Visual Perception of Surface Materials

Roland W. Fleming

Max Planck Institute
for Biological Cybernetics

Cognitive & Computational Psychophysics
Max Planck Institute for Biological Cybernetics

Visual Perception of Material Qualities

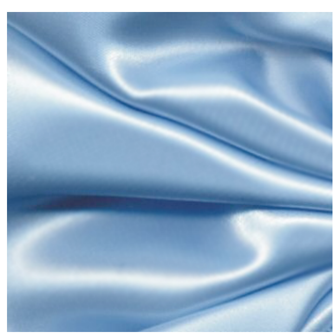

- Different materials, such as quartz, satin and chocolate have distinctive visual appearances
- Without touching an object we usually have a clear impression of what it would feel like: *hard* or *soft*; *wet* or *dry*; *rough* or *smooth*

2

Everyday life

3

Material-ism

4

The state of things

Edibility

Usability

7

Dung

8

Materials in Western Art

- The appearance of “stuff” was a major preoccupation in C17th Dutch art:
  - Sensual quality of things
  - Stasis, and contemplation in Still Life
- Contrast with Italian Renaissance:
  - Perspective and the spatial arrangement of things
  - Drama, events and dynamism.

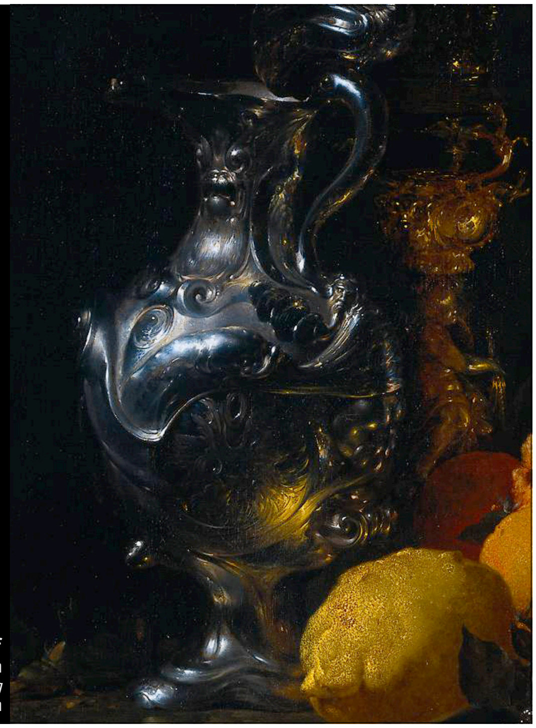

Willem Kalf
detail from
Still Life with Silver Jug
Rijksmuseum, Amsterdam

Materials in Computer Graphics

- Hollywood and the games industry know that photorealism means getting materials right
- Henrik Wann Jensen, Steve Marschner and Pat Hanrahan won a technical Oscar in 2004 for their work on Subsurface scattering

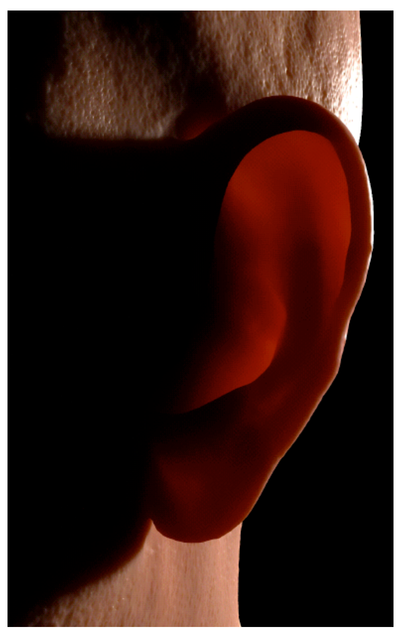

Henrik Wann Jensen and Craig Donner

10

The visual brain

- Towards a neuroscience of material perception

Research Questions

- What gives a material its characteristic appearance?
- What cues does the human brain use to identify materials across a wide variety of viewing conditions?

Outline

- Visual estimation of surface reflectance properties: *gloss*
- Perception of materials that transmit light
  - *Refraction*
  - *Sub-surface scattering*
- Exploiting the assumptions of the visual system to edit the material appearance of objects in photographs

Surface Reflectance

- These spheres look different because they have different **surface reflectance properties**.
- Everyday language: *colour*, *gloss*, *lustre*, etc.

Images: Ron Dror

Surface Reflectance

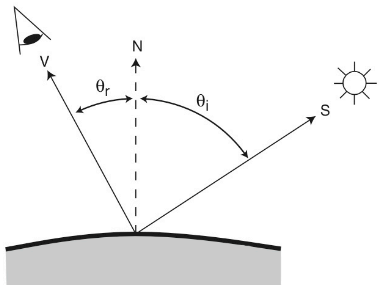

BRDF

$f(\theta_i, \phi_i; \theta_r, \phi_r)$

- **Visual system's goal:** Estimate the BRDF

Images: Ron Dror

Confounding Effects of Illumination

- Identical materials can lead to very different images
- Different materials can lead to very similar images

Ambiguity between Reflectance and Illumination

$a(f) = 1 / f^{\mu}$

$\mu = 0$ $\mu = 0.4$ $\mu = 0.8$ $\mu = 1.2$

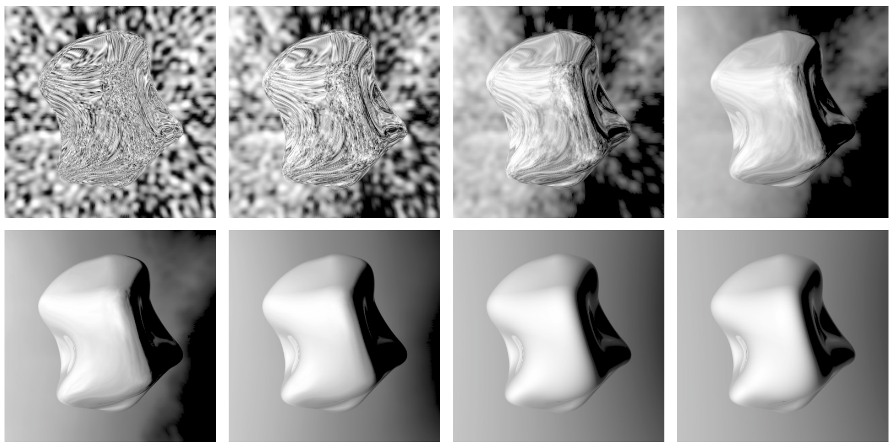

$\mu = 1.6$ $\mu = 2.0$ $\mu = 4.0$ $\mu = 8.0$

17

Ambiguity between Reflectance and Illumination

Under arbitrary combinations of reflectance and illumination, the problem is ill-posed (unsolvable)

18

Hypothesis

Humans exploit *statistical regularities of real-world illumination* in order to eliminate unlikely image interpretations

19

Real-world Illumination

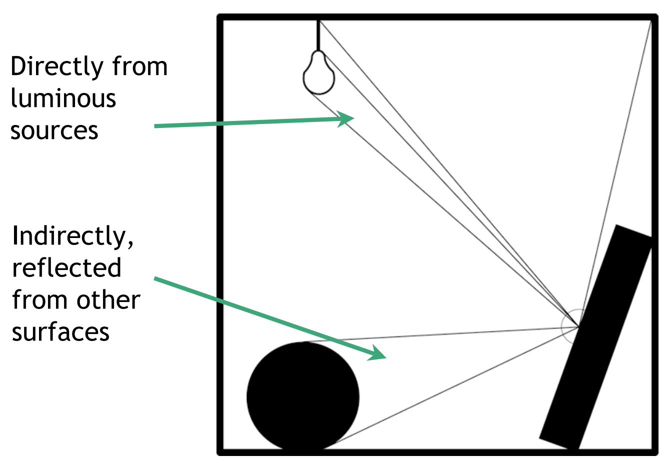

Illumination at a point in space: amount of light arriving from every direction.

20

Photographically captured light probes (Debevec *et al.*, 2000)

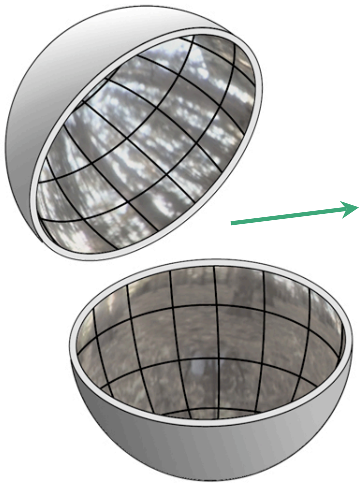

Panoramic projection

21

Statistics of typical Illumination (Dror et al. 2004)

- Intensity histogram is heavily skewed
  - Few direct light sources

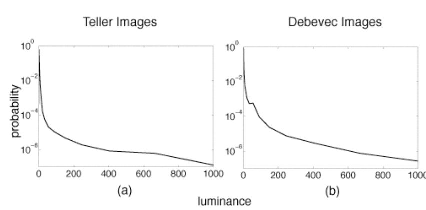

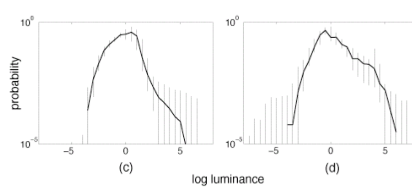

22

Statistics of typical Illumination (Dror et al. 2004)

- Typical 1/f amplitude spectrum

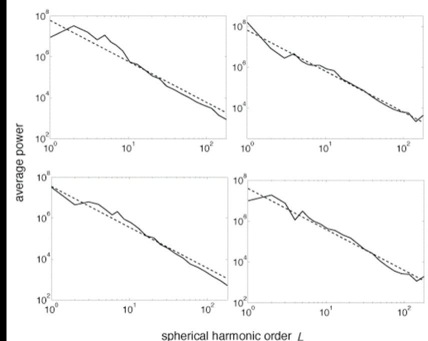

23

Hypothesis

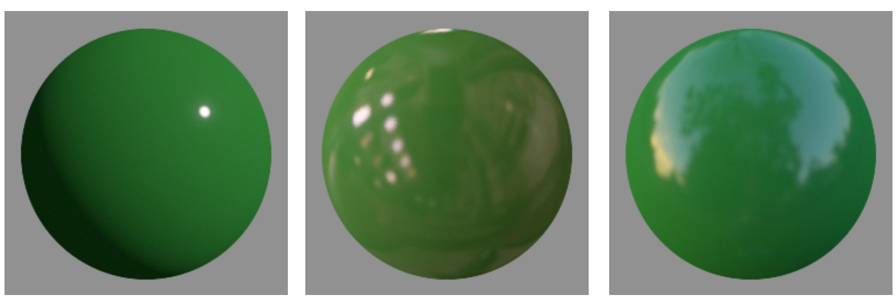

Humans exploit statistical regularities of real-world illumination in order to *eliminate unlikely image interpretations*

24

Dismissing unlikely interpretations

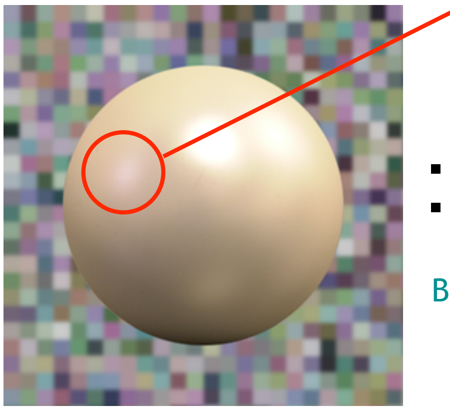

Blurry feature

2 interpretations:

- Sharp reflection, blurry world
- Blurry reflection, sharp world

But the world *usually* isn't blurry!

