Computational Vision
U. Minn. Psy 5036
Daniel Kersten
Lecture 25: Perceptual integration, Cooperative Computation

Initialize

- **Spell check off**

```
Off[General::spell1];
```

Outline

Last time

Object recognition

Today

- **Use Bayesian framework to better understand how perceptual information gets put together. Look at examples of actualy calcuations.**
- **Modular vs. cooperative computation**

For the most part, we've treated visual estimation as if it is done in distinct "modules", such as, surface-color-from-radiance (Land, 1959), shape-from-shading (Horn, 1975), optic flow (Hildreth, 1983) or structure-from-motion (Ullman, 1979).

In contrast to the modularity theories of vision, it is phenomenally apparent that visual information is integrated to provide a strikingly singular description of the visual environment. By looking at how human perception puts integrates scene attributes, we may get some idea of how vision modules in the brain interact, and what they represent.

Review of basic graph types in vision (from Lecture 6)

See: Kersten, D., & Yuille, A. (2003) and Kersten, Mamassian & Yuille (2004)

Basic Bayes & vision modules

$$p[S \mid I] = \frac{p[I \mid S]\, p[S]}{p[I]}$$

Usually, we will be thinking of the **Y** term as a random variable over the hypothesis space, and **X** as data. So for visual inference, **Y = S** (the scene), and **X = I** (the image data), and **I = f(S)**.

We'd like to have:

p(S|I) is the **posterior** probability of the scene given the image

-- i.e. what you get when you condition the joint by the image data. The posterior is often what we'd like to base our decisions on, because as we discuss below, picking the hypothesis **S** which maximizes the posterior (i.e. maximum a posteriori or **MAP** estimation) minimizes the average probability of error.

p(S) is the **prior** probability of the scene.

p(I|S) is the **likelihood** of the scene. Note this is a probability of **I**, but not of **S**.

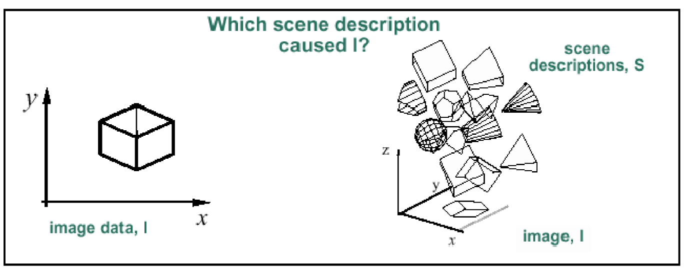

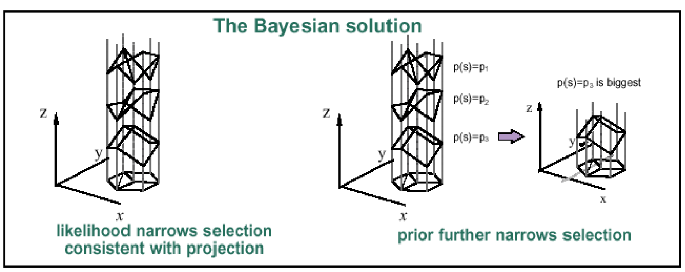

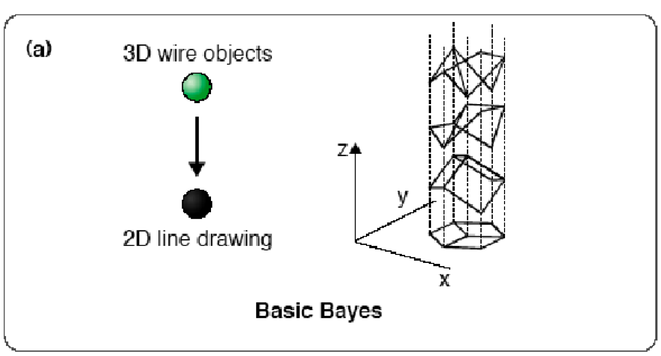

We've seen that the idea of prior assumptions that constrain otherwise underconstrained vision problems is a theme that pervades much of visual perception. Where do the priors come from? Some may be built in early on or hardwired from birth, and others learned in adulthood. See: Adams, W. J., Graf, E. W., & Ernst, M. O. (2004). Experience can change the 'light-from-above' prior. Nat Neurosci, 7(10), 1057-1058 for a recent example of learning the light from above prior for shape perception.

General Bayesian theory of low-level visual integration for separate "modules"

We've seen a number of applications of Basic Bayes, including the algorithms for shape from shading and optic flow.

In 1985, Poggio, Torre and Koch showed that solutions to many of computational problems of low vision could be formulated in terms of maximum a posteriori estimates of scene attributes if the generative model could be described as a matrix multiplication, where the image I is matrix mapping of a scene vector S:

$$I = \mathbf{A}S$$

$$E = (I - \mathbf{A}S)^T (I - \mathbf{A}S) + \lambda S^T \mathbf{B} S$$

Then a solution corresponded to minimizing a cost function E, that simultaneously tries to minimize the cost due to reconstructing the image from the current hypothesis S, and a prior "smoothness" constraint on S. λ is a (often free) parameter that determines the balance between the two terms. If there is reason to trust the data, then λ is small; but if the data is unreliable, then more emphasis should be placed on the prior, thus λ should be bigger.

For example, S could correspond to representations of shape, stereo, edges, or motion field, and smoothness be modeled in terms of nth order derivatives, approximated by finite differences in matrix B.

The Bayesian interpretation comes from multivariate gaussian assumptions on the generative model:

$$p(I \mid S) = k \times \exp\left[-\frac{1}{2\sigma_n^2}(I - \mathbf{A}S)^T (I - \mathbf{A}S)\right]$$

$$p(S) = k' \times \exp\left[-\frac{1}{2\sigma_s^2} S^T \mathbf{B} S\right]$$

Table 1 Regularization in early vision

Problem	Regularization principle
Edge detection	$\int [(Sf - i)^2 + \lambda (f_{xx})^2]\, dx$
Optical flow (area based)	$\int [i_x u + i_y v + i_t)^2 + \lambda (u_x^2 + u_y^2 + v_x^2 + v_y^2)]\, dx\, dy$
Optical flow (contour based)	$\int [(V \cdot N - V^N)^2 + \lambda ((\partial/\partial_s) V)^2]\, ds$
Surface reconstruction	$\int [S \cdot f - d)^2 + \lambda (f_{xx}^2 + 2f_{xy}^2 + f_{yy}^2)^2]\, dx\, dy$
Spatiotemporal approximation	$\int [(S \cdot f - i)^2 + \lambda (\nabla f \cdot V + ft)^2]\, dx\, dy\, dt$
Colour	$\lVert I^\nu - Az \rVert^2 + \lambda \lVert Pz \rVert^2$
Shape from shading	$\int [(E - R(f, g))^2 + \lambda (f_x^2 + f_y^2 + g_x^2 + g_y^2)]\, dx\, dy$
Stereo	$\int \{[\nabla^2 G * (L(x, y) - R(x + d(x, y), y))]^2 + \lambda (\nabla d)^2\}\, dx\, dy$

