Comp/Phys/APSc 715

Object recognition, Surface shape, Texture, Depth cues, Stereo, Combinations

Preview Videos

- Vis 2006: sreng.avi
  - Proximity and collision glyphs
- Digital ArtForms interface
  - Watch the video on the left side
- Video: Rendering text labels on visualizations

Administrative

- Questions about what you're turning in tonight?
- Questions about what you're presenting next week?

Discussion!

- When do we let the user choose?
- When don't we let the user choose?

Surprised?

- What information from Ware chapter 7 surprised you?

What is an Object?

- Any identifiable, separate, and distinct part of the visual world
- A visual object cognitively groups visual attributes
- Lesson: Representing data values as visual features and grouping them into visual objects can be a powerful tool for organizing related data.

How are Objects Recognized?

- Image based?
  - The mind as a huge movie reel
- Structure based?
  - Breaking object into 3D parts

Image-Based Object Recognition

- Have I seen this before? (2560 images, 90%)
- Rapid serial presentation, with attentional blink
- Memory may be indexed by images (recognition), which then fire other related memories (recall).
- Selective Priming: Visual, not verbal
- Canonical views in monkey brains.

Uses of Image Display

- Icons in user interface can cause recognition and then recall of function.
- Priming can be helpful when the user is searching for a pattern or image.
- It may be faster to present images in a "burst" at up to 10 frames per second – like flipping through a book – rather than side-by-side thumbnails.

Hint for Image Display

- Biederman and Cooper (1992) showed that the optimal size for recognizing visual objects is 4-6 degrees of visual angle.
- Mona Lisa from afar
- Gremlin and Raven nearby

Structure-Based Object Recognition

- We recognize new orientations of novel objects

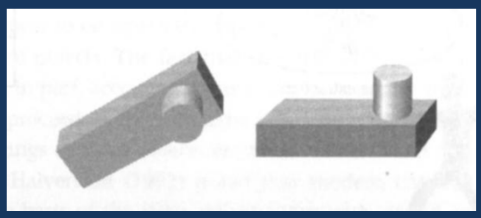

Geon Theory

- The whole is a sum of a set of basic primitive geometrical elements
- The way they are connected is also encoded
- Geon Man!

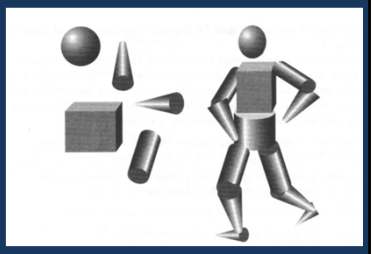

Silhouettes

- Especially important in object perception
  - Cave drawings based on this
  - Modern children draw this way
  - A clear diagram can be more effective than a photo
- Canonical silhouettes
  - Sideways Man!

- Concave sections break object into parts

Object-Based display of Data

- Grouping six variables into one object makes them more comprehensible

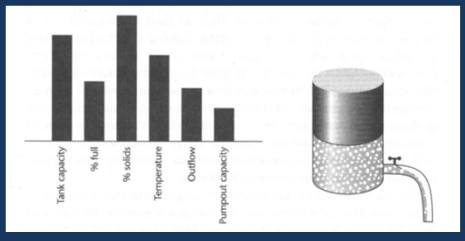

Object Display Characteristics

- Benefits
  - Can consolidate multiple related data sets into one object
  - Can map onto familiar objects
- Issues
  - Requires specific design for each application
  - Requires a meaningful metaphor

Surface Shape Perception

- Used for: digital elevation maps
  - Ocean floors
  - Molecular-scale surfaces
  - Mathematical functions
  - Other 2-dimensional scalar fields
- Important perceptual characteristics
  - Surface shading models and contours
  - Surface texture
  - (Stereo and Motion described elsewhere)

Surface Shading

- Basic components (from before)
  - Lambertian shading: diffuse reflection
  - Specular shading: glossy highlights
  - Ambient: Hack to simulate radiosity
  - Low-contrast texture with linear elements
  - Cast shadows: On itself or another object

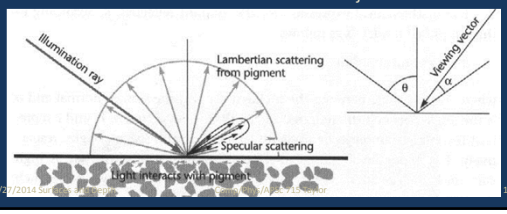

Surface Shading

- Basic components (from before)
  - Lambertian shading: diffuse reflection
  - Specular shading: glossy highlights
  - Ambient: Hack to simulate radiosity
  - Low-contrast texture with linear elements
  - Cast shadows: On itself or another object
- Goal is revealing shape, not realism
  - Visual system assumes a single light source from above
  - Multiple light sources may be confusing
  - Cast shadows inform relative positions (more later)