Therefore it is probably a blurry *reflection*

25

Dismissing unlikely interpretations

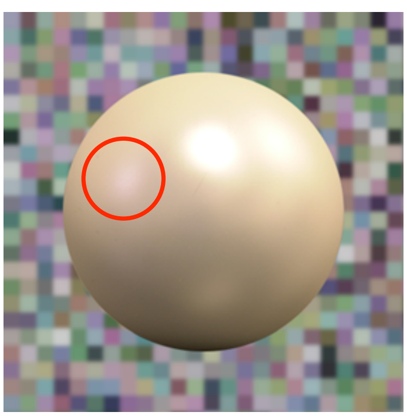

- In practice, unlikely image interpretations do not need to be explicitly entertained
- Under typical illumination conditions, different materials yield **diagnostic image features** that are responsible for their 'look'
- **The brain doesn't need to model the physics, it just needs to look out for tell-tale image features**

26

Observations

- Context has surprisingly little effect on apparent gloss

27

Findings

- Subjects are good at judging surface reflectance across real-world illuminations: '*gloss constancy*'
- Subjects are poorer at estimating gloss under illuminations with atypical statistics.

Beach Light probe | **Campus Light probe** | **Single point source** | **Gaussian 1/f Noise**

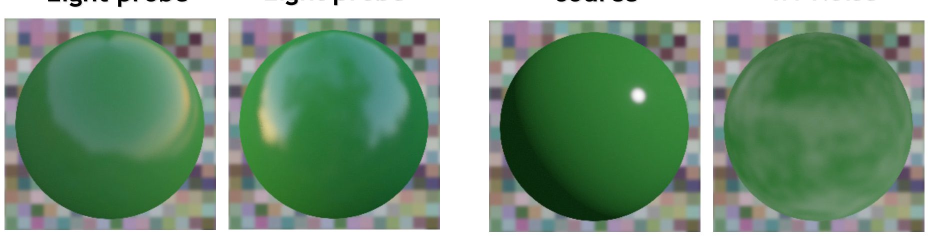

28

Illuminations have skewed intensity histograms

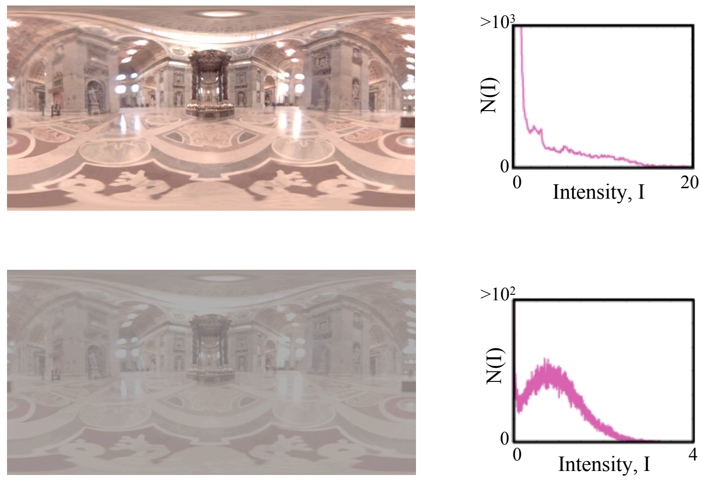

29

Illumination distribution is important …

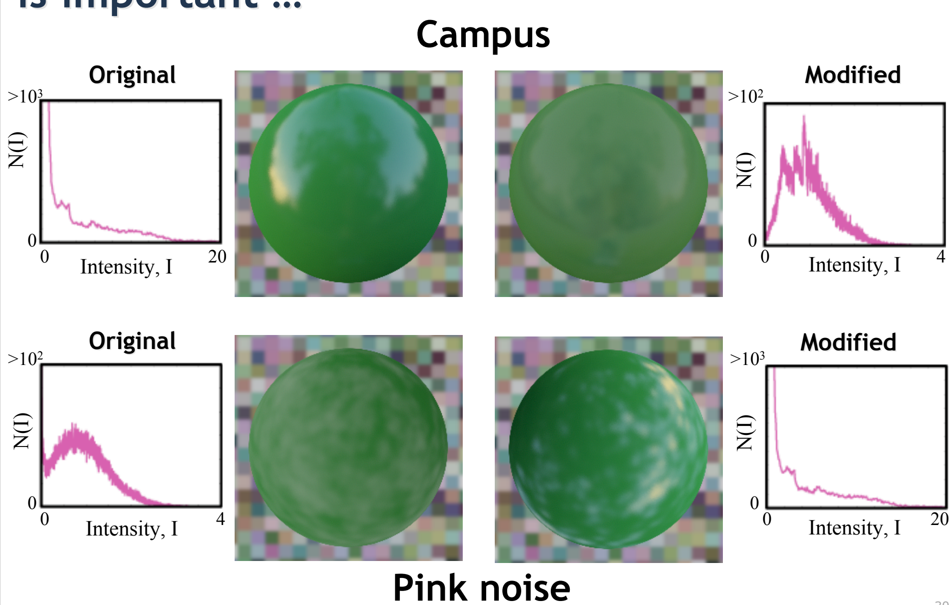

30

… but isn't everything

Campus original

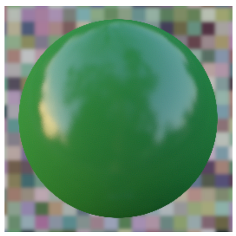

White noise with histogram of **Campus**

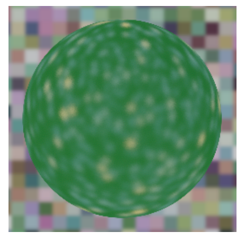

Heeger-Bergen texture synthesis

Input texture

Synthesized texture

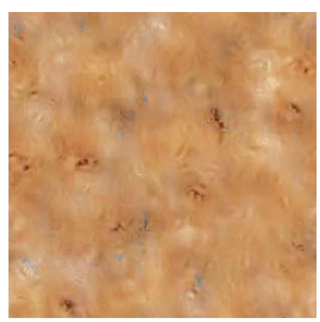

Taken from *Pyramid-Based Texture Analysis/Synthesis*

Treat illumination maps as if they are stochastic texture

Wavelet Statistics

Synthetic illuminations with same wavelet statistics as real-world illuminations

Beach | Building | Campus | Eucalyptus

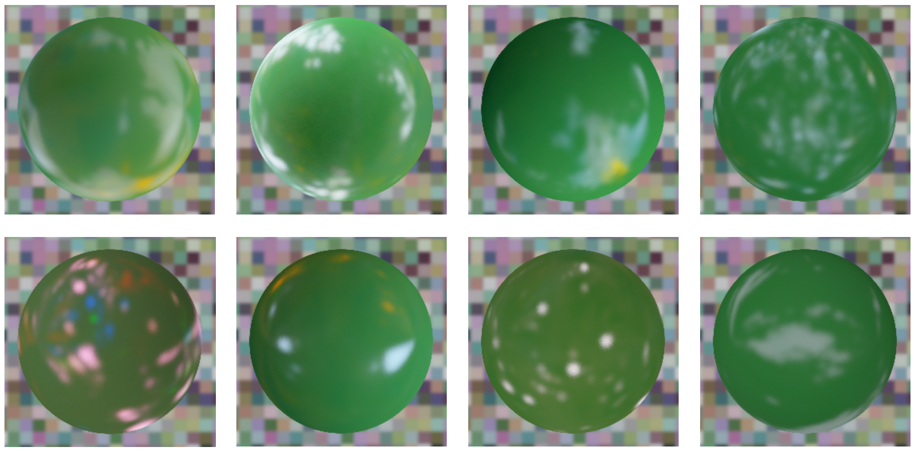

Grace | Kitchen | St. Peter's | Uffizi

33

- Image statistics are a powerful shortcut:
- Allow the brain to recognize glossy materials without explicitly estimating the BRDF
- However, when the statistics are infringed, perception can fail

34

Outline

- Visual estimation of surface reflectance properties: *gloss*
- Perception of materials that transmit light
  - *Refraction*
  - *Sub-surface scattering*
- Exploiting the assumptions of the visual system to edit the material appearance of objects in photographs