From Poggio, Torre & Koch, 1985

■ Discounting: Emphasing one model cause of image data over another

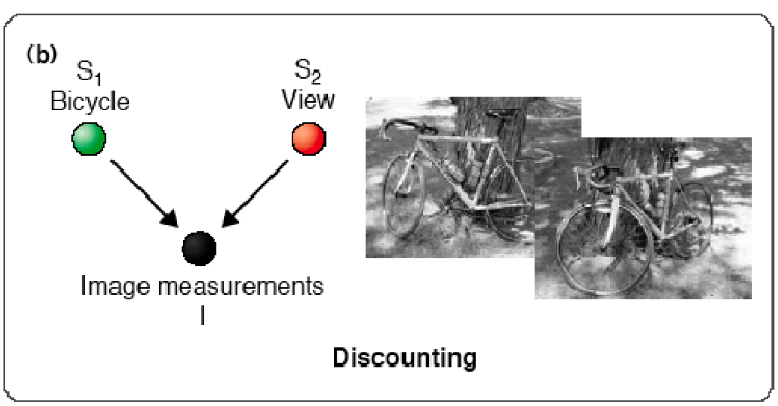

This Bayes net describes the case where the joint distribution can be factored as:

$p(s_1, s_2, I) = p(I|s_1, s_2) p(s_1) p(s_2)$

Optimal inference for this task requires that we calculate the marginal posterior:

$p(s_1|I) \propto \int_{S_2} p(s_1, s_2 \mid I)\, ds_2$

Liu, Knill & Kersten (1995) describe an example with:

I -> 2D x-y image measurements, s_1-> 3D object shape, and s_2-> view

Bloj et al. (1999) have an example estimating s_1-> surface chroma (saturation) with s_2-> illuminant direction.

■ Ideal observer for the "snap shot" model of visual recognition: Discounting views

Tjan et al. describe an application to object recognition in noise. Their idea was to measure how efficiently human observers recognize objects given view variation, and different types of image information. (See Tjan et al. (1995) Human efficiency for recognizing 3-D objects in luminance noise.Vision Research, 35, 3053-3069.)

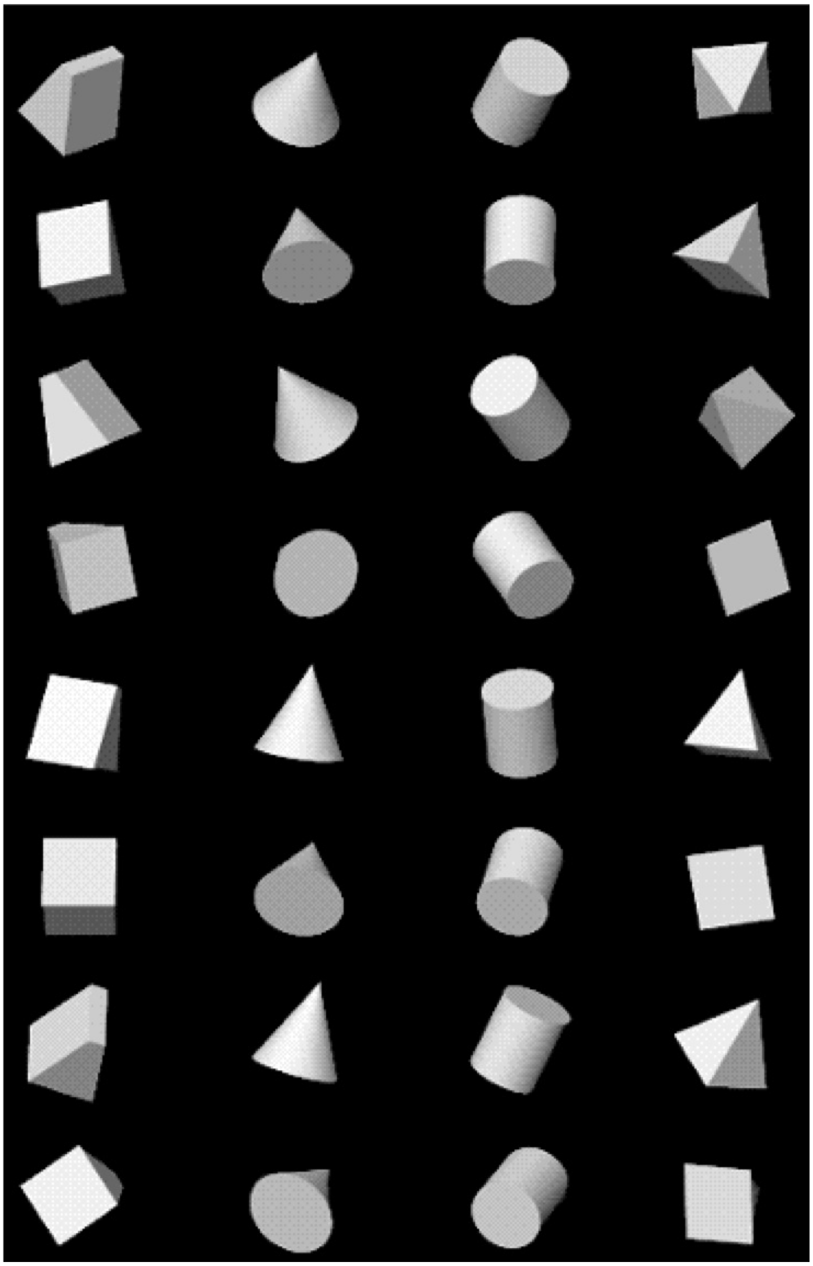

Eight views of four objects are shown above. Four different types of image information are shown below: shaded, large silhouette, line drawing, and small sillhouette. It had previously been shown that humans can name objects just as rapidly for line-drawing versions as for fully shaded versions of the same objects (Biederman & Ju, 1988). But it isn't clear whether this may be due to high efficiency for line-drawings, i.e. that the visual system is in some sense "well-tuned" for line drawings. Let's look specifically at how an ideal recognizer can be calculated that discounts views.

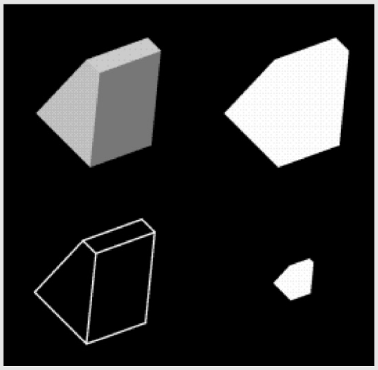

Let **X** = the vector describing the image data. Let $\boldsymbol{O_i}$ represent object i, where i = 1 to N. Suppose that $\boldsymbol{O_i}$ is represented in memory by M "snap shots" of each object, call them views (or templates) $\boldsymbol{V_{ij}}$, where j = 1, M. Given image data X, the posterior probability of $\boldsymbol{O_i}$ is computed by integrating out view, i.e. summing or "marginalizing" with respect to the M viewpoints.

$$p\,(O_i \mid X) = \sum_{j=1}^{M} p\,(V_{ij} \mid X)$$

$$= \sum_{j=1}^{M} \frac{p\,(X \mid V_{ij})\, p\,(V_{ij})}{p\,(X)}$$

Given image data, Ideal observer chooses object category *i* that maximizes the posterior $p\,(O_i \mid X)$. Given that p(X) is constant for a given observation X, the optimal strategy is equivalent to choosing i that maximizes:

$$L\,(i) = \sum_{j=1}^{M} p\,(X \mid V_{ij})\, p\,(V_{ij})$$

Tjan et al. used i.i.d additive gaussian noise in their experiment (as we did for the signal-known-exactly detection ideal near the beginning of the course), so the precise expression for the likelihood is:

$$p\,(X \mid V_{ij}) = \frac{1}{\left(\sigma \sqrt{2\,\pi}\right)^{p}} \exp\left(-\frac{1}{2\,\sigma^2} \|X - V_{ij}\|^2\right)$$

where p is the number of pixels in the image. The sum can then be computed given a particular pattern of image intensities X.