Expect Light from Above

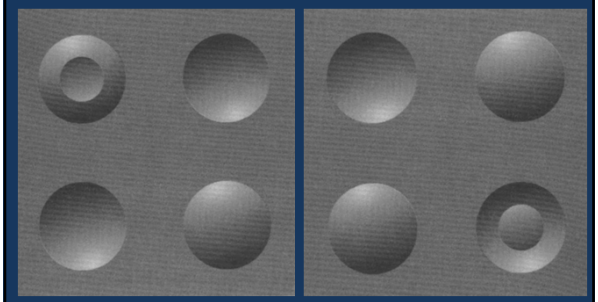

Shading and Contours

- Shape from shading is inherently ambiguous
  - Assumes a lighting direction, for one thing
- Different contours with same shading → different perceived shapes

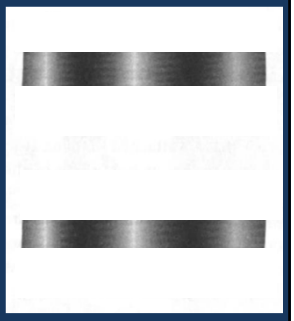

Shading and Contours

- Shape from shading is inherently ambiguous
  - Assumes a lighting direction, for one thing
- Different contours with same shading → different perceived shapes

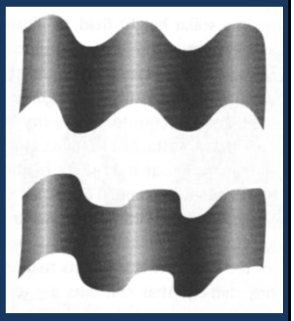

Shading and Internal Contours

- Internal contours also override shading information (apparent light direction shifts)

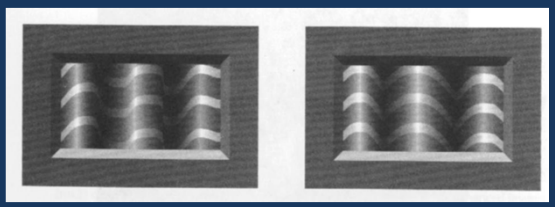

Shading and Internal Contours

- Equal-spaced lines enable gradient estimation

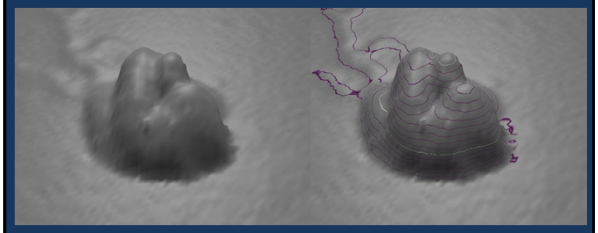

Surface Texture

- Gibson claims that a non-textured surface is just a patch of light
- Shape information comes from texture gradient

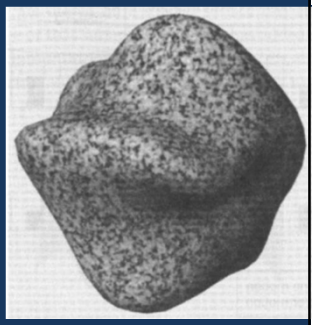

Surface Texture

- Kim, 2003
- a). 1st & 2nd PD
- b). 1st PD
- c). LIC on 1st PD
- d). No texture

UNC VPAW System

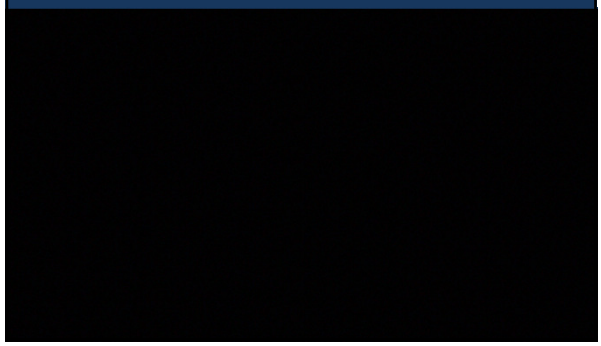

Surface Texture and Stereo

- Untextured polygons produce no *internal* stereoscopic correspondences
- Stereo correspondences reveal surface shape

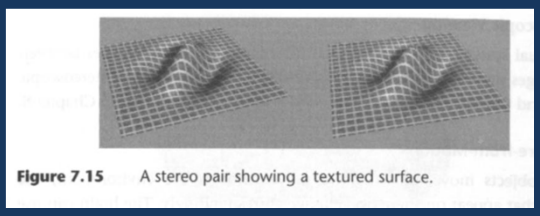

Figure 7.15 A stereo pair showing a textured surface.