Previous work on materials that transmit light

Adelson E.H. (1993). Perceptual organization and the judgment of brightness. *Science*. **262**(5142): 2042-2024.
Adelson, E. H. (1999). Lightness perception and lightness illusions. In In *The New Cognitive Neurosciences*. 2nd edn. (ed. Gazzaniga, M.S.) 339—351 (MIT Press, Cambridge, Massachusetts, 1999).
Adelson, E. H., and Anandan, P. (1990). Ordinal characteristics of transparency. *Proceedings of the AAAI-90 Workshop on Qualitative Vision*, 77- 81.
Anderson B.L. (1997). A theory of illusory lightness and transparency in monocular and binocular images: the role of contour junctions. *Perception*. **26**(4): 419-453.
Anderson, B.L. (1999). Stereoscopic surface perception. *Neuron*. **24**: 991-928.
Anderson, B.L. (2001). Contrasting theories of White's illusion. *Perception*, **30**: 1499-1501.
Anderson, B.L. (2003). The Role of Occlusion in the Perception of Depth, Lightness, and Opacity. *Psychological Review*, **110**(4): 785-801.
Anderson B.L. (2003). Perceptual organization and White's illusion. *Perception*. **32**(3): 269-84.
Beck J. (1978). Additive and subtractive color mixture in color transparency. *Percept Psychophys*. **23**(3): 265-267.
Beck J, Prazdny K, Ivry R. (1984). The perception of transparency with achromatic colors. *Percept Psychophys*. **35**(5): 407-422.
Beck J, Ivry R. (1988). On the role of figural organization in perceptual transparency. *Percept Psychophys*. **44**(6): 585-594.
Chen, J.V. & D'Zmura, M. (1998). Test of a convergence model for color transparency perception. *Perception*, **27**: 595-608.
D'Zmura, M., Colantoni, P., Knoblauch, K. & Laget, B. (1997). Color transparency, *Perception*, **26**, 471-492.
D Zmura, M., Rinner, O., & Gegenfurtner, K. R. (2000). The colors seen behind transparent filters. *Perception*, **29**, 911-926.
Gerbino W., Stultiens C.I., Troost J.M. and de Weert C.M. (1990). Transparent layer constancy. *J Exp Psychol Hum Percept Perform*. **16**(1): 3-20.
Gerbino, W. (1994). Achromatic transparency. In A.L. Gilchrist, Ed. *Lightness, Brightness and Transparency*. (pp 215-255). Hove, England: Lawrence Erlbaum.
Heider, G. M. (1933). New studies in transparency, form and color. *Psychologische Forschung*. **17**: 13—56.
Kersten, D. (1991) Transparency and the cooperative computation of scene attributes. In *Computational Models of Visual Processing*, M.I.T. Press, Landy M and Movshon A., Eds..
Kersten, D. and Bülthoff, H. (1991) Transparency affects perception of structure from motion. Massachusetts Institute of Technology, Center for Biological Information Processing *Tech. Memo, 36*.
Kersten, D., Bülthoff, H. H., Schwartz, B., & Kurtz, K. (1992). Interaction between transparency and structure from motion. *Neural Computation*, **4**(4), 573-589.
Khang BG, Zaidi Q. (2002). Cues and strategies for color constancy: perceptual scission, image junctions and transformational color matching. *Vision Res*. 42(2):211-26. Erratum in: *Vision Res*. 2002. **42**(24): 2729.
Khang, B.-G., & Zaidi, Q. (2002). Accuracy of color scission for spectral transparencies. *Journal of Vision*, *2*(6), 451-466, doi:10.1167/2.6.3.
Kanizsa, G. (1955). Condiziono ed effetti della transparenza fenomica. *Revista di Psicologia*. **49**: 3—19.
Koffka, K. (1935). *Principles of Gestalt Psychology*. Cleveland: Harcourt, Brace and World.
Lindsey DT, Todd JT. (1996). On the relative contributions of motion energy and transparency to the perception of moving plaids. *Vision Res*. **36**(2): 207-22.
Metelli, F. (1970). An algebraic development of the theory of perceptual transparency. *Ergonomics*, **13**(1): 59-66.
Metelli, F. (1974a). Achromatic color conditions in the perception of transparency. In Eds. Macleod, R. B., & Pick, H. L. (Eds.) *Perception* (pp. 95—116). Ithaca, NY: Cornell University Press.
Metelli F. (1974b). The perception of transparency. *Scientific American*, **230**(4): 90-98.
Metelli, F. (1985). The simulation and perception of transparency. *Psychological Research*, **47**(4): 185-202.
Metelli, F., Da Pos, O. and Cavedon, A. (1985). Balanced and unbalanced, complete and partial transparency. *Perception & Psychophysics*, 38(4): 354-66.
Masin SC. (1999). Color scission and phenomenal transparency. *Percept Mot Skills*. **89**(3 Pt 1): 815-23.
Masin SC. (1999) Phenomenal transparency in achromatic checkerboards. *Percept Mot Skills*. **88**(2): 685-92.
Masin SC. (1997). The luminance conditions of transparency. *Perception*. **26**(1): 39-50.
Nakayama, K., & Shimojo, S. (1992). Experiencing and perceiving visual surfaces. *Science*. **257**: 1357-1363.
Nakayama, K., Shimojo, S., & Ramachandran, V. (1990). Transparency: Relation to depth, subjective contours, luminance, and neon color spreading. *Perception*, **19**: 497-513.
Plummer DJ, Ramachandran VS. (1993). Perception of transparency in stationary and moving images. *Spat Vis*. **7**(2): 113-23.
Robilotto R, Khang BG & Zaidi Q. (2002). Sensory and physical determinants of perceived achromatic transparency. *JOV*. **2**(5): 388-403.
Robilotto, R., & Zaidi, Q. (2004). Perceived transparency of neutral density filters across dissimilar backgrounds. *Journal of Vision*, *4*(3), 183-195, doi:10.1167/4.3.5.
Singh M, Anderson BL. (2002). Perceptual assignment of opacity to translucent surfaces: the role of image blur. *Perception*. **31**(5): 531-52.
Singh M, Anderson BL. (2002). Toward a perceptual theory of transparency. *Psychol Rev*. **109**(3): 492-519.
Somers, D. C., and Adelson, E. H. (1997). Junctions, Transparency, and Brightness. Investigative Ophthalmology and Vision Science (supplement), 38, S453.
Stoner GR, Albright TD, Ramachandran VS. (1990). Transparency and coherence in human motion perception. *Nature*. **344**(6262): 153-155.
Tommasi M. (1999) A ratio model of perceptual transparency. *Percept Mot Skills*. **89**(3 Pt 1): 891-7.
Watanabe, T. & Cavanagh, P. (1992). The role of transparency in perceptual grouping and pattern recognition. *Perception*, **21**,131-139.
Watanabe, T. & Cavanagh, P. (1993). Transparent surfaces defined by implicit X junctions. *Vision Research*, **33**: 2339-2346.
Watanabe, T. & Cavanagh, P. (1993). Surface decomposition accompanying the perception of transparency. *Spatial Vision*, **7**, 95-111.

Transparent Materials

Metelli's episcotister

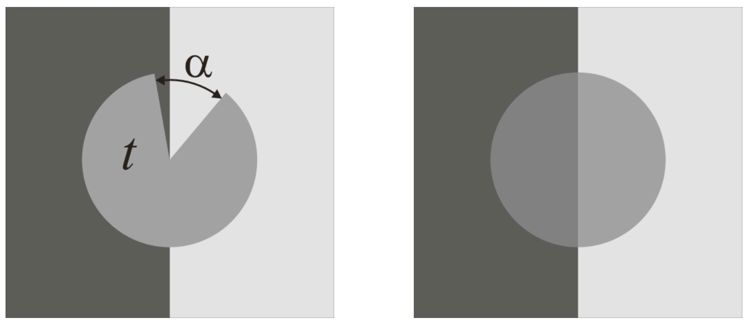

37

Real transparent objects

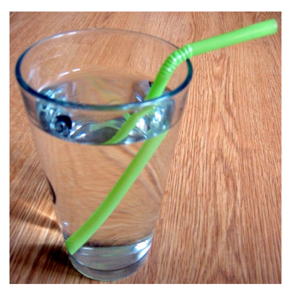

38

Real transparent objects

- … are not ideal infinitesimal films
- … obey Fresnel's equations:
  - Specular reflections
  - Refraction
- … can appear vividly transparent without containing the image cues traditionally believed to be important for the perception of "transparency".

39

Real transparent objects

- **Questions:**
- What image cues do we use to tell that something is transparent?
- How do we estimate and represent the refractive index of transparent bodies?

40

Refractive Index

- Possibly the most important property that distinguishes real chunks of transparent stuff from Metelli-type materials
- Snell's Law:

$$\frac{\sin(\theta_i)}{\sin(\theta_t)} = \frac{n_2}{n_1}$$

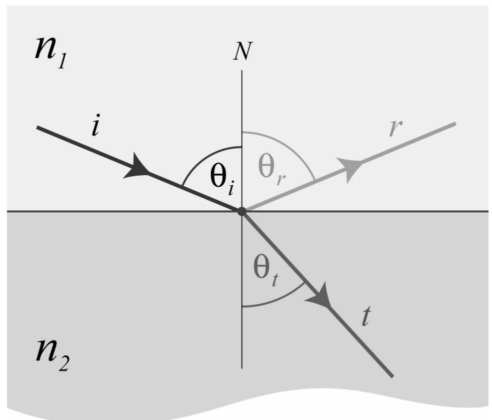

Refractive Index

- Varying the refractive index can lead to the distinct impression that the object is made out a different material