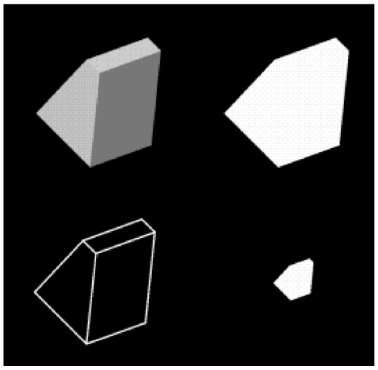

Tjan et al.'s experimental results showed that size, spatial uncertainty and detection efficiency played large roles in accounting for human object recognition efficiency.

Interestingly, highest recognition efficiencies (~7.8%) were found for small silhouettes of the objects, compared with (2.7%) fir line-drawings. (The small silhouettes were 0.7 deg, vs. 2.4 deg for the large silhouettes). *Detection* efficiencies, however, were about the same for both.

TABLE 1. Human object-recognition efficiency under different rendering conditions

	Subject BT		Subject WB	
	Sampling efficiency (%)	Total efficiency (%)	Sampling efficiency (%)	Tot efficie (%
Shaded objects	3.33 ± 0.07	3.09 ± 0.09	3.96 ± 0.24	3.47 ±
Large silhouettes	5.15 ± 0.32	4.68 ± 0.17	4.71 ± 0.08	4.33 ±
Line drawings	3.03 ± 0.06	2.82 ± 0.06	2.98 ± 0.06	2.57 ±
Small silhouettes	8.47 ± 0.16	8.38 ± 0.16	7.99 ± 2.06	7.30 ±

± intervals indicate ± 1 SE.

- **Now more on Cue integration and "Explaining away"...**

Cue integration: When a single model causes more than one measurement

- **Weak fusion**

Recall the generative graph for cue integration

Clark & Yuille, Landy & Maloney, Knill & Kersten, Schrater & Kersten.

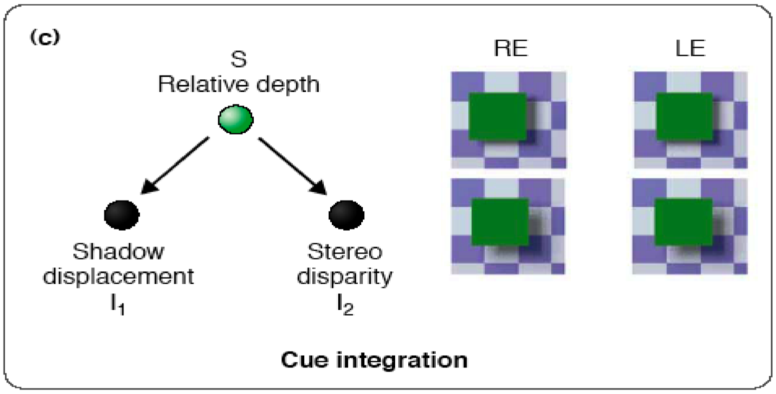

This Bayes net describes the factorization:

p(S,I1,I2) = p(I1|S)p(I2|S)p(S)

Maximum a posteriori observer for cue integration: conditionally independent cues

We'll change notation, and let x_1 and x_2 be image measurements or cues. The simple Bayes net shown above describes the case where the two cues are *conditionally independent*. In other words, p(x_1, x_2 | s) = p(x_1| s) p(x_2 | s).

Let's consider the simple Gaussian case where $x_i = s + n_i = \mu + n_i$. We'll show that optimal combined cue estimate is a weighted average of the cues.

p(s|x_1, x_2)=p(x_1, x_2 | s) p(s) /p(x_1, x_2) ∝ p(x_1| s) p(x_2 | s) = $e^{-(x_1-s)^2/2\sigma_1{}^2} e^{-(x_2-s)^2/2\sigma_2{}^2}$

$$\mathbf{PowerExpand}\left[\mathbf{Log}\left[E^{-(x_1-\mu)^2/(2\sigma_1{}^2)}\,E^{-(x_2-\mu)^2/(2\sigma_2{}^2)}\right]\right]$$

$$-\frac{(x_1-\mu)^2}{2\sigma_1^2}-\frac{(x_2-\mu)^2}{2\sigma_2^2}$$

What value of the variable μ = s gives the biggest posterior probability? We find the maximum by taking the derivative, and then finding the mean value of μ that give us zero.

$$D\left[-\frac{(x_1-\mu)^2}{2\sigma_1^2}-\frac{(x_2-\mu)^2}{2\sigma_2^2},\mu\right]$$

$$\frac{x_1-\mu}{\sigma_1^2}+\frac{x_2-\mu}{\sigma_2^2}$$

$$\text{Solve}\left[\frac{x_1-\mu}{\sigma_1^2}+\frac{x_2-\mu}{\sigma_2^2}==0,\mu\right]$$

$$\left\{\left\{\mu\to\frac{x_2\,\sigma_1^2+x_1\,\sigma_2^2}{\sigma_1^2+\sigma_2^2}\right\}\right\}$$

$$\left\{\left\{\mu\to\frac{x_2\,\sigma_1^2+x_1\,\sigma_2^2}{\sigma_1^2+\sigma_2^2}\right\}\right\}\,/.\,\left\{\sigma_1^2\to 1/r_1,\sigma_2^2\to 1/r_2\right\}$$

$$\left\{\left\{\mu\to\frac{\frac{x_1}{r_2}+\frac{x_2}{r_1}}{\frac{1}{r_2}+\frac{1}{r_1}}\right\}\right\}$$

where $r_i\left(=\frac{1}{\sigma_i{}^2}\right)$, is called the reliability.

$$\mu\to\frac{r_1\,x_1}{r_1+r_2}+\frac{r_2\,x_2}{r_1+r_2}$$

$$\mu\to\frac{r_1\,x_1}{r_1+r_2}+\frac{r_2\,x_2}{r_1+r_2}$$

In general, one can show that the combined estimate is the weighted sum of the separate estimates, μ_{cue1} and μ_{cue2} of the scene variable s measured under separate conditions. The weights w_i are determined by the relative reliabilities :

$$\hat{\mu}_{\text{combined}}=\hat{\mu}_{\text{cue1}}\,w_1+\hat{\mu}_{\text{cue2}}\,w_2=\hat{\mu}_{\text{cue1}}\frac{r_1}{r_1+r_2}+\hat{\mu}_{\text{cue2}}\frac{r_2}{r_1+r_2}.$$

■ An application to integrating cues from vision and haptics (touch)

When a person looks and feels an object, vision often dominates the integrated percept--e.g. the perceived size of an object is typically driven more strongly by vision than by touch. Why is this? Ernst and Banks showed that the reliability of the visual and haptic information determines which cue dominates. They first measured the variances associated with visual and haptic estimation of object size. They used these measurements to construct a maximum-likelihood estimator that integrates both cues. They concluded that the nervous system combines visual and haptic information in a fashion that is similar to a maximum-likelihood ideal observer. Specifically, visual dominance occurs when the variance associated with visual estimation is lower than that associated with haptic estimate.