Surface Display Guidelines

- One light at infinity, from above and to one side
- Lambertian + moderate specular lighting
  - Specular lighting is important to reveal details
  - Specular lighting is local, so enable control over light
- Surfaces should be textured with low-contrast textures that have linear features
- Cast shadows if they don't interfere: soft edges on the shadows
- Rotation and stereo (and head tracking) helpful

Self-Occlusion of Light

- Tarini, Cignoni, Montani, IEEE TVCG 12(5), 2006

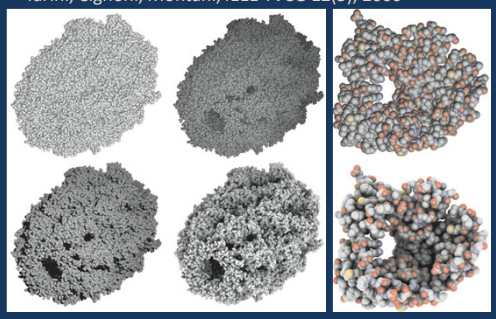

Image-based Relighting

- Akers et al., IEEE Vis. 2003

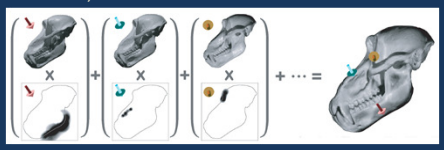

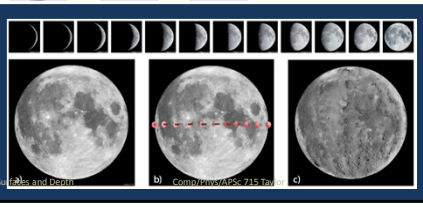

Living in a 3D World

- Recent hardware advances make it possible to view things in 3D easily and cheaply
- Early Powerpoint users taught us that
  - Just because you **can** do something
  - doesn't mean that you should!
- It can be helpful when used appropriately

Depth Cues

- Monocular cues
  - Seen with one eye
  - Static
    - Picture not moving
    - Like a photo on the wall
  - Dynamic
    - Picture is moving
    - Like on TV or at the movies
- Binocular cues
  - Toy Story in 3D, Virtual Reality
- Artificial cues
  - Not like in the real world, but they work

Monocular Static Cues

- Occlusion: King of the depth cues!
- Linear Perspective
  - Size Gradient
- Texture Gradient
- Depth of Focus
- Cast Shadows
- Shape-from-Shading

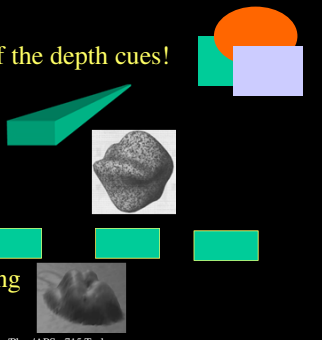

Occlusion

- Strongest cue, but binary
- Don't mess with occlusion →
- Occluding object looks closer

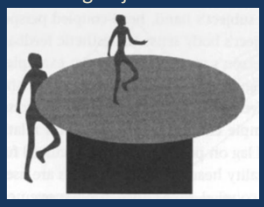

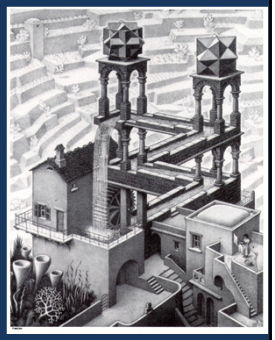

Linear Perspective

- 3rd person view
- 1st person view

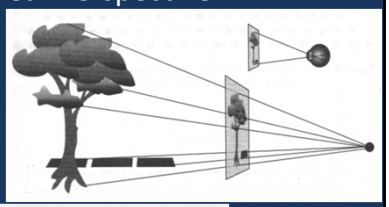

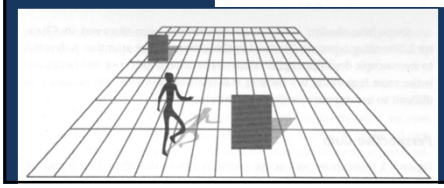

Linear Perspective Characteristics

- Parallel lines converge to a single point
- Objects that are further away appear smaller
- Size constancy
  - Objects of known size (e.g. people) effectively scale the whole scene
- Can perceive objects in pictures even though perspective is incorrect for where we view from
  - Can perceive both "picture-plane" size and "3D" size of objects shown in pictures
  - Visual system overrides some aspects of perspective
  - Perhaps built-in assumption of objects as rigid bodies causes this
- Hint: Simulate long-focal-length lenses for extreme off-axis viewing (less perspective effect)