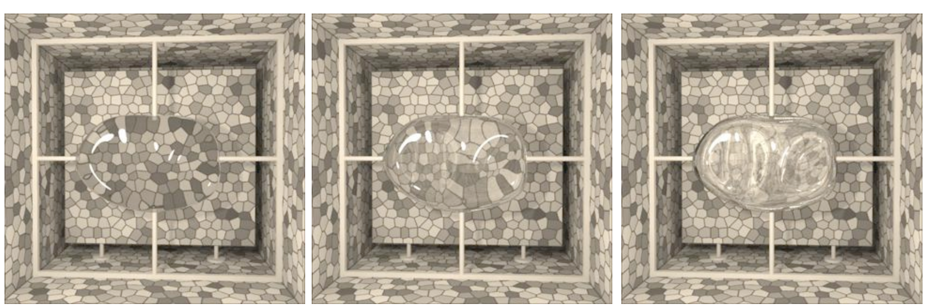

1.2 1.5 2.3

refractive index →

Refractive Index

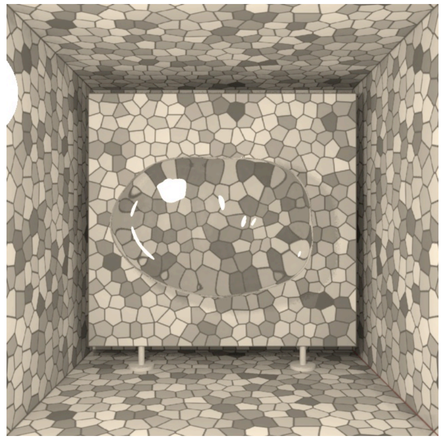

43

Refractive Index

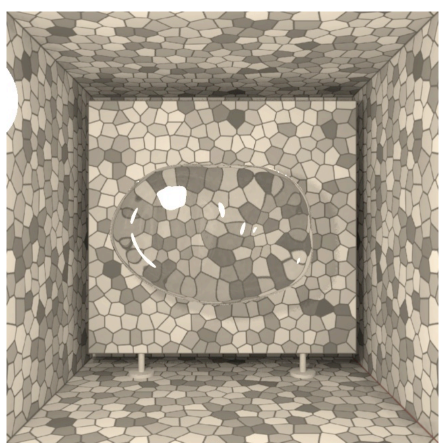

44

Refractive Index

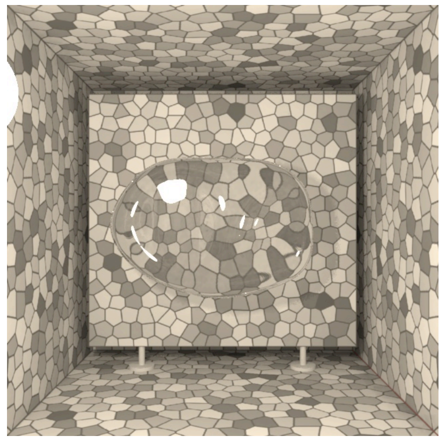

45

Refractive Index

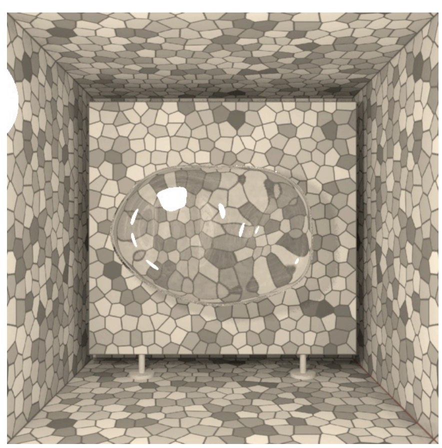

46

Refractive Index

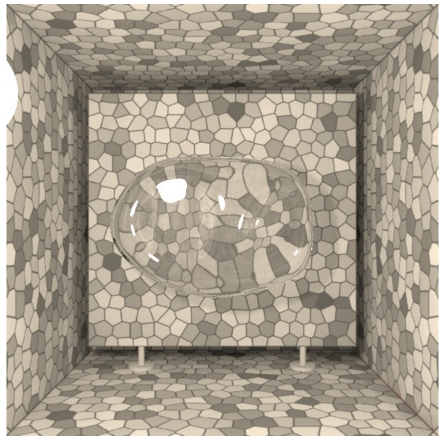

47

Refractive Index

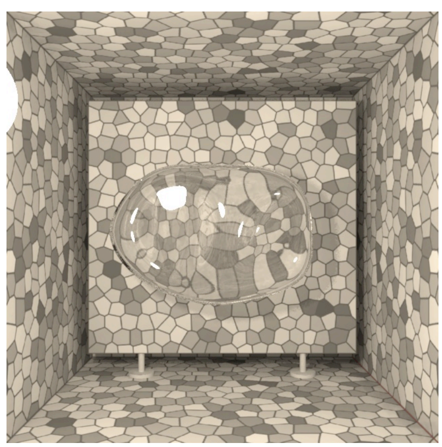

48

Refractive Index

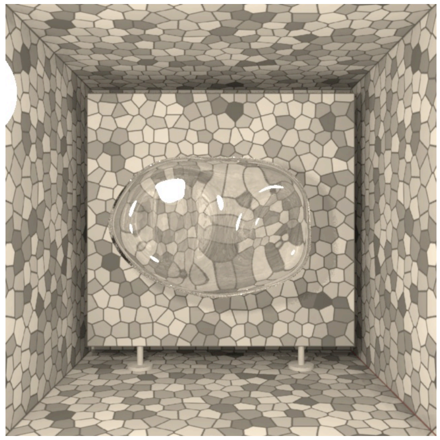

49

Refractive Index

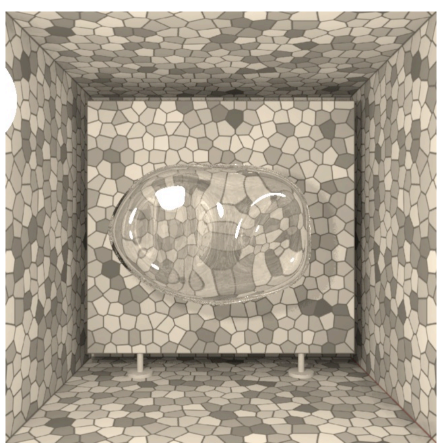

50

Refractive Index

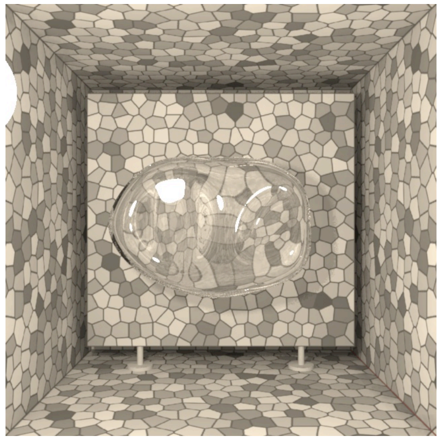

51

Refractive Index

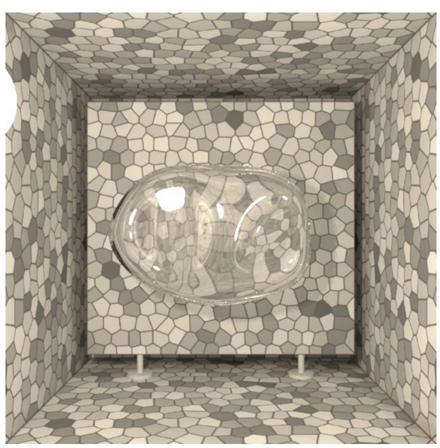

52

Refraction and image distortion

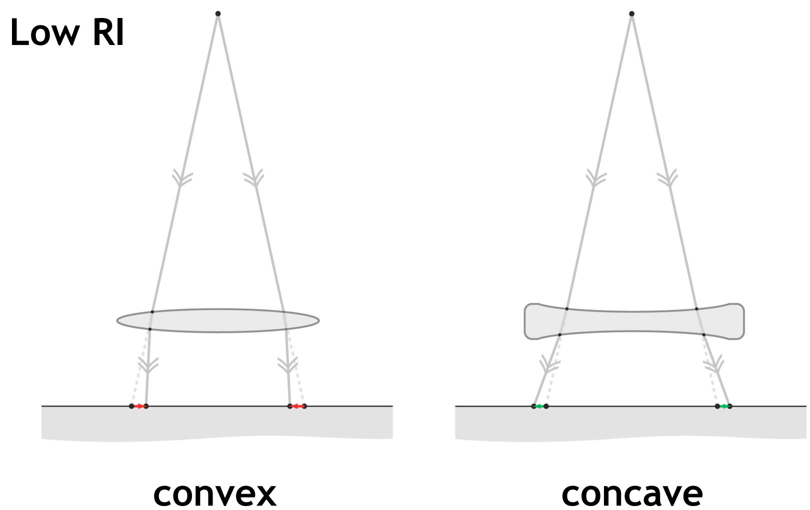

Refraction and image distortion

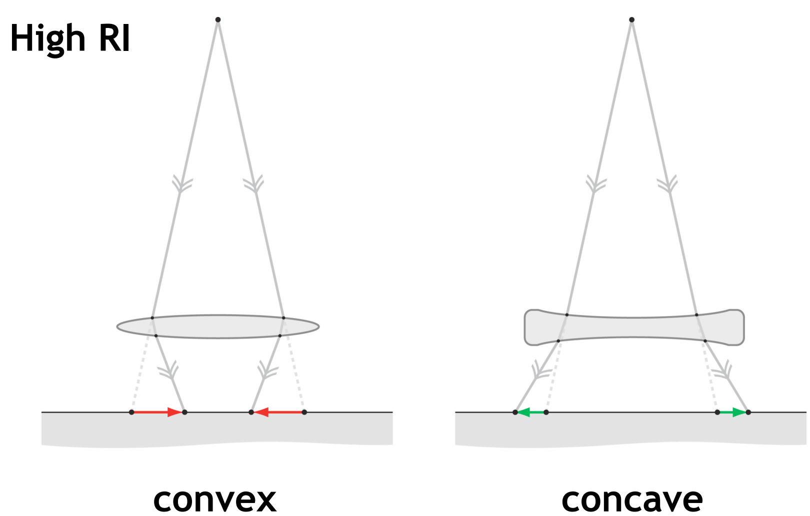

Displacement Field

- The perturbation of the texture caused by refraction can be captured by the **displacement field.**

$$\vec{d} = \vec{p}_{refracted} - \vec{p}_{actual}$$

- Measures the way the transparent object displaces the position of refracted features in the image.

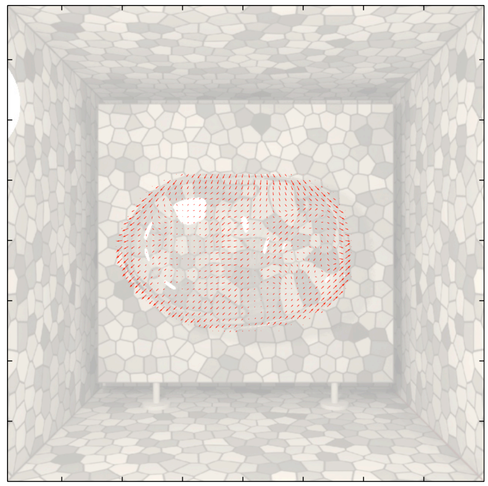

55

Displacement Field

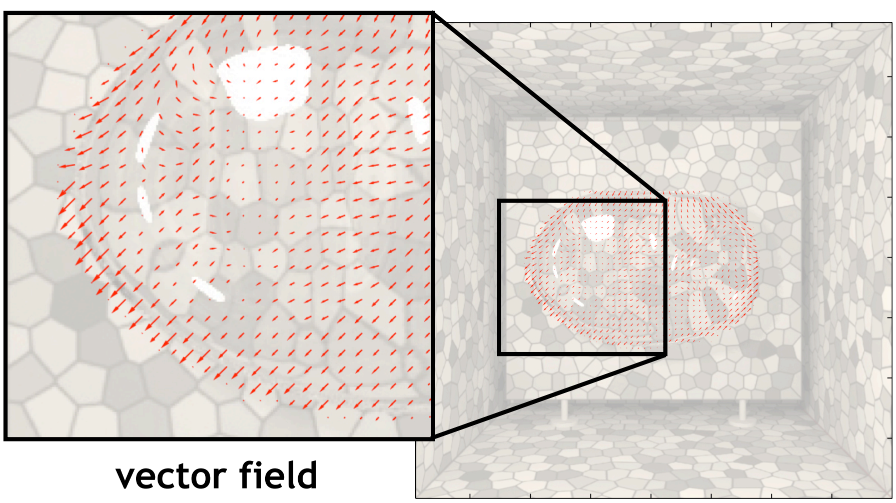

vector field

56

Distortion Fields

- **But,** to compute the displacement field the visual system would need to know the positions of non-displaced features on the backplane.
- Let us assume, instead, that the visual system can estimate the **relative distortions** of the texture (compression or expansion of texture)