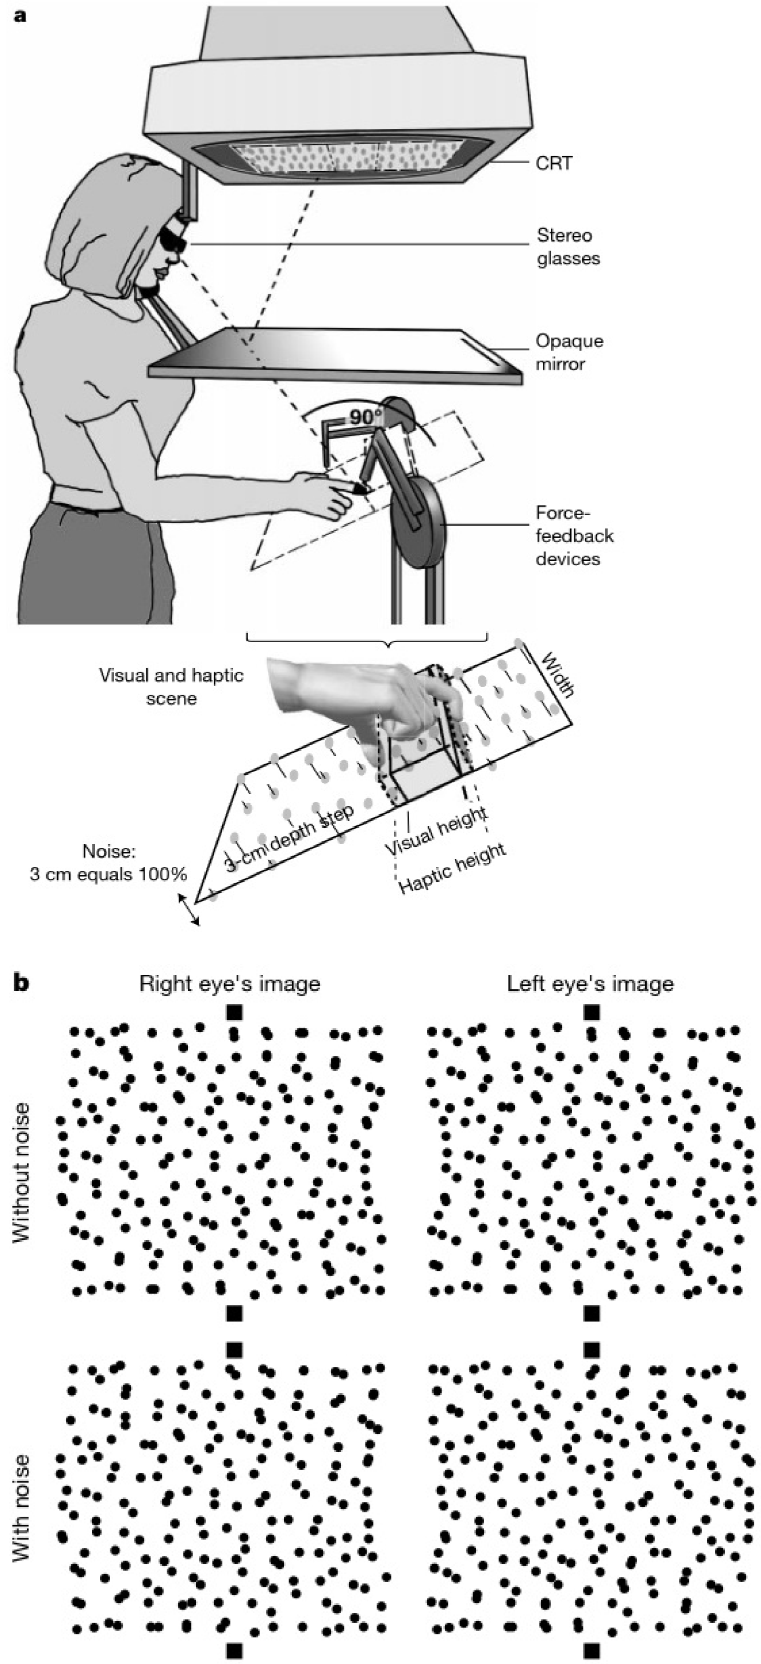
a
CRT
Stereo glasses
Opaque mirror
90°
Force-feedback devices
Visual and haptic scene
Width
3-cm depth step
Visual height
Haptic height
Noise: 3 cm equals 100%
b
Right eye's image
Left eye's image
Without noise
With noise

See Ernst MO, Banks MS (2002) Humans integrate visual and haptic information in a statistically optimal fashion. Nature 415:429-433.

Perceptual explaining away: More that one model working cooperatively to “explain” the data

Perception as puzzle solving

Rock, I. (1983). *The Logic of Perception*. Cambridge, Massachusetts: M.I.T. Press.

Strong fusion

In their 1990 book Data Fusion for Sensory Information Processing, Clark & Yuille distinguished weak from strong fusion. Strong fusion uses more sophisticated generative knowledge for how the image data arises. At several times during this course, we seen examples of “perceptual explaining away”, where the visual system uses its knowledge of how several processes interact to produce a complex image. For example, the ambiguous translating diamond. How the motion information gets integrated can be understood in the context of “cue integration”, as was described in Weiss, Adelson & Simoncelli. However, the selection process for what information should be integrated can be influenced by “auxiliary cues” for other processes, such as occlusion, that explains away missing information, such as the vertices of the diamond below.

Perceptual explaining away

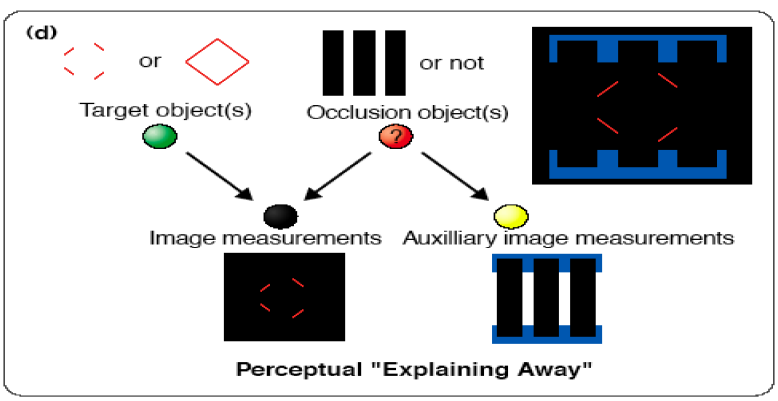

Both causes S1 and S2 can be primary variables.

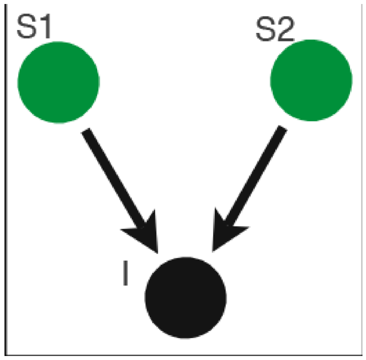

The above Bayes net describes the factorization:

p(S1,S2,I) = p(I|S2,S2) p(S1)p(S2)

If we average over I, S1 and S2 are independent. However, knowledge of I makes S1 and S2 conditionally dependent. The two causes S1 and S2 can behave like competing hypotheses to explain the data I.

In general, "explaining away" is a phenomenon that occurs in probabilistic belief networks in which two (or more) variables influence a third variable whose value can be measured (Pearl, 1988). Once measured, it provides evidence to infer the values of the influencing variables.

