

# CS 498SL: Virtual Reality

Prof: Steve LaValle, Office: 3318 SC - 2-3pm TuTh  
 TAs: Apollo Ellis - computer graphics  
 Aldi Sipolins - psychology, VR

## About myself

- Ugrad & PhD at UIUC ECE
- Prof since 2001 in CS; taught 373 a lot
- Robotics researcher: motion planning, sensors, robots 3D geometry, simulation, physical world
- Oculus VR - 2012-2015
  - started by Palmer Luckey
  - I joined a few days after Kickstarter
  - Head tracking, sensor calibration & characterization, perceptual psychology, health & safety
  - Chief Scientist March 2013 - FB acquires

★ Slides: Palmer, Finland, Ciscia Babcia, Grandma Oculus  
 Course Administration

<http://courses.engr.illinois.edu/cs498sl/>

- Two midterm exams, no final
  - Machine Problems - 4 or 5
  - Final project - Teams coming
- work in pairs  
 → faculty pairing possible

VR Lab: 4240 Stebel Center  
 12 PCs w TitanBlack graphics cards 66B, 2 monitors, 27" 2  
 Oculus Rift DK 2s

VR Lab

Windows 8

Visual Studio

Unity Free game engine

Programming: - Unity scripts in C# or Java

- C++ with raw Oculus SDK

- Not in CS? Meet people outside!

Text books: - Mather, Foundations of Sensation &amp; Perception, 2nd ed. 2009

- Shirley et al., Fundamentals of Computer Graphics, 2009

→ Why these? - No modern book on VR  
 - Want technology invariant coverage  
 - Pieces exist in other fields  
 - Check Extra Material

Lectures: Blackboard + pictures

Class forum: Piazza

Social  
experiencesGoals:


- Learn how to build a good VR experience
- Understand how VR works (engineering + psychology)
- Learn how to criticize VR
- Learn fundamentals to shape future of VR

good = comfortable + adequate for task

↑  
no nausea

Task? Game, write code, maintain relationship, relax, watch film, virtual travel

5)



- learn how to build a good VR system
- understand how VR works (theory + application)
- learn how to create VR
- learn fundamentals to create future of VR

goal: comfortable + adaptive for task

Task: Done with the entire relationship between work and life

Definition of VR: (lots of bad defs out there)

Inducing targeted behavior in an organism by using artificial sensory stimulation, while the organism has little or no awareness of the interference.

Note: Feeding the organism. Not consciously aware

Note: Does not depend on particular technology, but "artificial" → technology of some kind

Common terms: immersion  
presence

\* orbil, treadmill, cockroach, omnitread

Other Concepts

Augmented reality: Like Google Glass

Telepresence: Imagine robot + omnidirectional camera + VR headset

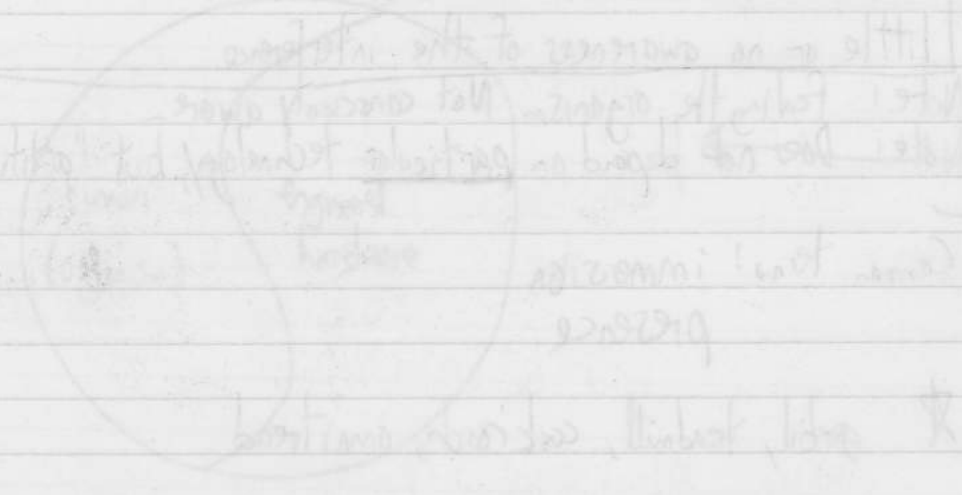
Teleoperation: Remote control

For me, the world inside the headset does not need to be virtual.