$$\boldsymbol{D} = \operatorname{div}(\vec{\boldsymbol{d}})$$

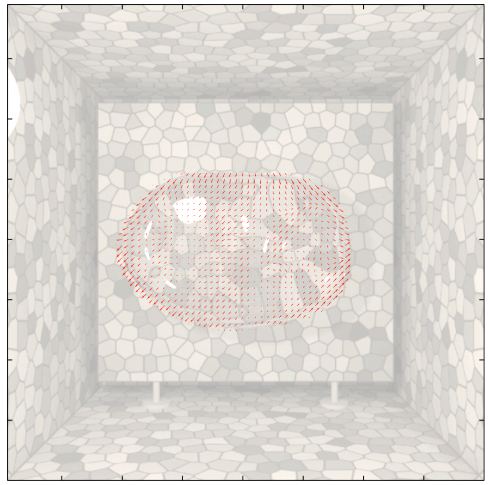

Distortion Fields

image

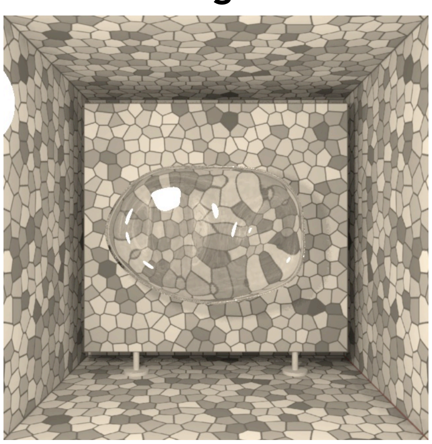

distortion field

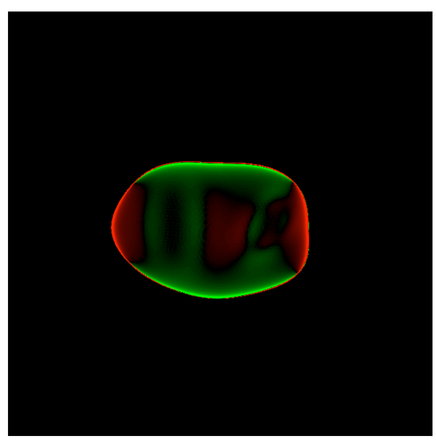

- Red = magnification
- Green = minification

Distance to backplane

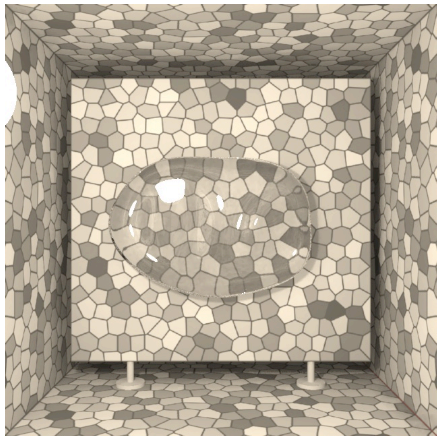

59

Distance to backplane

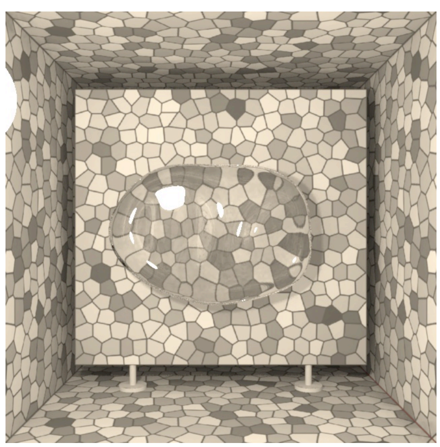

60

Distance to backplane

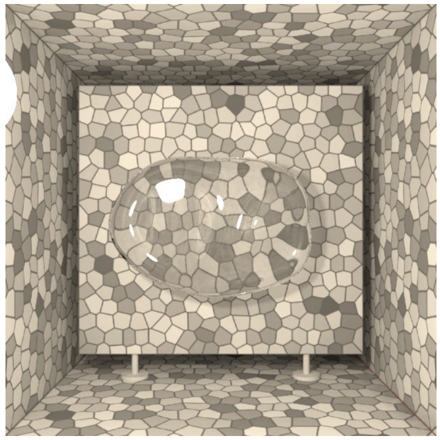

61

Distance to backplane

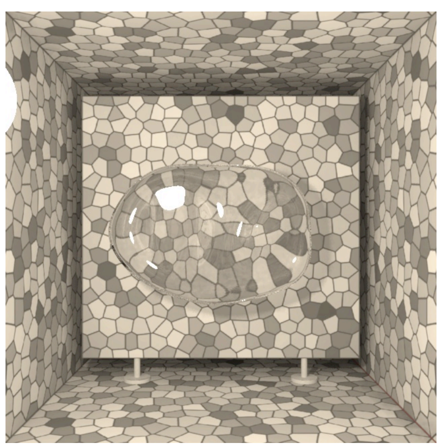

62

Distance to backplane

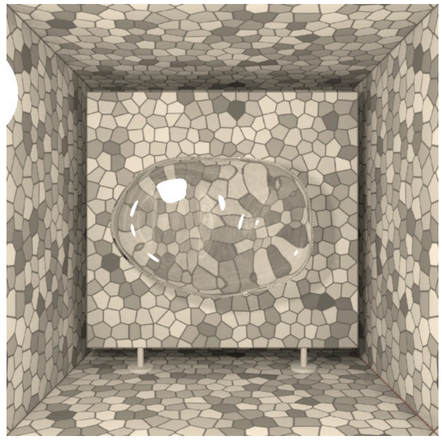

63

Distance to backplane

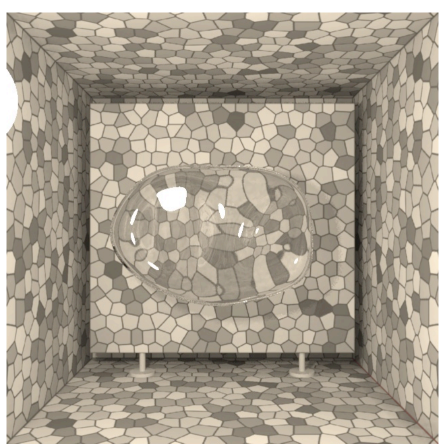

64

Distance to backplane

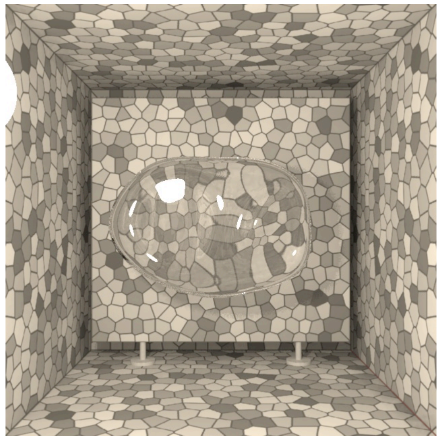

65

Distance to backplane

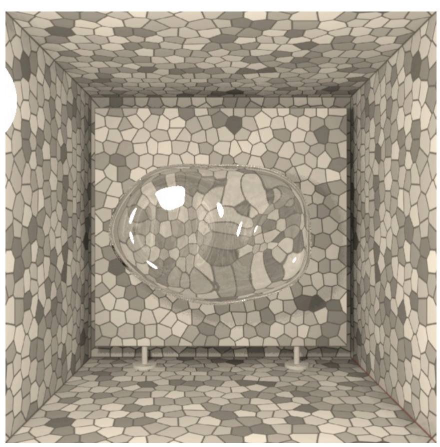

66

Distance to backplane

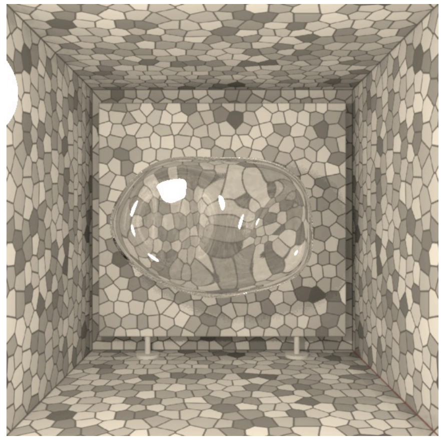

67

Distance to backplane

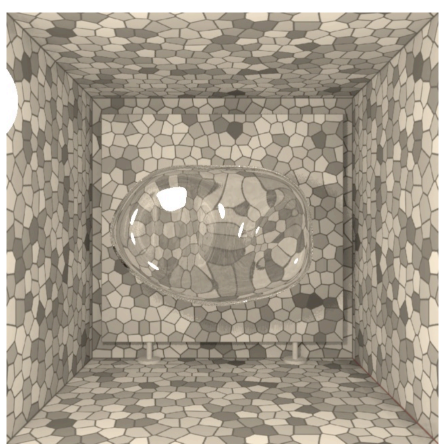

68

Distance to backplane

near

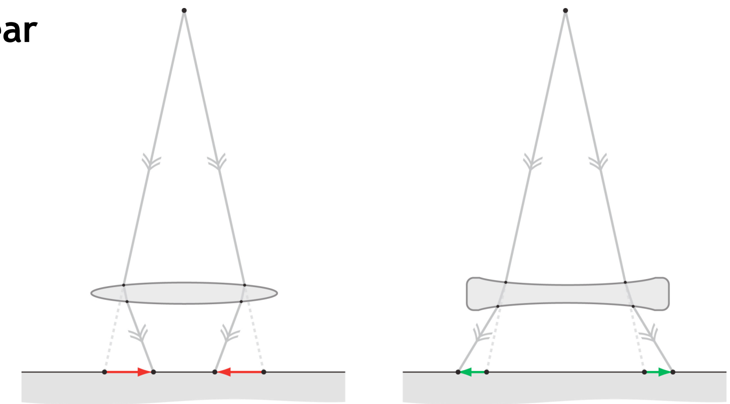

convex concave

Distance to backplane

far

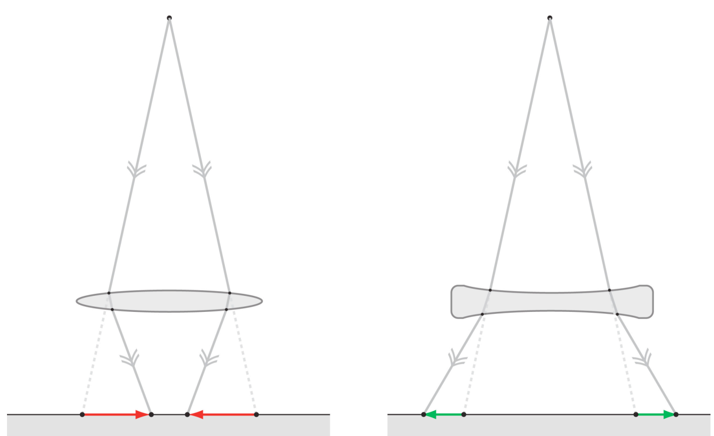

convex concave

Object thickness

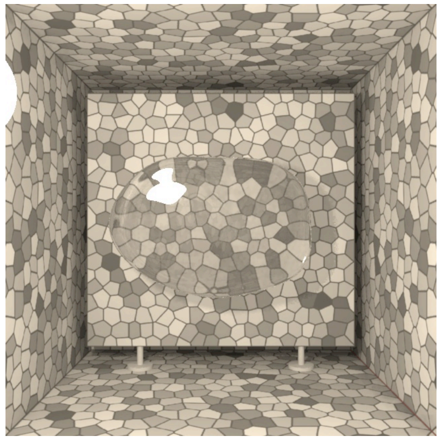

Object thickness

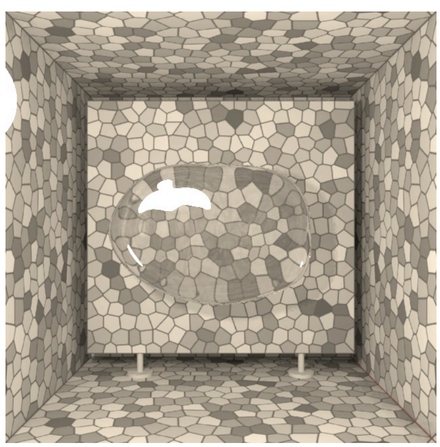

Object thickness

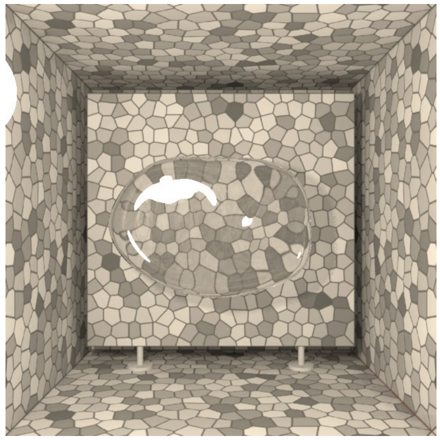

Object thickness

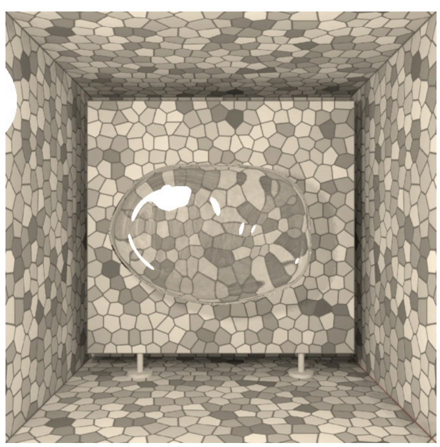

Object thickness

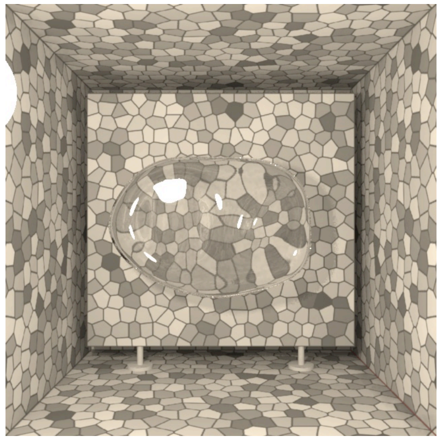

75

Object thickness

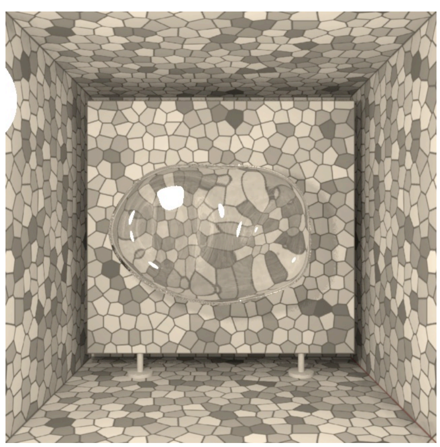

76

Object thickness

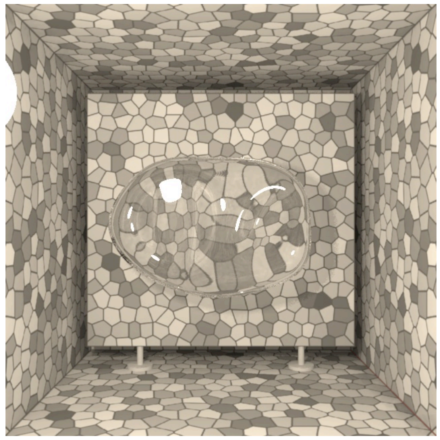

Object thickness

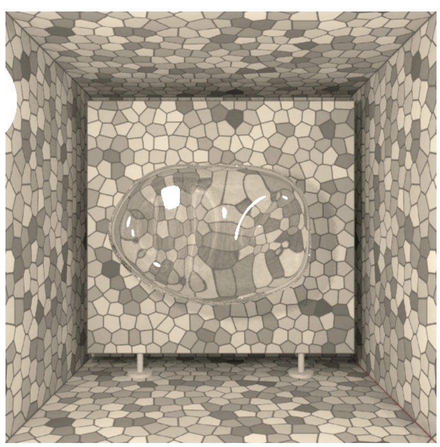

Object thickness

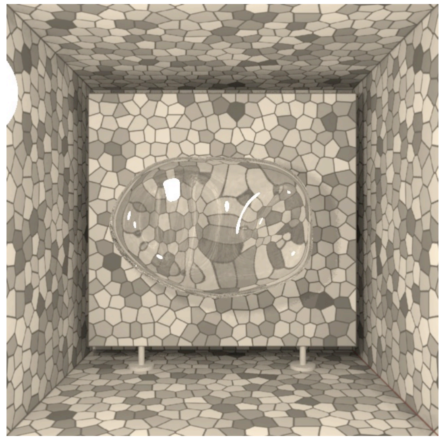

Object thickness

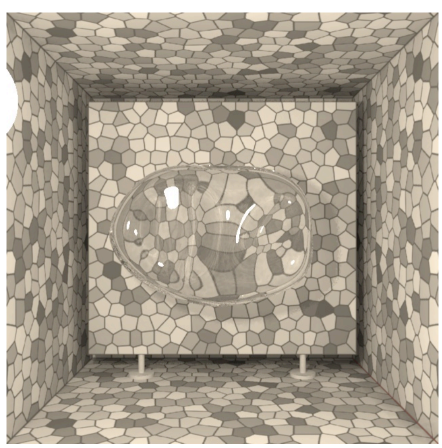

Asymmetric Matching task

- The distortions in the image depend not only on the RI but also on:
  - Geometry of object (curvatures, thickness)
  - Distance between object and background
  - Distance between viewer and object
- So, if the observers base their judgments of RI primarily on the pattern of distortions (rather than correctly estimating the physics), then they should show systematic errors in their estimates when these irrelevant scene factors vary.

81

Asymmetric Matching task

test

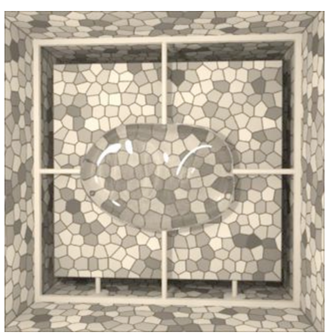

probe

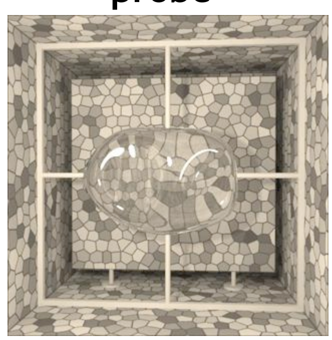

- Subject adjusts RI of standard probe stimulus to match the appearance of the other stimulus.