Imagine two coins that can be flipped independently, and the results (heads or tails) have an influence on a third variable. For concreteness, assume the third variable's value is 1 if both coins agree, and 0 if not (NOT-XOR). If we are ignorant of the value of the third variable, knowledge of one influencing variable doesn't help to guess the value of the other—the two coin variables are independent. (This is called marginal independence, "marginal" with respect to the third variable, I)

But if the value of the third variable is measured (suppose it is 1), the two coin variables become coupled, and they are said to be *conditionally dependent*. Now knowing that one coin is heads guarantees that the other one is too.

The phrase "explaining away" arises because coupling of variables through shared evidence often arises in human reasoning, when the influences can be viewed as competing causes.

Suppose we have a prior reason to believe that both coins are heads, but we believe the second coin is even more likely than the first to be heads. Now we make a measurement, and discover the evidence is 0. The evidence explains away the "competing hypothesis" that both coins are heads, and in particular that the first coin is heads, because the first coin must be tails if the second coin is heads.

Human reasoning is particularly good at these kinds of inferences.

"Explaining away" is also a characteristic of perceptual inferences, for example when there are alternative perceptual groupings consistent with a set of identical or similar sets of local image features.

Demonstrations of cooperative computation and explaining away in perception

In addition to occlusion and the ambiguous translating diamond, many other perceptual phenomena that we've seen before can be interpreted as "explaining away".

One of the challenges is to quantitatively test the extent to which human vision uses knowledge of the generative processes to solve the more complex “puzzles” of perception.

One would like to know, for example, when and the extent to which human vision may use uncertainty in an auxiliary cue to change beliefs about the primary hypothesis of interest. Whether it does or not requires quantitative tests.

Recently, Battaglia has worked out the theory for Gaussian processes. Battaglia PW (2010). Bayesian perceptual inference in linear Gaussian models, MIT Technical Report, MIT-CSAIL-TR-2010-046. http://hdl.handle.net/1721.1/58669

Dependence of shape on perceived light source direction

Explaining away can be interpreted as cooperation through “module” competition.

Dependence of shape on perceived light source direction

This can be illustrated with shape and illumination direction. There may be no auxiliary cue, but both types of variables interact to determine the shading pattern.

It may be the case that p(shape,illumination direction) ~ p(shape) p(illumination direction). However, this doesn’t imply that the variables remaing independent given a specific image. They may become conditionally dependent given shading.

Conditional dependence means that: p(shape, illumination | shading) can **not** be factored into separate terms, p(shape | shading) p(illumination | shading). This is because shape and illumination are tied via a functional relationship with image measurements: shading = f(shape, illumination).

If the internal belief that light source direction changes from left to right, then the most probably shape estimate changes from convex to concave (or vice versa, depending on the shading direction).

Brewster (1926), Gibson, Ramachandran, V. S. (1990), crater illusion and the single light source assumption

Adams, W. J., Graf, E. W., & Ernst, M. O. (2004). Experience can change the 'light-from-above' prior. Nat Neurosci, 7(10), 1057-1058.

■ Vertical light direction

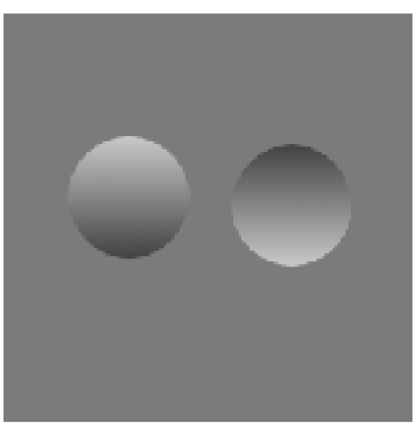

■ Horizontal light direction

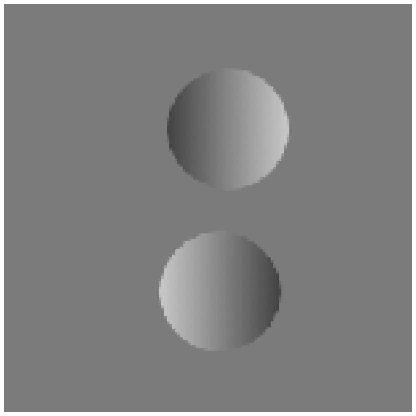

Review: Occlusion & motion: Lorenceau & Shiffrar, Sinha

Recall translating diamond used to illustrate the aperture problem.

When the diamond is seen as coherently translating, one often also interprets the vertices as being covered by rectangular occluders.

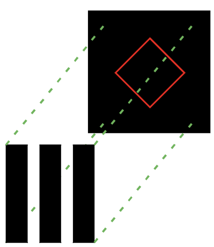

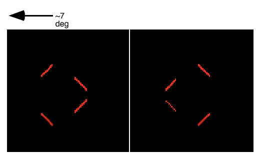

■ Translating diamond with "occluding occluders"

Occlusion as explaining away:

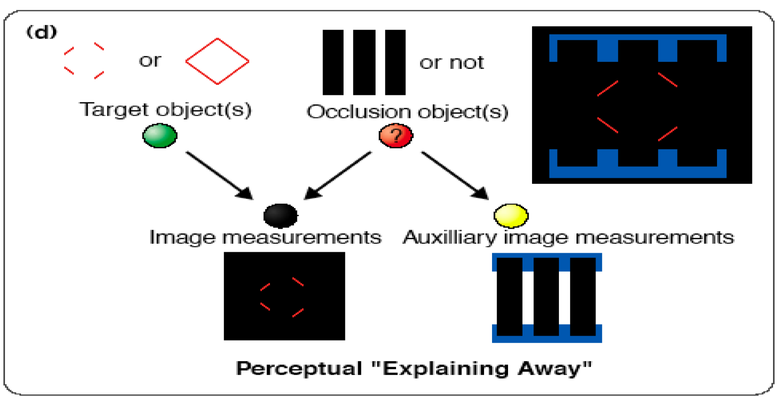

Lightness & surface geometry

■ Mach card

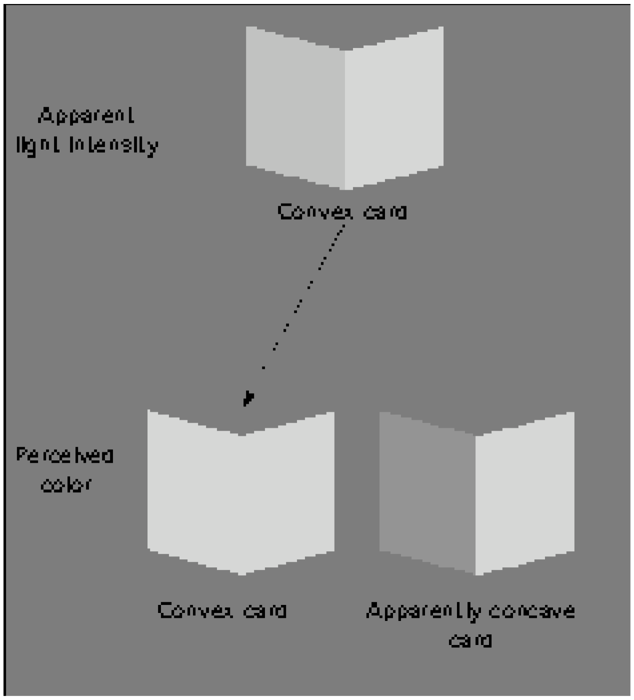

■ Lightness and shape

Recall the lightness demonstration that is similar to the Craik-O'Brien-Cornsweet effect, but difficult to explain with a simple filter mechanism (Knill, D. C., & Kersten, D. J., 1991). The idea is that the lightness of a pair of luminance gradients on the left of the figure below look different, whereas they look similar for the pair luminance gradients on the right. The reason seems to be due to the fact that the luminance gradients on the right are attributed to smooth changes in shape, rather than smooth changes in illumination.

http://vision.psych.umn.edu/www/kersten-lab/demos/lightness.html

These demonstrations suggest the existence of scene representations in our brains for shape, reflectance and light source direction.