Texture Gradient

- Textures with uniform statistics show shape by their distortion
  - May be uniform in projection
  - Better: uniform on the surface

Rainbow color map suboptimal

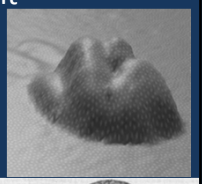

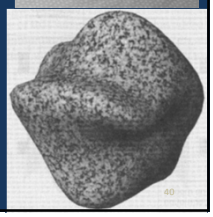

Depth of Focus

- Objects at same plane as fixated objects in focus
  - Objects in front or behind are out of focus
    - Objects behind → sharp boundary with fixated object
    - Objects in front → blurry boundary with fixated object
  - Separates foreground object from background

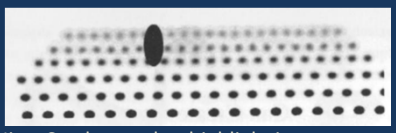

- Hint: Can be used to highlight important parts by blurring non-critical portions of the display

Cast Shadows

- Show relative heights of objects above plane
- Even used in game consoles, where polygon count is *the* critical resource

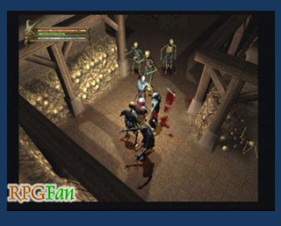

Cast Shadows

- Strong cue for relative height, especially in motion

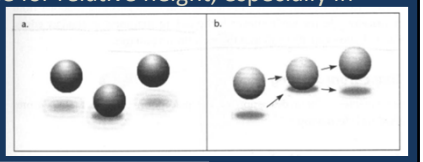

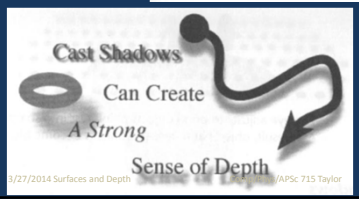

Cast Shadow Characteristics

- Potent at showing height above a plane
- Especially valuable in combination with motion
- In some cases can be stronger than texture, projection type, frames of reference, and motion

- Hints
  - Shadow shape does not have to be correct
  - Fuzzy-edged shadows lead to less confusion

Hovering Man?

Shape-From-Shading (Recap)

- Basic components
  - Lambertian shading: diffuse reflection
  - Specular shading: glossy highlights
  - Ambient: Hack to simulate radiosity
  - Low-contrast texture with linear elements
- Goal is revealing shape, not realism
  - Visual system assumes a single light source from above
  - Multiple light sources may be confusing

Monocular Dynamic Cues: Structure from Motion

- Motion Parallax
  - Sideways out a car window
  - Forwards out the car windshield
  - Head-Motion Parallax
- Kinetic Depth Effect
  - Objects in the environment moving or rotating

Motion Parallax

- Vehicle / linear motion

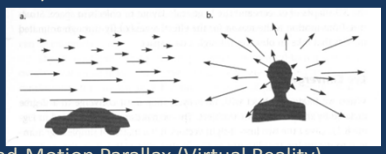

- Head-Motion Parallax (Virtual Reality)
  - Combination of directions coupled to head motion
  - Powerful effect if done with low latency
  - Especially powerful when combined with stereo

Kinetic Depth Effect

- Objects in the scene moving and/or rotating

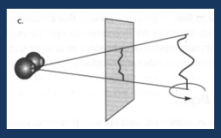

- Kinetic Depth vs. Perspective
  - Rotating trapezoidal window
  - Appears to swing back and forth
  - www.exploratorium.edu