- One scene variable (thickness, distance to back-plane) is clamped at a different value for **test** stimulus.
- Measures ability of observer to 'ignore' or 'discount' the differences that are caused by irrelevant scene variables

82

Asymmetric Matching task

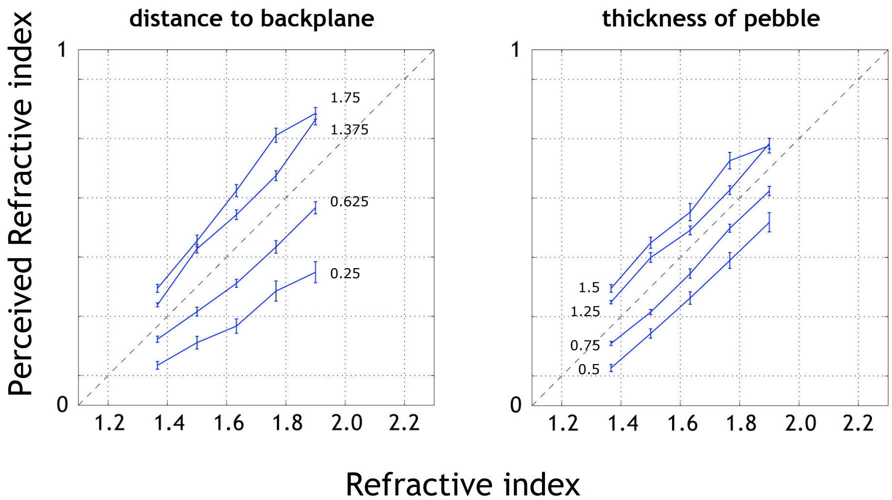

83

84

Observations

- Observers don’t appear to estimate the physics correctly.
- Instead, they probably use some **heuristic image measurements** (e.g. distortion fields) that are affected by refractive index, but also by other scene factors.
- This can lead to illusions (mis-perceptions)

Outline

- Visual estimation of surface reflectance properties: *gloss*
- Perception of materials that transmit light
  - *Refraction*
  - *Sub-surface scattering*
- Exploiting the assumptions of the visual system to edit the material appearance of objects in photographs

Image based material editing

input

output

- **Goal**: given single (HDR) photo as input, change appearance of object to completely different material
- Physically accurate solution would require fully reconstructing illumination and 3D geometry.
  - Beyond state-of-the-art computer vision capabilities

87

Image based material editing

input

output

- **Alternative**: “Perceptually accurate” solution
- Series of simple, highly approximate manipulations, each of which is provably incorrect, but whose *ensemble* effect is visually compelling.

88

Processing Pipeline

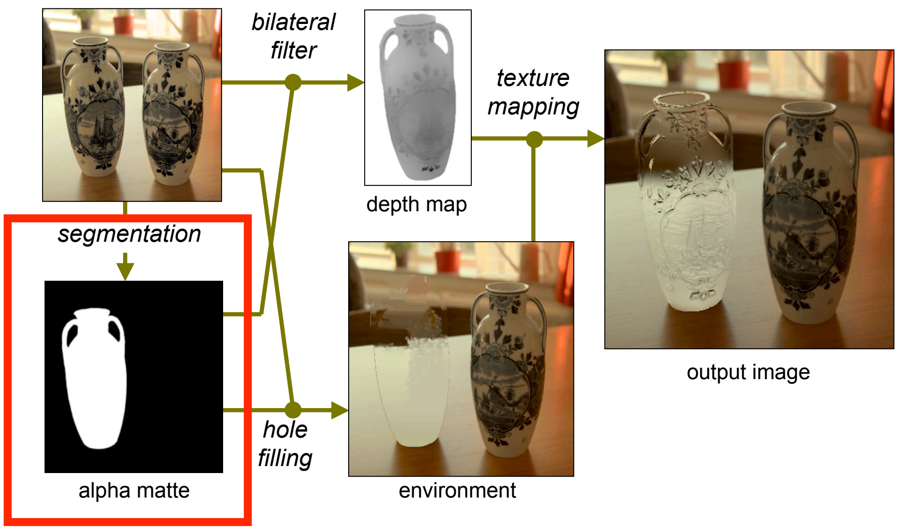

89

Alpha Mattes

- By restricting our image modifications to the area within the boundary of the object, we can create the illusion of a transformed material.

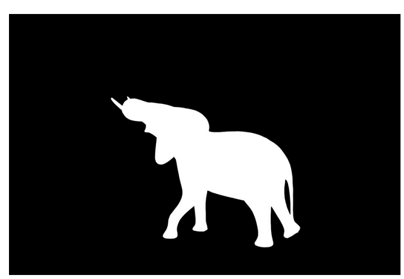

- **Assumption**: addition of new non-local effects (e.g. additional reflexes, caustics, etc.) is not crucial.
  - Approximate shadows and reflections are already in place in original scene

90

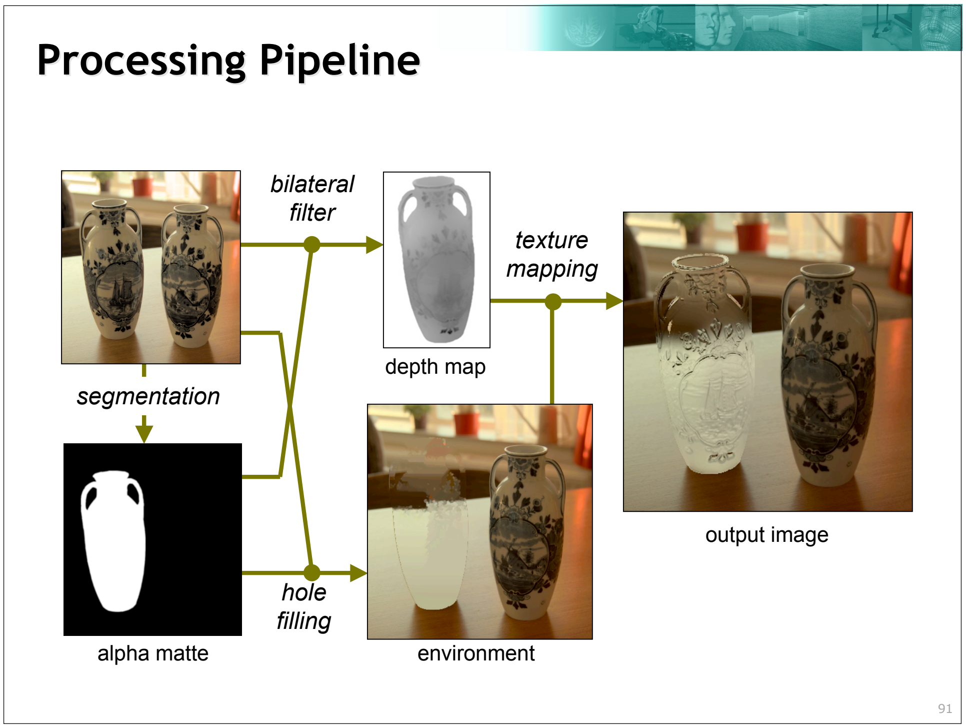
Processing Pipeline
bilateral filter
depth map
texture mapping
segmentation
output image
alpha matte
hole filling
environment
91

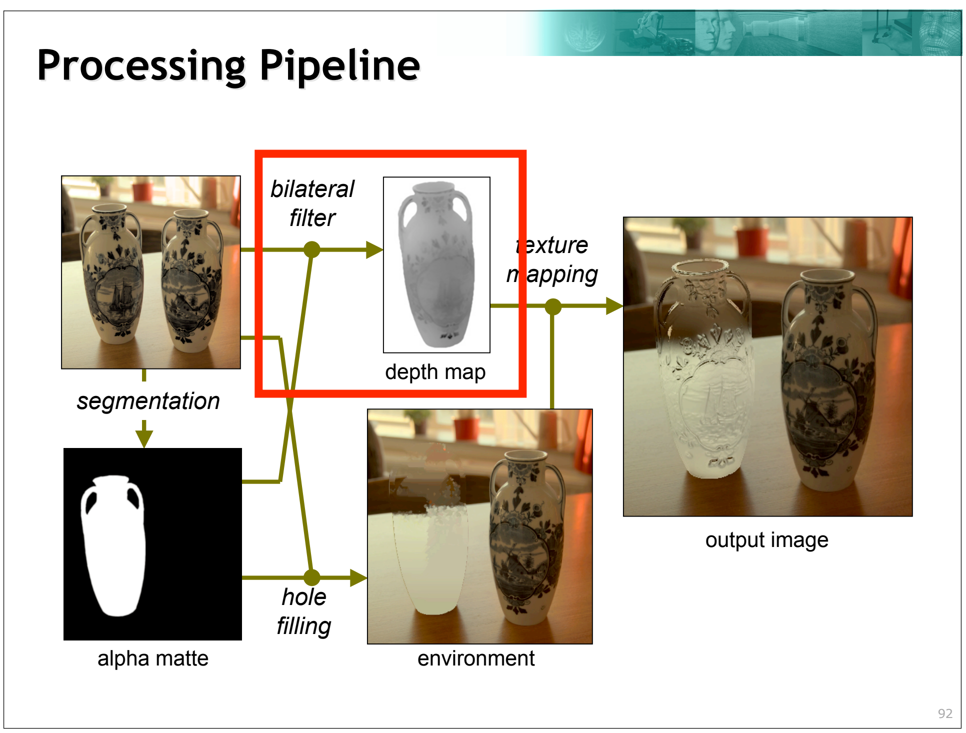
Processing Pipeline
bilateral filter
depth map
texture mapping
segmentation
output image
alpha matte
hole filling
environment
92

How **not** to do shape-from-shading

Try using the state-of-the-art algorithms and you will generally be disappointed!