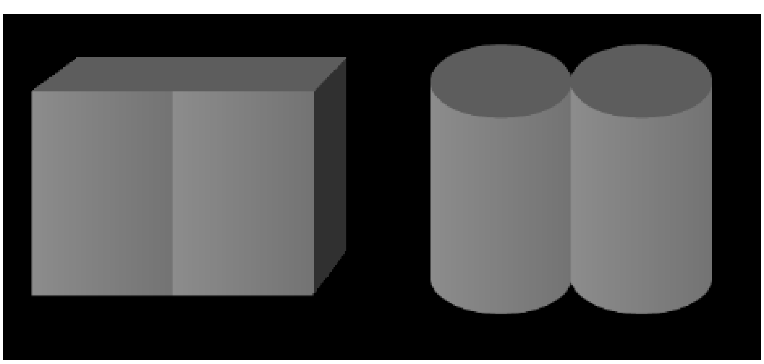

Draw a diagram to illustrate the above illusion in terms of "explaining away"

■ Dependence of lightness on spatial layout

In the 1970's, Alan Gilchrist was able to show that the lightness of a surface patch may be judged either dark-gray, or near-white with only changes in perceived spatial layout (Gilchrist, A. L. (1977). How did he do this? What is going on? Interpret lightness as reflectance estimation.

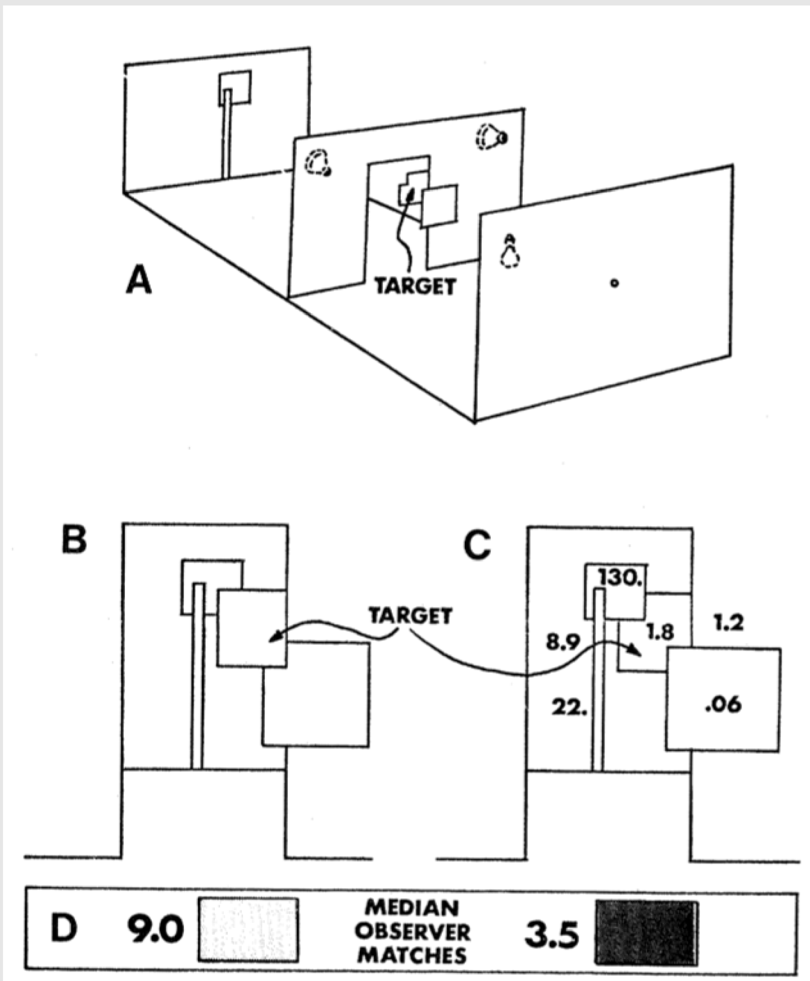

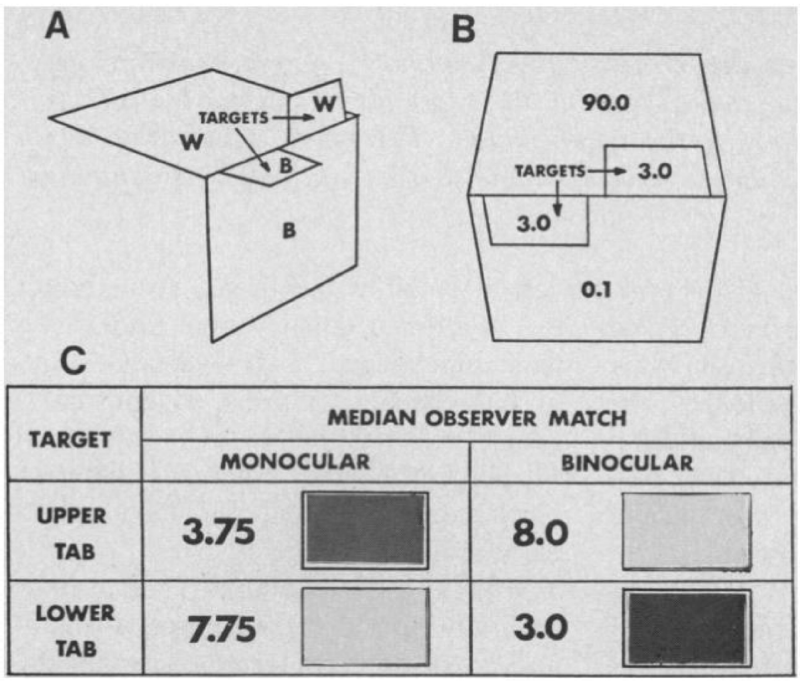

o The Room-in-a-Shoe-Box experiment

o Coplanar card experiment

The left and right inner gray disks in the above figure are the same intensity. In classic simultaneous contrast, the brighter annulus on the right makes the inner disk appear darker.