- Think also about:
- reading a book
  - talking on phone
  - playing a videogame
  - watching a movie
  - video conferencing
  - listening to music
  - Second Life
- } sort better

Amorents

- 1) Copy Ch1 Mather - scan online
- 2) Try Unity tutorials
- 3) Exam dates on Piazza
- 4) Project lot coming in day out



- reading a book
- talking a game
- looking at a video
- watching a movie
- listening to music
- playing a game

\* Some historical perspective:


- ☆
- oldest cave painting El Castillo hands 40,800 yrs old
  - Staring at rectangles: Medieval art, bad perspective (1470)
  - Perspective fixed: Vredeman de Vries (~1600)
  - Old movie: L'Arrivée D'un Train En Gare De La Ciotat, 1895, Lumière
  - old movie: A trip to the moon 1902
  - Gravity clip
  - Albert Pratt helmet gun — Like TF2 by Valve for Oculus 1916
  - Sensorama, Heilig 1957 stereo 3D, smells, seat vibrations, stereo sound  
↳ Made a headset in 1960 for 3D movies - Teleophore
  - Ivan Sutherland - 1965-8 Sword of Damocles - Video
  - CAVE - Projection system U. Illinois - Chicago 1991
  - VFX1 by Forte, mid 90s
  - Virtuality gaming systems
  - Lawnmower Man 1992
  - Nintendo Virtual Boy, 1995 - red lines & headaches, death of VR

Beware of hype!

Assignment: - Read over Unity Interface Overview & Manual  
"Extra material"

Right - Read Mather, Ch 1

VR system



2051  
1804 J-187

BIRD'S-EYE VIEW: HARDWARE

Forgotten from last time:

- Books: Mather, Foundations of Sensation and Perception, 2nd Ed., 2009  
 Shirley et al, Fundamentals of Computer Graphics, 2nd Ed., 2009
- Assignment: Read Mather, Ch1 by next Thu class; read Unity tutorials  
 - How to make good VR experiences  
 ↳ comfortable & adequate for task

\* Show slides and videos, to complete the history

BIRD'S-EYE VIEW

Pick a sensor on your body: eye, ear, finger, tongue, nose,  
↳ sense? def?

stereosmell  
with two nostrils??

Each sensor moves through space or changes in some way

- Controlled by brain (nervous system?)
- Degree Of Freedom (DOFs)
- Each sensor has a configuration space: set of all configurations

Ex Left ear

Position - 3DOFs } → 6DOFs total  
Orientation - 3DOFs

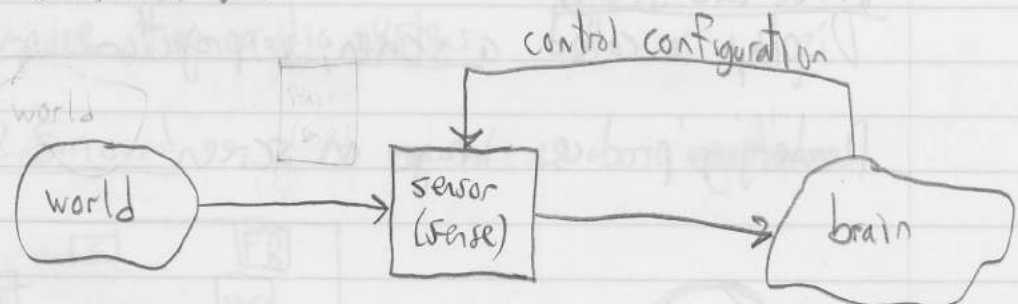
Right eye

More complicated: - 6DOF  
- Rotates with head fixed  
- Refocuses (accommodation)




# BIRDS EYE VIEW: HARDWARE


An abstract view!



Alter the world!



The brain is fooled!