Binocular Cues

- Eye convergence
  - Not good at absolute judgment

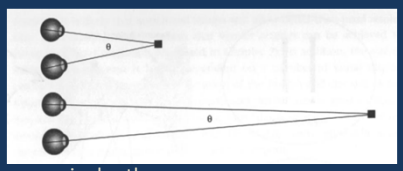

- Stereoscopic depth
  - …

Binocular Cues

- Stereoscopic depth
  - Disparity between images in each eye

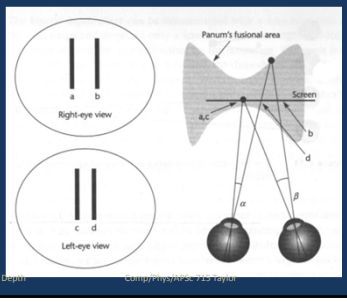

Nausea, Headache and Other Hazards of Stereoscopic Display

- Everything is in focus all the time
  - at a fixed distance that may not match fixated object distance
  - convergence and focus are out of alignment
- Objects cut off at the edge of the display
  - even if they are in front of it
- Stereo itself stops working after 30 meters
- Hint: Visual system is flexible about stereo cues
  - Enable user to adjust scaling and eye separation to suit

Integration of Cues

- All cues are useful (specular, lambertian, texture, stereo, motion parallax)
- Relative importance varies from person to person, and depends on the task
  - 2-4% of population is stereo blind
- Motion and Stereo reduced errors in combination with *any* of the others
- Lambertian shading with either stereo or motion was nearly the best for all subjects
- Others found that texture beat Lambertian or specular
- Stereo + head motion is much better than either alone

Artificial Spatial Cues

- Dropping a line to the ground plane
  - Directly shows height
  - A bit like shadows

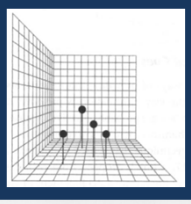

- Proximity Luminance Covariance
  - Fades into background
  - A bit like fog

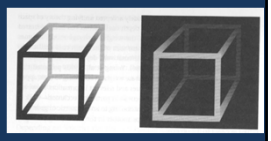

Depth Cue Relative Importance

James Lane

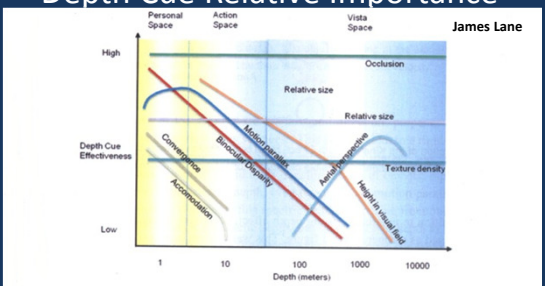

Figure 2.14 Effectiveness of depth cues at various distances, adapted from [Lyness, 2004]. The X-axis indicates distance in meters from an observer. The Y-axis represents the effectiveness in terms of visual sensitivity to the cues. Lines in the graphs indicate depth cue effectiveness at different distances.

Depth Cues in Combination

- Quest for the relative values of cue combinations
  - It turns out to depend on the task (adding certain cues actually make some tasks harder)
- Quest for a small set of elementary tasks: here are some
  - Tracing data paths in 3D graphs
  - Judging the relative positions of objects in space
  - Judging the relative movement of self within the environment
  - Judging the "up" direction
  - Feeling a sense of "presence"
- We'll focus on these two in this lecture
  - Judging the shape of surfaces
  - Finding patterns of points in 3D space

Task: Understand Surface Shape

- Rule of Thumb:
  - Use stereo, kinetic depth, shape-from-shading, and texture if you can
  - Test for yourself which matter most when you can't do them all
- Studies (sometimes conflicting)
  - Judging heights of cones:
    - stereo >> structure-from-motion
  - Judging gradient of textured surface:
    - Structure-from-motion > stereo
  - How long have you been seeing spots?
    - Kinetic depth effect dominates for 4-6 seconds
    - Then stereoscopic depth became dominant

Task: Understanding Patterns of Points in 3D Space

- 3D scatter plot of points
  - Little perspective information available
  - Weak depth information from size gradient
  - Occlusion won't help for very small points
  - Cast shadows won't work – which shadow for a point?
  - Shape-from-shading missing for unlit points

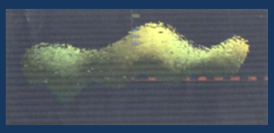

- What might work?
  - Stereoscopic depth
  - Structure-from-motion
  - Orient the points near boundaries, light the points to show cloud surface shape (based on gradient strength)

Immersive Display Systems Hints

- Head motion must change display naturally
  - Head-coupled perspective > stereo
  - HCP + stereo >> HCP
- Eye-hand relationship can vary
  - So long as there is not excessive lag
  - Think of the mouse (forward/back → up/down)

Credits

- This lecture is drawn from Chapters 7 and 8 of Colin Ware's "Information Visualization" book.
- Cast Shadows: RPGFan.com.