psychol Rev* 110:785-801.
- Anderson BL (2003) Perceptual organization and White's illusion. *Perception* 32:269-284.
- Beck, J. (1972). *Surface Color Perception*. Ithaca, NY: Cornell University Press.
- Bloj, M. G., Kersten, D., & Hurlbert, A. C. (1999). Perception of three-dimensional shape influences colour perception through mutual illumination. *Nature*, 402(6764), 877-879.
- Brainard, D. H., Brunt, W. A., & Speigle, J. M. (1997). Color constancy in the nearly natural image. I. Asymmetric matches. *J Opt Soc Am A*, 14(9), 2091-110.
- Brainard, D. H. (1998). Color constancy in the nearly natural image. 2. Achromatic loci. *J Opt Soc Am A Opt Image Sci Vis*, 15(2), 307-25.
- Brainard, D. H., & Freeman, W. T. (1997). Bayesian color constancy. *J Opt Soc Am A*, 14(7), 1393-411.
- Brainard, D. H., & Wandell, B. A. (1986). Analysis of the retinex theory of color vision. *J Opt Soc Am [A]*, 3(10), 1651-61.
- Brainard, D. H., & Freeman, W. T. (1994). *Bayesian Method for Recovering Surface and Illuminant*

Properties from Photosensor Responses. Paper presented at the Human Vision, Visual Processing and Digital Display V, San Jose, California.

Boyaci H, Maloney LT, Hersh S (2003) The effect of perceived surface orientation on perceived surface albedo in binocularly viewed scenes. *J Vis* 3:541-553.

Buckley, D., Frisby, J. P., & Freeman, J. (1994). Lightness Perception can be Affected by Surface Curvature from Stereopsis. *Perception*, 23, 869-881.

Doerschner K, Boyaci H, Maloney LT (2004) Human observers compensate for secondary illumination originating in nearby chromatic surfaces. *J Vis* 4:92-105.

Doerschner, K., Fleming, R. W., Yilmaz, O., Schrater, P. R., Hartung, B., & Kersten, D. (2011). Visual motion and the perception of surface material. *Current Biology*, 21(23), 2010–2016.

doi:10.1016/j.cub.2011.10.036

Drew M, Funt B (1990) Calculating surface reflectance using a single-bounce model of mutual reflection. In: *Proceedings of the 3rd International Conference on Computer Vision*, IEEE,, pp 393-399. Osaka.

D'Zmura, M. (1992). Color constancy: surface color from changing illumination. *Journal of the Optical Society of America, A*, 9(3), 490-493.

Fleming RW, Dror RO, Adelson EH (2003) Real-world illumination and the perception of surface reflectance properties. *J Vis* 3:347-368.

<http://journalofvision.org/3/5/3/Fleming-2003-jov-3-5-3.pdf>

Fleming, R. W., Jäkel, F., & Maloney, L. T. (2011). Visual Perception of Thick Transparent Materials. *Psychological Science*, 22(6), 812–820. doi:10.1177/0956797611408734

Freeman, W. T. (1992). The generic viewpoint assumption in a Bayesian framework. In D. C. Knill & W. Richards (Eds.), *Perception as Bayesian Inference*.: Cambridge University Press.

Hoffman DD (1998) Visual Intelligence. New York: W. W. Norton & Company.

Horn, B. K. P. (1974). Determining Lightness from an Image. *Computer Graphics and Image Processing*, 3((1)), 277-299.

Hurlbert, A. (1986). Formal connections between lightness algorithms. *Journal of the Optical Society of America, A*, 3, 1684-1693

Knill, D. C., & Kersten, D. (1991). Apparent surface curvature affects lightness perception. *Nature*, 351, 228-230.

Kraft, J. M., & Brainard, D. H. (1999). Mechanisms of color constancy under nearly natural viewing [In Process Citation]. *Proc Natl Acad Sci U S A*, 96(1), 307-12.

Land, E.H., 1983, Recent advances in retinex theory and some implications for cortical computations: Color vision and the natural image. *Proceedings of the National Academy of Science*, 80, 5163-5169.

Larson GW, Shakespeare R (1998) Rendering With Radiance: The Art and Science of Lighting Visualization.

Liu, Y. (2010). Computational Symmetry in Computer Vision and Computer Graphics. *Foundations and Trends® in Computer Graphics and Vision*, 5(1-2), 1–195. doi:10.1561/06000000008

Maloney, L.T., & Wandell, B.A. (1986). Color Constancy: A Method for Recovering Surface Spectral Reflectance..3, 29-33.

O'Brien, V. (1958). Contour Perception, Illusion and Reality. *Journal of the Optical Society of America*, 48((2)), 112-119.

Radonjić, A., Cottaris, N. P., & Brainard, D. H. (2015). Color constancy in a naturalistic, goal-directed task. *Journal of Vision*, 15(13), 3–21. <http://doi.org/10.1167/15.13.3>

Singh M, Anderson BL (2002) Toward a perceptual theory of transparency. *Psychol Rev* 109:492-519.

Van Tuijl, H. F. J. M. (1975). A new visual illusion: neonlike color spreading and complementary color induction between subjective contours. *Acta Psychologica*, 39, 441-445.