# BIRDS EYE VIEW: HARDWARE

## Displays and rendering

## Audio and hearing

Display is called a speaker





Rendering produces sounds for the ear(s)

Need to track head to  
 show proper images  
 (or de Zeiss Cinema old glasses)

Video and seeing

Display is called a screen, or projector

Rendering produces images on screen

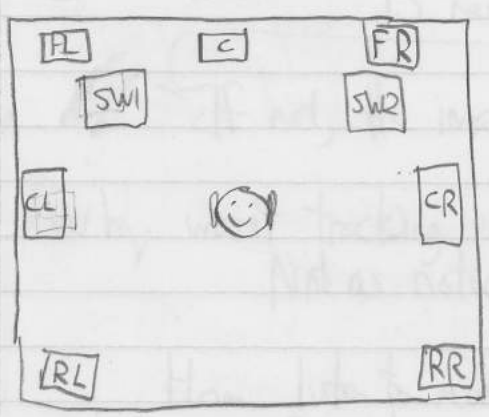


No tracking needed if:  
1) rotation only  
2) monoscopic



### Compare two audio systems

#### 7.2 Surround



#### earphones




- displays are far fr. eas
- localized audio for fixed room.

displays are close separated by R/L, follows head rotation

Note: Could use head tracking to simulate 7.2 using headphones

### Compare two VR visual systems

#### CAVE



#### HMD ( Rift, Morpheus )

Like "eyephones"



Need to track head to show proper images (or else Zeiss Cinemizer OLED glasses)

Make comparison to audio:

- Head tracking was not needed for 7.2 surround and CAVE
- " " IS needed for HMDs. Why?

↳ If not, the image will follow your head, unlike the real world

Why wasn't tracking used for headphones?

Not as noticeable for audio, as for video.

Hmm - listen to music with headphones and turn your head.  
Does it seem like the band is rotating around you?

They are fixed in 7.2

Same is true for CAVE and HMD

- To fool the brain, we need to undo the sensor transformations caused by head and body motion!

- Advantages to keeping display as close as possible to sensor!

- Lower cost,
- Lower power consumption
- Increased portability

Hey, why not neural interfaces? Cut the optic nerve!

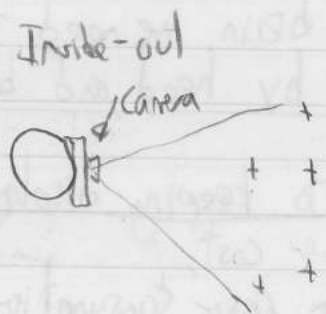
- Drawbacks: - Tracking needed  
- More work for the alternate world generator

Example: Body tracking for HMD, not CAVE


8)

Make (camera) to audio  
 - How tracking was not needed for 2.5 surround of CAVE  
 - 12 needed for HMDs Why  
 - Why want tracking up for headphones  
 - Not as noticeable for audio as for video  
 - How like to move with headphones as the headphones  
 - Do it seem like the sound is relating around you  
 - They are fixed in 2.5  
 - Same is the for CAVE as HMD

- To feel the door to undo the same tracking?



Outside-in



Regarding CAVE: Tracking not needed unless want to change viewpoint based on head motion. Includes stereo.  
Multiple viewers! Time multiplexing frames

## Buds-Eye View: VR Hardware

Tracking hardware components:

### - Inertial Measurement Unit (IMU):

- Originally designed for aircraft, spacecraft
- Mainly provide data for orientation estimation
- Now available in MEMS thanks to smartphones  
(InvenSense IMU-3000; ST, Microelectrics)
- From a few dollars to hundreds of thousands; calibration crucial
- Measure: 3D angular velocity (rad/sec) - gyroscope  
3D linear acceleration (m/sec<sup>2</sup>) - accelerometer

### - Magnetometer: 3D magnetic field (Gauss)

- Camera: VGA, 60FPS, etc → 1) Inside-out: On headset  
2) Outside-in: In room, fixed

### - Depth camera - like Kinect

What to track?

Head, eyes

Hands, whole body

Entire surrounding environment

Display components:

→ including vector display (skipped plasma)

- Visual: - (CRT, LCD/OLED) + lenses
- Light field display
- Virtual retina display


- Audio: - Standard speaker / earphones
- Bone conduction

- Touch: - Haptic (force feedback) devices
- keyboard, Xbox controller, mouse

Smell? Taste? Vestibular?