Review: Color & shape

■ Bloj, Kersten & Hurlbert

Demo

http://gandalf.psych.umn.edu/users/kersten/kersten-lab/Mutual_illumination/BlojKerstenHurlbertDemo99.pdf

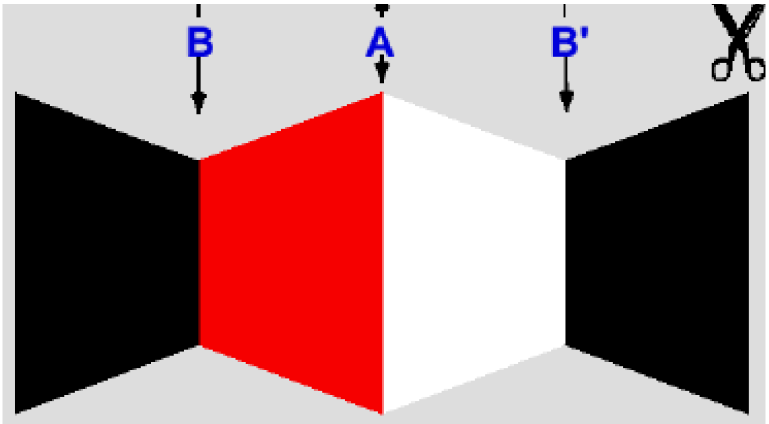

Interpretation

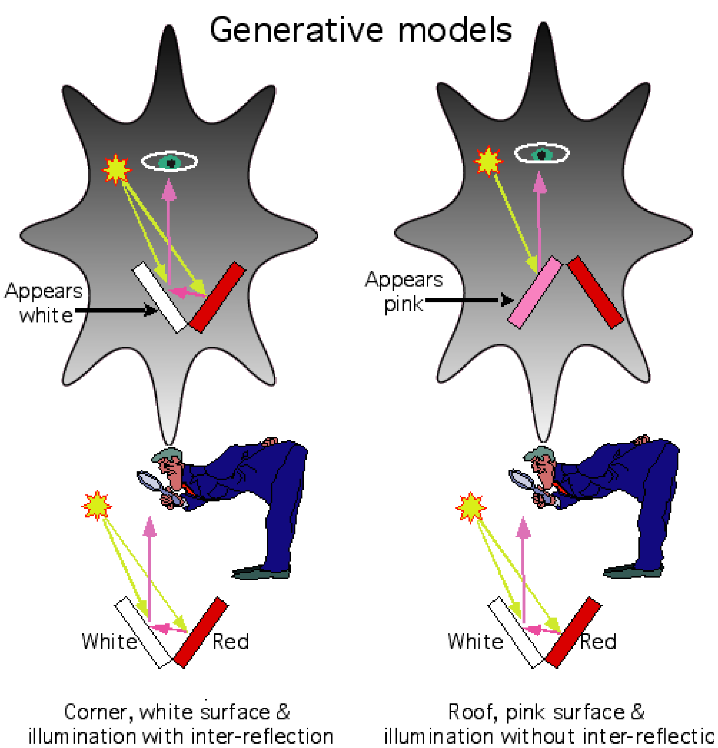

Interreflection as explaining away. Stereo can be used as an auxiliary cue to change the perceived shape from concave to convex.

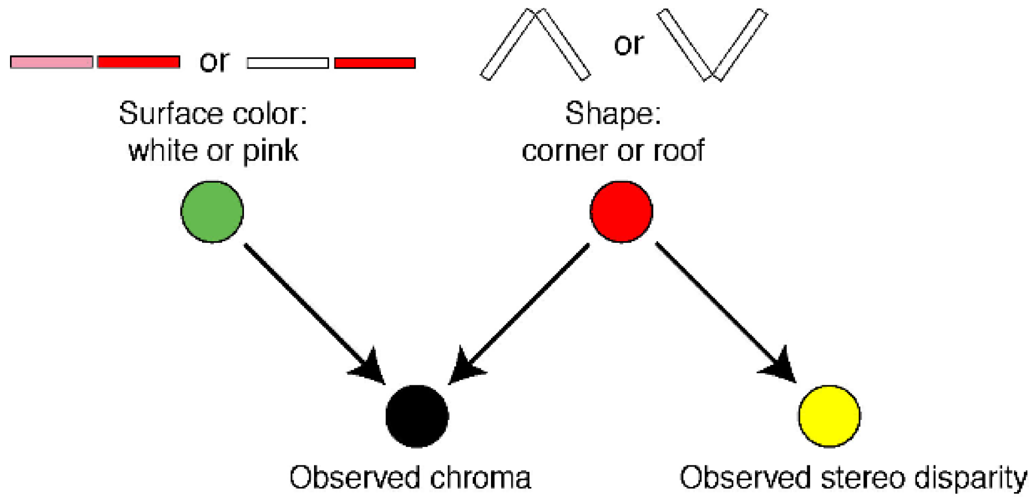

Review: Transparency and structure-from-motion

■ Motion and transparency (Kersten et al., 1992)

Dependence of transparency on perceived depth

Kersten and Bülthoff

o orientation and transparency

o transparency and depth from motion--computer demo

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

Nakayama, Shimojo (1992)

o transparency and depth from stereo demos, neon color spreading

Material shininess and surface curvature

Bruce Hartung, showed that visual perception takes into account 3D shape (curvature) when inferring whether a surface is matte or shiny. This inference requires a prior assumption that the statistics of the pattern of illumination in a complex environment is approximately radially uniform. In other words, illumination sources rarely have high aspect ratios (e.g. long flourescent lights would be considered rare compared with other sources). Hartung, B., & Kersten, D. (2003). How does the percep- tion of shape interact with the perception of shiny ma- terial? [Abstract] *Journal of Vision*, *3*(9), 59a, http://journalofvision.org/3/9/59/, doi:10.1167/3.9 .59.

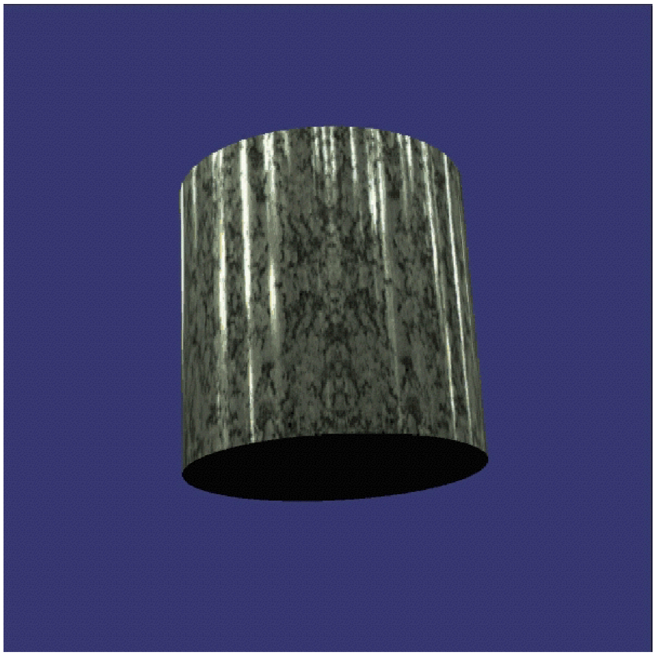

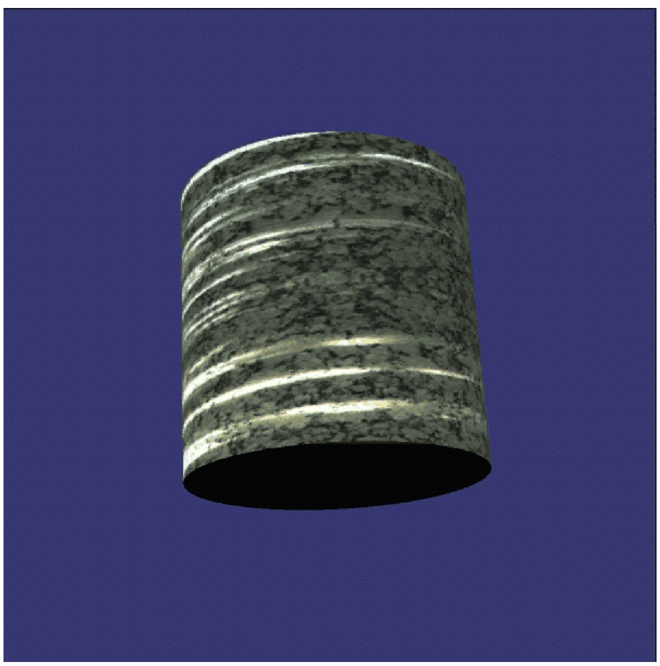

Application to image parsing, object recognition

■ **Incorporating higher-level knowledge--Image parsing and recognition using cooperative computation**