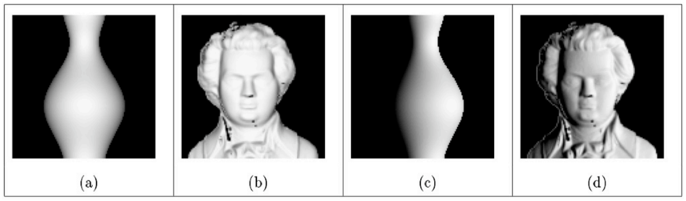

input

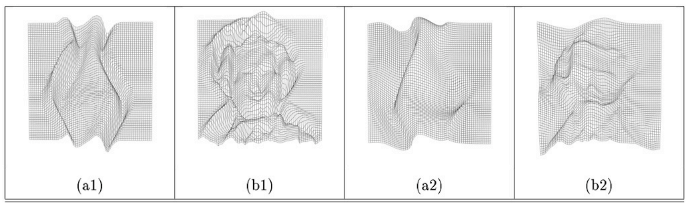

reconstruction

93

How **not** to do shape-from-shading

Try using the state-of-the-art algorithms and you will generally be disappointed!

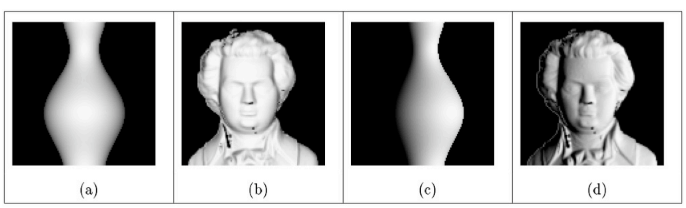

input

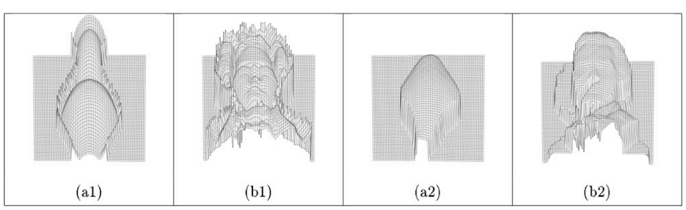

reconstruction

94

How **not** to do shape-from-shading

Try using the state-of-the-art algorithms and you will generally be disappointed!

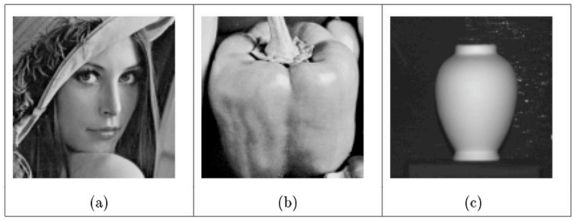

input

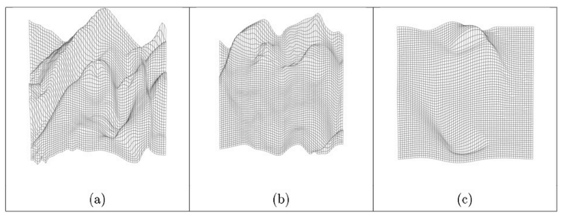

reconstruction

What is the alternative ?

We use a simple but suprisingly effective heuristic:

Dark is Deep

In other words …

$$z(x, y) = L(x, y)$$

Complicated shape from shading algorithm

input **initial depth estimate**

Bilateral Filter

- The ‘recovered depths’ are conditioned using a **bilateral filter** (Tomasi & Manduchi, 1998; Durand & Dorsey, 2002).
- Simple non-linear **edge-preserving filter** with kernels in space and intensity domains.

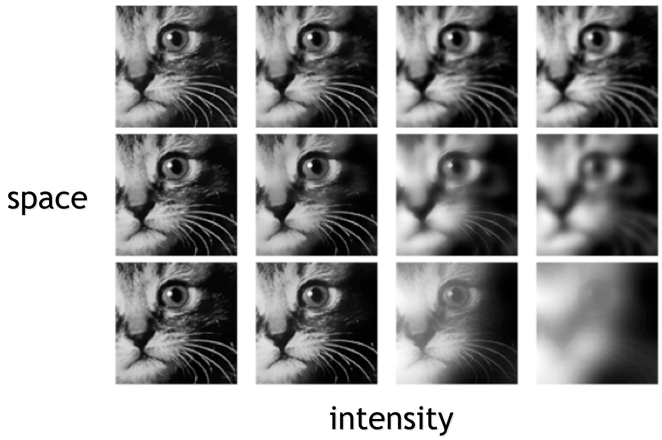

97

Bilateral Filter:
3 main functions

- 1. De-noising depth-map
  - **Intuition**: depths are generally smoother than intensities in the real world.
- 2. Selectively enhance or remove textures for embossed effect

98

Bilateral Filter:
3 main functions

- 3. **Shape-from-silhouette**, like level-sets shape 'inflation' (e.g. Williams, 1998)
  - **Intuition**: values outside object are set to zero, so blurring across boundary makes recovered depths smooth and convex.

99

Forgiving case

- Diffuse surface reflectance leads to clear shading pattern
- Silhouette provides good constraints

original

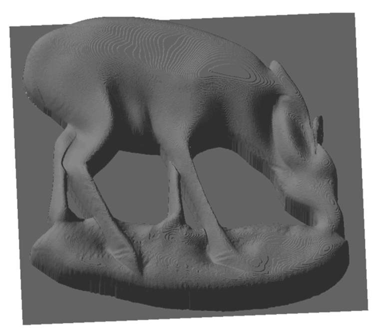

reconstructed depths

100

Difficult case

- Strong highlights create large spurious depth peaks
- Silhouette is relatively uninformative

original

reconstructed depths

101

Light from the side

- Shadows and intensity gradient leads to substantial distortions of the face

original

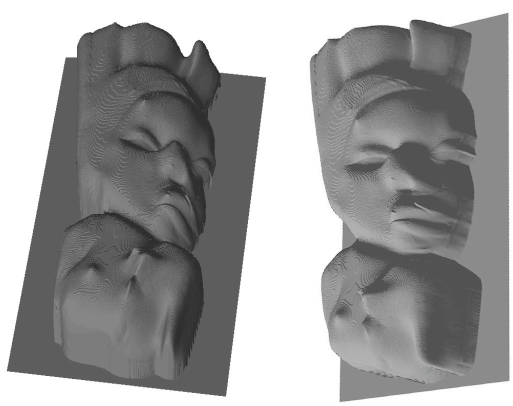

reconstructed depths

102

Importance of viewpoint

- Substantial errors in depth reconstruction are not visible in transformed image

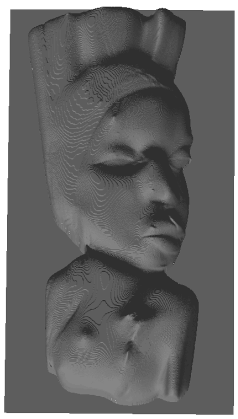

correct viewpoint **transformed image**

103

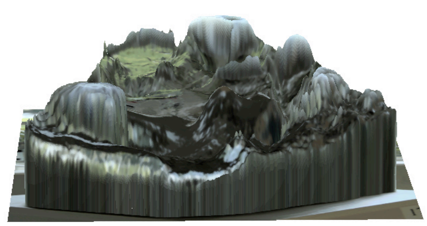

104

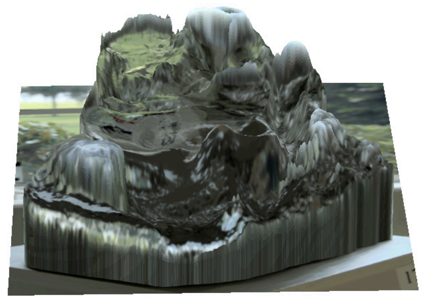

105

106

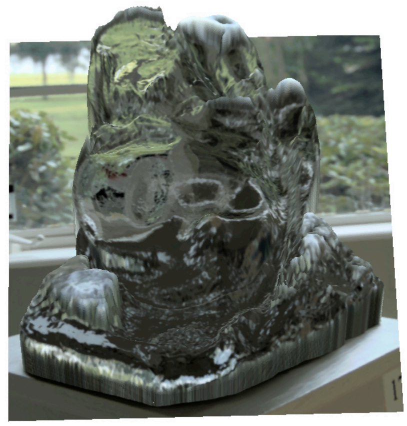

109

Why does it work ?

- Generic viewpoint assumption (Koenderink & van Doorn, 1979; Binford, 1981; Freeman, 1994)

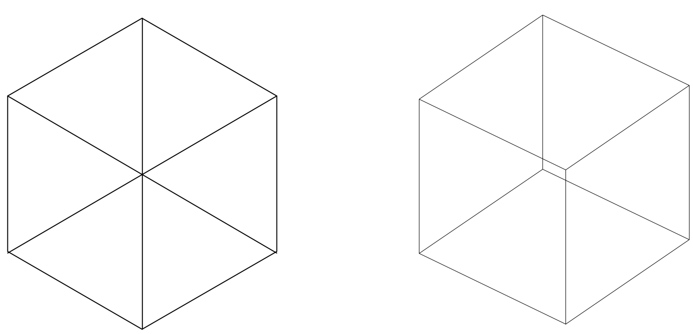

110

Why does it work ?

- Masking effect of patterns that are subsequently placed on surface (e.g. highlights).