Visual rendering: Computer graphics (HW+SW) (GPU)

→ [Visual display] Resolution: # of pixels in each axis  
 → [Frame rate: Displayed frames per second (Hz)  
 Pixel switching speed  
 [Audio is similar with sampling rate & quantization levels]

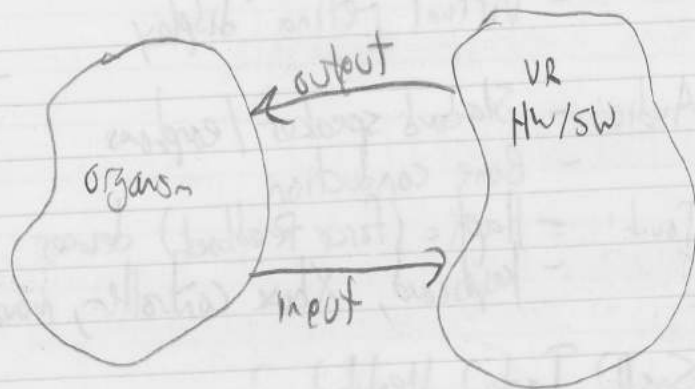


Computer: CPU, GPU

Lens - Palmer's law Fra Fry's

Input: Xbox controller, keyboard, mouse

10)



Handwritten notes in a box, possibly describing the components or context of the diagram.

Handwritten notes below the box, including the phrase "This is similar to sampling rate & quantization level".