The material in Lectures 24 and 25, described computer vision work showing how competing object class models could compete to “explain away” false positives, see: Tu Z, Zhu S-C (2002), Zhu and Tu (2000). For a review, see: Yuille and

Kersten (2006). Such a mechanism, perhaps as argued involving analysis-by-synthesis mechanisms, may be particularly important given occlusion and a high amount of clutter.

http://gandalf.psych.umn.edu/users/kersten/kersten-lab/courses/Psy5036W2008/Lectures/24.%20ObjectRecognitionClutter/Lecture24TopDown.pdf

References

Cue integration, cooperative computation

Adams, W. J., Graf, E. W., & Ernst, M. O. (2004). Experience can change the 'light-from-above' prior. Nat Neurosci, 7(10), 1057-1058.

Barrow, H. G., & Tenenbaum, J. M. (1978). Recovering Intrinsic Scene Characteristics from Images. In A. R. Hanson, & E. M. Riseman (Ed.), Computer Vision Systems (pp. 3-26). New York: Academic Press.

Battaglia PW (2010). Bayesian perceptual inference in linear Gaussian models, MIT Technical Report, MIT-CSAIL-TR-2010-046. http://hdl.handle.net/1721.1/58669

Bergstrom, S. S. (1977). Common and Relative Components of Reflected Light as Information About the Illumination, Colour and Three-Dimensional Form of Objects. 18, 180-186).

Biederman, I., & Ju, G. (1988). Surface versus edge-based determinants of visual recognition. Cogn Psychol, 20(1), 38-64.

Bishop, C. M. (2006). Pattern recognition and machine learning. New York: Springer.

Brewster, D. (1826). On the optical illusion of the conversion of cameos into intaglios and of intaglios into cameos, with an account of other analogous phenomena. Edinburgh Journal of Science, 4, 99-108.

Clark, J. J., & Yuille, A. L. (1990). Data Fusion for Sensory Information Processing . Boston: Kluwer Academic Publishers.

Ernst MO, Banks MS, Bulthoff HH (2000) Touch can change visual slant perception. Nat Neurosci 3:69-73.

Ernst MO, Banks MS (2002) Humans integrate visual and haptic information in a statistically optimal fashion. Nature 415:429-433.

Gilchrist, A. L. (1977). Perceived Lightness Depends on Perceived Spatial Arrangement. Science, 195, 185-187.

Gibson, J. J. (1950). The Perception of the Visual World. Boston, MA: Houghton Mifflin.

Hartung, B., & Kersten, D. (2003). How does the percep- tion of shape interact with the perception of shiny ma- terial? [Abstract] Journal of Vision, 3(9), 59a, http://journalofvision.org/3/9/59/, doi:10.1167/3.9 .59.

Hillis JM, Ernst MO, Banks MS, Landy MS (2002) Combining sensory information: mandatory fusion within, but not between, senses. Science 298:1627-1630.

Humphrey, K. G., Goodale, M. A., Bowen, C. V., Gati, J. S., Vilis, T., Rutt, B. K., & Menon, R. S. (1996). Differences in Perceived Shape from Shading Correlate with Activity in Early Visual Areas.,, 1-16.

Jacobs, R. A., Jordan, M. I., Nowlan, S. J., & Hinton, G. E. (1991). Adaptive mixtures of local experts. Neural Computation, 3, 79-87.

Jacobs RA (2002) What determines visual cue reliability? Trends Cogn Sci 6:345-350.

Kersten, D. J. (1991). Transparency and the Cooperative Computation of Scene Attributes. In M. Landy, & A. Movshon (Ed.), Computational Models of Visual Processing. Cambridge, Massachusetts: M.I.T. Press.

Kersten, D., Bülthoff, H. H., Schwartz, B., & Kurtz, K. (1992). Interaction between transparency and structure from motion. Neural Computation, 4(4), 573-589.

Kersten D, Yuille A (2003) Bayesian models of object perception. Current Opinion in Neurobiology 13:1-9.

Kersten D, Mamassian P, Yuille A (2004) Object perception as Bayesian Inference. Annual Review of Psychology 55:271-304.

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

Mach, E. (1886, 1959). The Analysis of Sensations . New York: Dover.

Nakayama, K., & Shimojo, S. (1992). Experiencing and perceiving visual surfaces. Science, 257, 1357-1363.

Pearl J (1988) Probabilistic reasoning in intelligent systems : networks of plausible inference, Rev. 2nd printing. Edition. San Mateo, Calif.: Morgan Kaufmann Publishers.

Poggio, T., Torre, V., & Koch, C. (1985). Computational vision and regularization theory. Nature, 317, 314-319.

Poggio, T., Gamble, E. B., & Little, J. J. (1988). Parallel integration of vision modules. Science, 242, 436-440.

Ramachandran, V. S. (1990). Visual perception in people and machines. In A. Blake, & T. Troscianko (Ed.), A.I. and the Eye John Wiley & Sons Ltd.

Ripley, B. D. (1996). Pattern Recognition and Neural Networks . Cambridge, UK: Cambridge University Press.

Schrater PR, Kersten D (2000) How optimal depth cue integration depends on the task. International Journal of Computer Vision 40:73-91.

Tjan B., Braje, W., Legge, G.E. & Kersten, D. (1995) Human efficiency for recognizing 3-D objects in luminance noise.Vision Research, 35, 3053-3069.

Todd, J. T., & Mingolla, E. (1983). Perception of Surface Curvature and Direction of Illumination from Patterns of Shading. Journal of Experimental Psychology: Human Perception & Performance, 9(4), 583-595.

Tu Z, Zhu S-C (2002) Parsing Images into Region and Curve Processes. In: Proc. of the 7th European Conference on Computer Vision, p 393 ff. Copenhagen, Denmark: Springer-Verlag, Berlin Heidelberg.

von der Heydt R, Friedman H, Zhou HS (2003) Searching for the neural mechanisms of color filling-in. In: Filling-in: From Perceptual Completion to Cortical Reorganization (Pessoa L, P DW, eds), pp 106-127. Oxford: Oxford University Press.

Yuille, A. L., & Bülthoff, H. H. (1996). Bayesian decision theory and psychophysics. In K. D.C., & R. W. (Ed.), Perception as Bayesian Inference Cambridge, U.K.: Cambridge University Press.

Yuille, A., & Kersten, D. (2006). Vision as Bayesian inference: analysis by synthesis? Trends Cogn Sci, 10(7), 301-308. http://gandalf.psych.umn.edu/~kersten/kersten-lab/papers/yuillekerstenTICs2006.pdf

Zhu S-C, Zhang R, Tu Z (2000) Integrating Bottom-up/Top-Down for Object Recognition by Data Driven Markov Chain Monte Carlo. In: Proc. of Int'l Conf. on Computer Vision and Pattern Recognition. SC.

© 2000, 2004,2006, 2008, 2010 Daniel Kersten, Computational Vision Lab, Department of Psychology, University of Minnesota.
kersten.org