Processing Pipeline

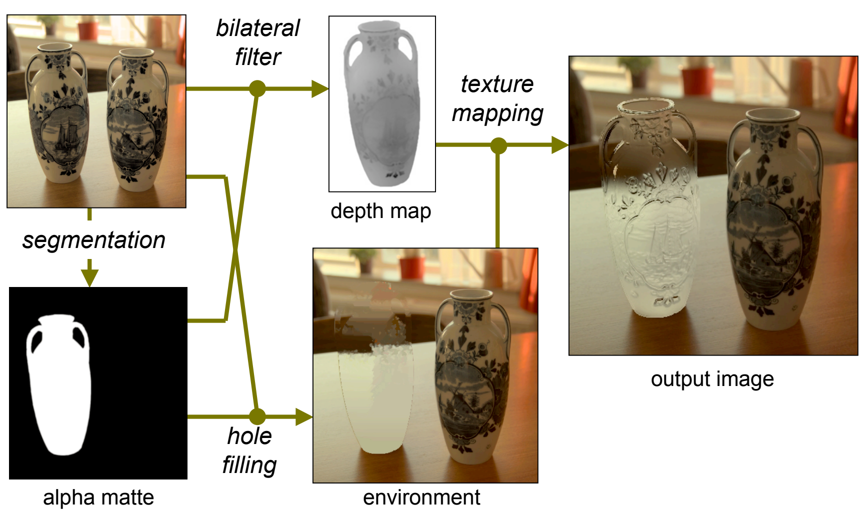

Processing Pipeline

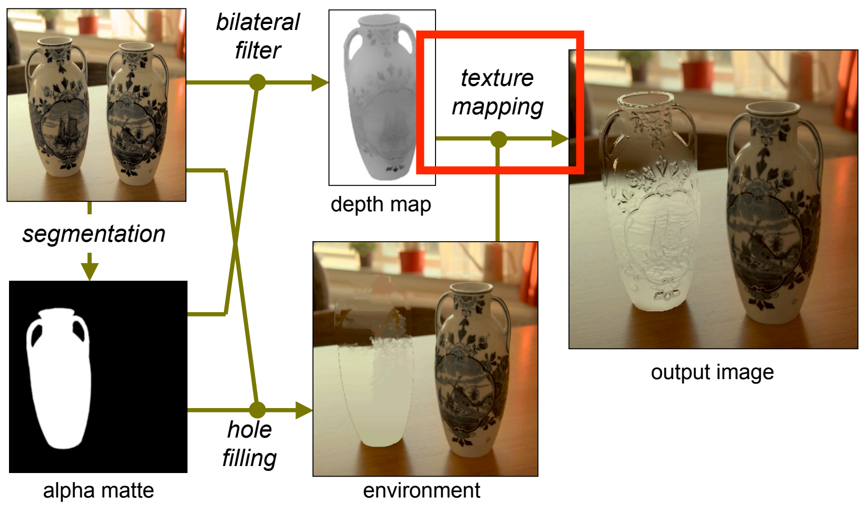

113

Re-texturing the object

- Use recovered surface normals as indices into a texture map
- **The most important trick:** blend original intensities back into image, for correct shading and highlights

114

Re-texturing the object

- Use recovered surface normals as indices into a texture map
- **The most important trick**: blend original intensities back into image, for correct shading and highlights

115

Re-texturing the object

- **The most important trick**: blend original intensities back into image, for correct shading and highlights
- **TextureShop** (Fang & Hart, 2004) tacitly uses the same trick:

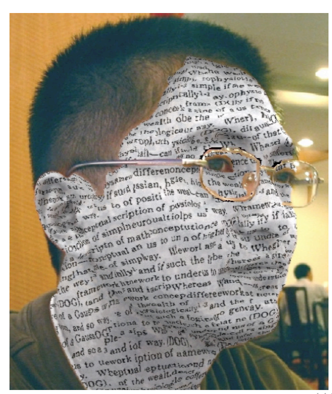

116

Other material transformations

- To create the illusion of transparent or translucent materials, we map **a modified version of the background** image onto the surface.
- To apply arbitrary BRDFs, we use
  - the recovered surface normals, and
  - an approximate reconstruction of the environment

  to evaluate the BRDF.

117

Processing Pipeline

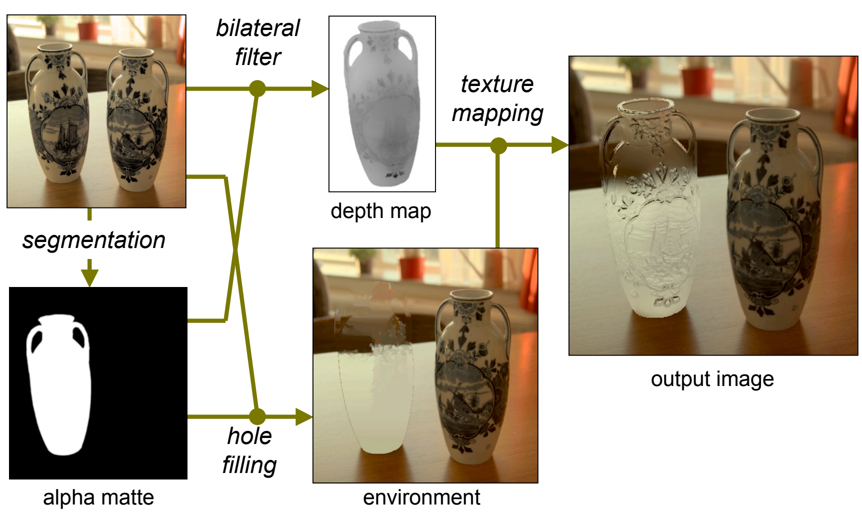

118

Processing Pipeline

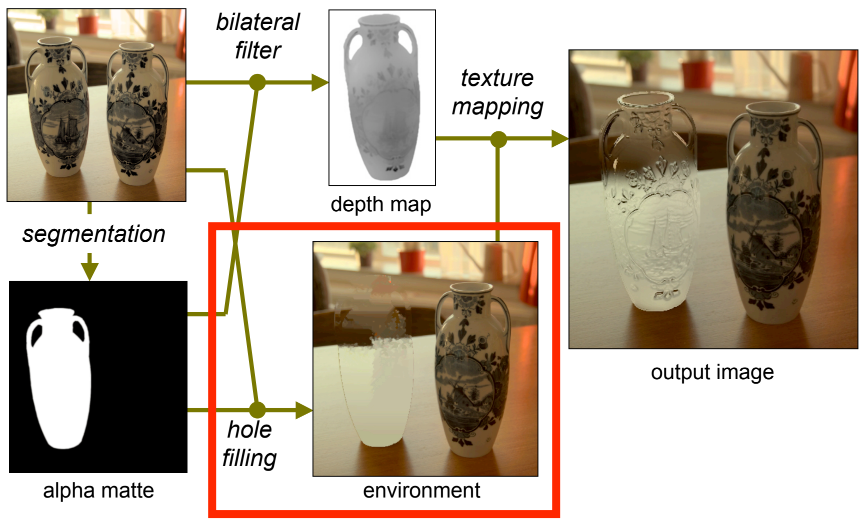

119

Hole filling the easy way

- A number of sophisticated algorithms exist for object removal (e.g. Bertalmio et al. 2000; Drori et al. 2003; Sun et al. 2005)
- Crude but fast alternative: cut and paste!

120

Hole filling the easy way

- A number of sophisticated algorithms exist for object removal (e.g. Bertalmio et al. 2000; Drori et al. 2003; Sun et al. 2005)
- Crude but fast alternative: cut and paste!

121

Hole filling the easy way

- A number of sophisticated algorithms exist for object removal (e.g. Bertalmio et al. 2000; Drori et al. 2003; Sun et al. 2005)
- Crude but fast alternative: cut and paste!

122

Fake transparency

- The human visual system appears does not recognize transparency by correctly using **inverse optics**.
- Instead, it seems to rely on the consistency of the **image statistics** and **patterns of distortion**.

123

Environment map

- Background image is used to generate full HDR light probe for image-based lighting

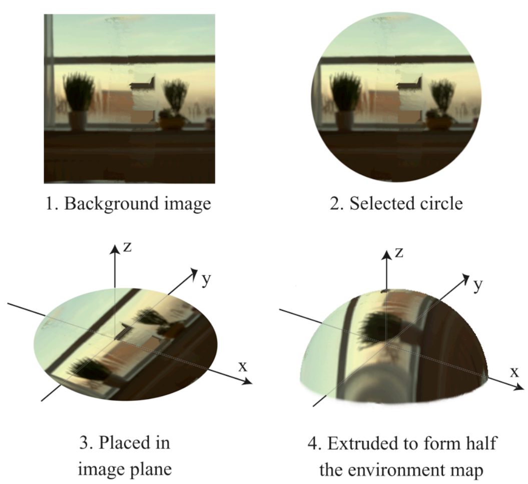

124

Arbitrary BRDFs

- Given surface normals and complete HDR light probe, we can evaluate empirical or parametric BRDFs, as in standard image based lighting (local effects only).
- We used Matusik's BRDFs.

original

blue metallic

nickel

125

Arbitrary BRDFs

- Given surface normals and complete HDR light probe, we can evaluate empirical or parametric BRDFs, as in standard image based lighting (local effects only).
- We used Matusik's BRDFs.

original

nickel

126

General Principles

- **Complementary heuristics.** Normally, errors accumulate as one approximation is added to another. The key to our approach is choosing heuristics such that the errors of one approximation visually **compensate** for the errors of another.
- Exploit visual tolerance for ambiguities to achieve solutions that are **perceptually acceptable** even though they are **physically wrong**.

127

Using shape to infer material properties

128

Using shape to infer material properties

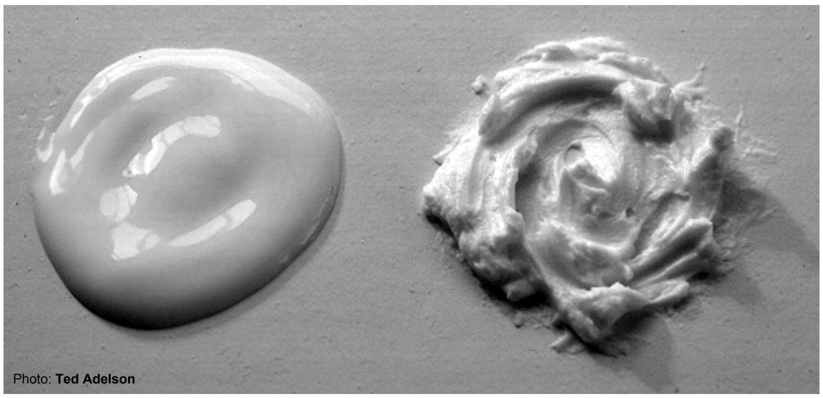

129

Final thoughts

- 'Stuff' adds emotional meaning to our visual environment and can even play a role in our biological survival

130

Final thoughts

- **Paradox:** the brain can be exquisitely sensitive to subtle differences in material properties, so to do good renderings you need to get them right

No subsurface scattering **With subsurface scattering**

Nvidia advanced skin rendering demo

Final thoughts

- **On the other hand:** The brain makes assumptions, so you can sometimes get the physics hopelessly wrong as long as you get the statistics roughly right.

Thank You

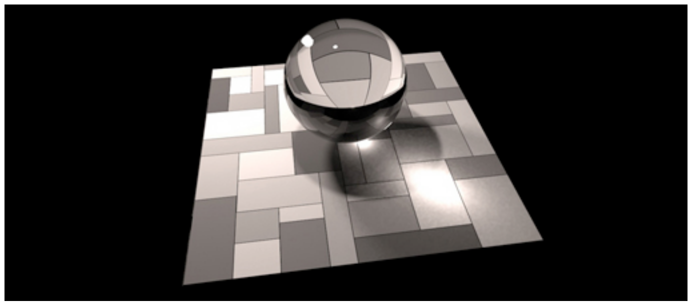

Co-authors:

Ted Adelson, Ron Dror, Frank Jäkel, Larry Maloney, Erum Kahn, Erik Reinhard, Heinrich Bülthoff

Funding: DFG FL 624/1-1

Some renderings generated using Henrik Wann Jensen's DALI

133