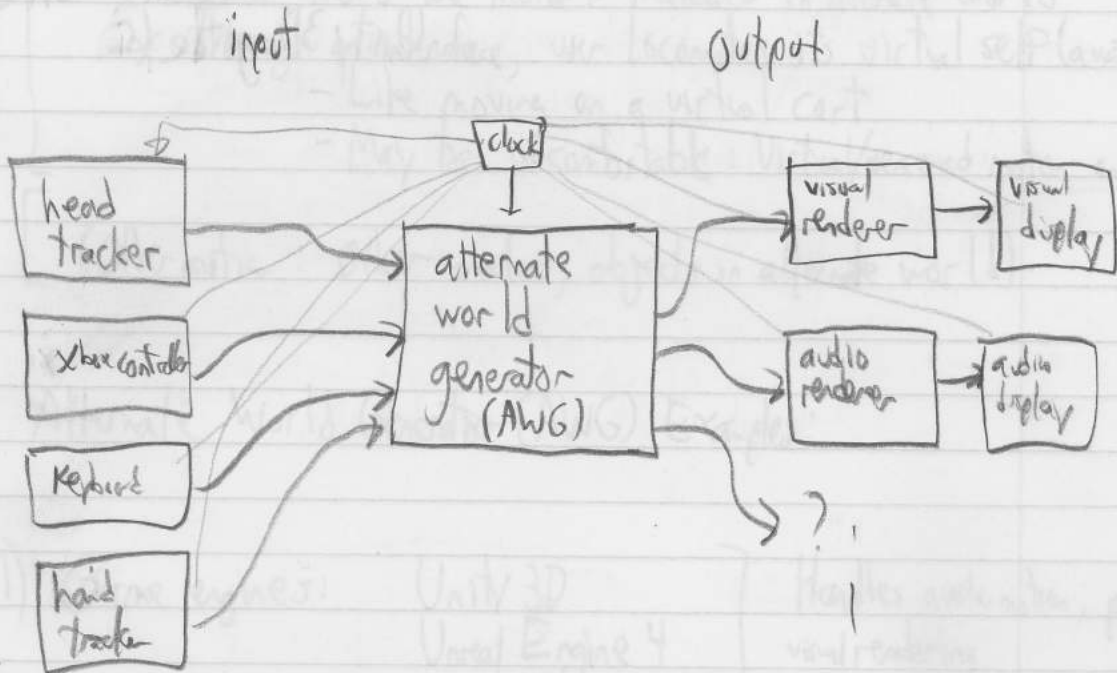
Computer: CPU, GPU

Low - Power for laptops

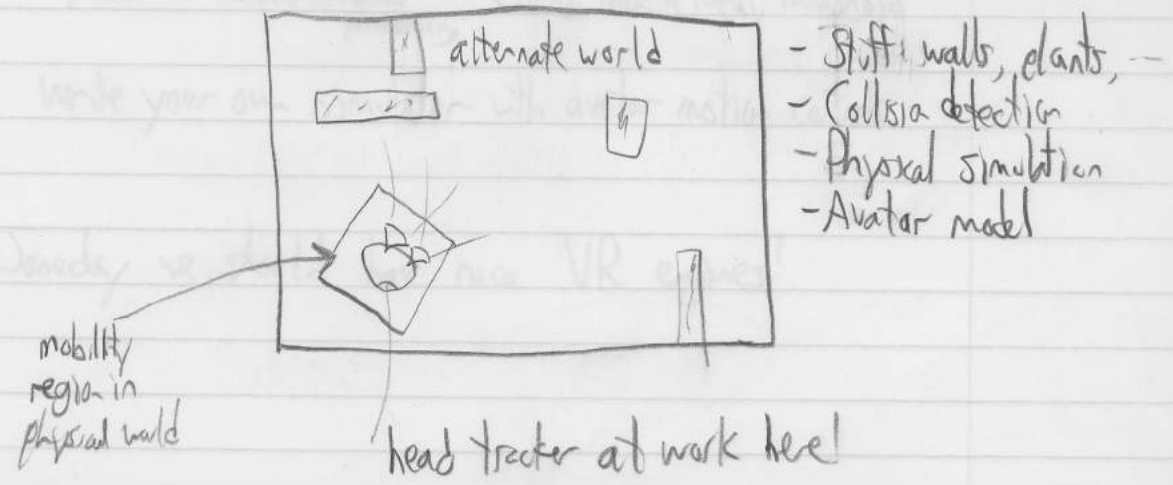
Light: Xeon, Intel, AMD, NVIDIA



# Birds-Eye View Software




Some geometry: Superimposing alternate and <sup>(real)</sup> physical worlds




(11)

Assignment: Read/Review Shirley Ch 2, 5, 6  
Lakalle Ch 3, 4  
Unity 3D guide



Unity 3D guide, Unity 3D guide, Unity 3D guide

- Unity 3D guide
- Unity 3D guide
- Unity 3D guide
- Unity 3D guide



Unity 3D guide

Unity 3D guide

## Two main kinds of motions:

- self
- 1) User moves and motion is matched in alternate world
  - 2) Through an interface, user locomotes its virtual self (avatar)
    - Like moving on a virtual cart
    - May be uncomfortable: Virtual/perceived motion only
- others
- (Other motion: other avatars, objects in alternate world)

## Alternate World Generator (AWG) Examples:

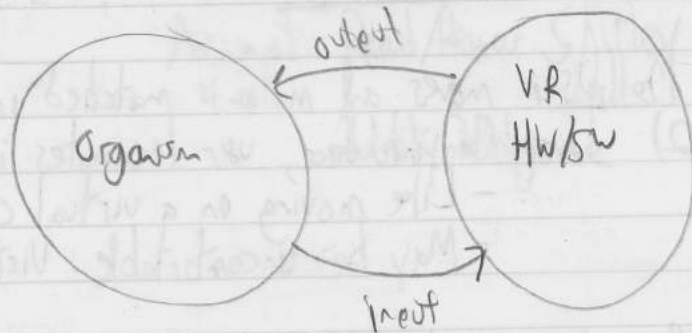
- 1) Game engines: Unity 3D, Unreal Engine 4 } Handles avatar motion, physics, visual rendering

Powerful, but not optimized for VR.

- 2) Google Street View server + interface for navigation
- 3) robot + camera + video processing → drone, mobile robot, humanoid
- 4) Write your own simulator with avatar motion controls.

Someday, we should have nice VR engines!

(13)



Alternate World (Avatar) (AW) Example

- 1) Game engine - Univ 3D
- 2) Unreal Engine 4
- 3) Unreal Engine 5

Playable, but not optimized for VR

- 1) Game Start New scene + interface for navigation
- 2) Level + camera + controls
- 3) Write your own simulator with custom controls +


Ready to start building VR engine!

Bird's-Eye View: Sensation & Perception  
 (Perceptual psychology; vision science)

VR questions

- How do we perceive depth? (3D)
- What causes nausea, fatigue?
- How do we perceive an audio source?
- How much resolution is enough in a display?
- What is presence?

Our perception process appears to be simple:  
 Immediate, effortless, direct




Human brain

- Cerebral cortex:
- ~ 10 billion neurons
  - outer layer, folded sheet 2.5mm thick
  - most advanced part of brain
  - most neurons devoted to vision perception

Computers can't beat cerebral cortex!

↳ really?

Like a big directed graph:



\* Show optical illusions + Balance (Vestibular)

Classification of Senses

<u>Sense</u>	<u>Stimulus</u>	<u>Receptor</u>	<u>Sensory structure</u>	(optional) Robot Equivalent
Vision	Electromagnetic energy	Photoreceptors	Eye	Camera
Hearing	Air pressure waves	Mechanoreceptors	Ear	microphone
Touch	Tissue distortion	Mechanoreceptors Thermoreceptors	Skin, muscle, -	pressure, force sensors thermometer
Balance (Vestibular)	Gravity, acceleration	Mechanoreceptors	Vestibular organs	IMU
Taste/Smell	Chemical compounds	Chemoreceptors	Mouth/Nose	ph meter?

Order of importance for our class:

- 1) VISION!
- 2) Balance
- 3) Hearing
- 4) Touch
- 5) Taste/Smell



The brain fuses these to make a coherent experience

Most important for us: Vision + Balance (Vestibular)

- 1) Vection → Optical Flow disagrees with vestibular
- 2) Vestibulo Ocular Reflex (VOR)
  - Affects sense of stationarity
  - Eye counter rotates with head rotation to fixate eyes on a target.

Some fundamental concepts in sensation and perception


Receptors: Transducers that convert stimulus energy to neural impulses (Much like an engineered sensor to an analog or digital signal)



Hierarchical Processing!


Sensory system selectivity

- Each cell has a receptive field
- Limited frequency, color, shape, -- for each (tiny "activation zone")
- Electrodes to perform single-unit recordings (Display particular stimulus, determine whether cell activates)



Organization of neurons (network spatial topology!)

Topographic Map

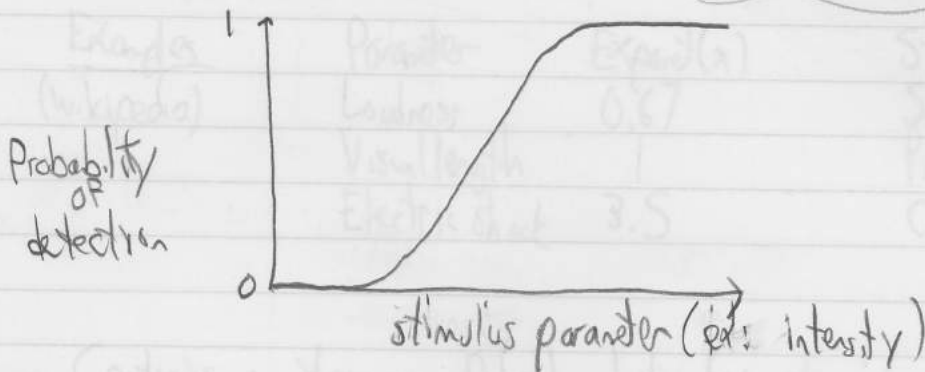


striate (visual) cortex  
left hemisphere

Psychophysics: Scientific study of:

physical stimuli → perceptual phenomena (conscious mental state)

What range of stimuli produce "red"? "sour"? "dark"? ...?



→ "qualia"

in Criticism: You can fit the data by changing  $c$  and  $\lambda$ , but is the law really "correct"?

Related: Just noticeable difference (JND)




(21)

# Joke: Stigler's Law of Eponymy

No scientific law is named after its original discoverer

The discoverer: Robert Merton



(un)scientific  
what  
judicial Po

law  
the

Psychophysics: scientific studies

psychophysics → (scientific studies)

What task is it? (e.g. 'back')



Probability  
detection

Stimulus quantity (e.g. intensity)

Magnitude estimation:

Steve's power law (1961) - see wikipedia




$m$  = magnitude or intensity of stimulus  
 $p$  = perceived magnitude

$$m \mapsto cm^x$$

$$p = cm^x$$

$x$  = exponent that depends on particular experiment  
 (could have  $x < 1$ ,  $x > 1$ ,  $x = 1$ )

$c$  = constant based on units, etc.



Examples (wikipedia)	Parameter	Exponent ( $x$ )	Stimulus
	Loudness	0.67	Sound 3000Hz tone
	Visual length	1	Projected line
	Electric shock	3.5	Current through fingers

Criticism: You can fit the data by choosing  $c$  and  $x$ , but is the law always "correct"?

Related: Just-Noticeable Difference (JND)

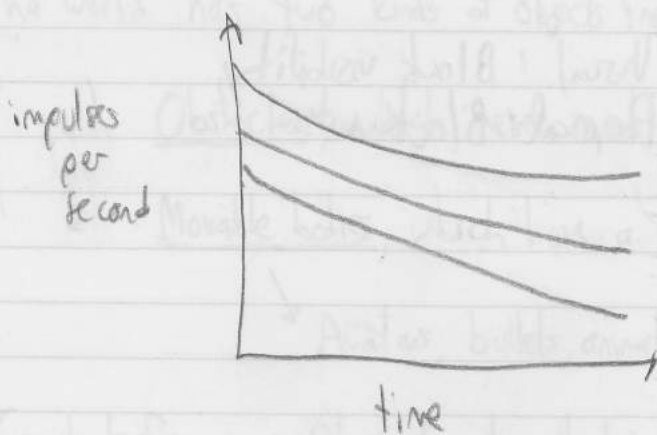
Weber's law:

$$\frac{\Delta m}{m} = c$$

(in other words, only percent changes matter)

$\Delta m$  is the JND of stimulus magnitude

Plasticity & adaptability



Coordinate frame = description, specified in world frame  
 - Each rigid body described in its own body frame

Let's model obstacles, but some principles apply for bodies

- Two choices:
- 1) Solid representation
  - 2) Boundary representation

Polynomial Time-Non-Reduction (PTN)

Worst case

(more tangible show-stops in action)

$$c = \frac{m\Delta}{m}$$


Maybe better ones:

- 1) Stationary models
- 2) Movable models

Δ is the TND of study magnitude

Physicality & complexity



- Obstacles:
- 1) Visual: Block visibility
  - 2) Physical: Block motion



exit

## Geometric Modeling

Let  $W$  be a 3D world:  $W \subseteq \mathbb{R}^3$



Every point in  $W$  described by coordinates  $(x, y, z) \in W$

The world has two kinds of objects inside!


- 1) Obstacles, which never move: walls, furniture,
- 2) Movable bodies, which have a space of possible transformations (configuration space)  
 ↓  
 Avatars, bullets, animab, blowing leaves, coffee cup

Coordinate frames: = Obstacles, described in world frame  
 - Each (rigid) body described in its own body frame

Let's model obstacles, but some principles apply for bodies

- Two choices:
- 1) Solid representation
  - 2) Boundary representation


Boundary representation is easier, but almost always hard to distinguish "inside" from "outside"



Worst-case model: Polygon soup.

Think about real-world data from cameras, Kinect,

Let's use triangle primitives



Graphics problem: How to render the surface?

- Compute normal — ambiguity problem



Strip

- approximate any surface
- saves vertices



Mesh

- might have numerical holes
- GPUs like these!

## Transforming Rigid Bodies


Transform each vertex of each triangle

		DOFs	
		2D	3D
3 cases:	Easy: Translation	2	3
	Harder: Rotation	1	3
	Hardest: Rotation + Translation	3	6

### Translation

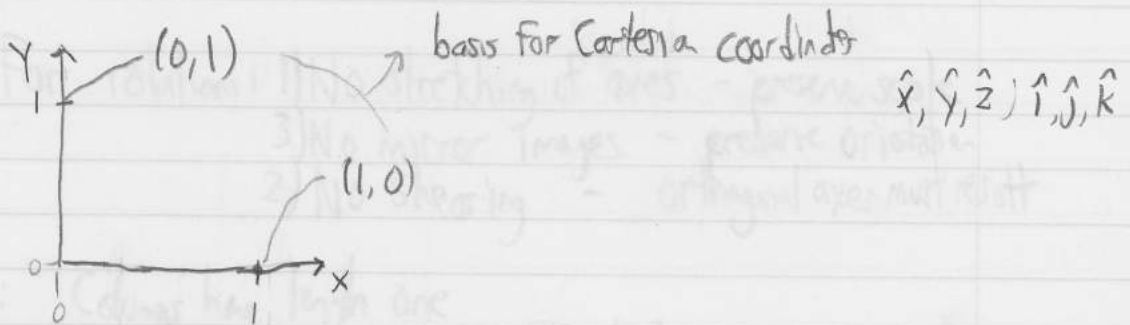
Shift triangle by  $(x_t, y_t, z_t)$

$$\begin{aligned} (x_1, y_1, z_1) &\rightarrow (x_1 + x_t, y_1 + y_t, z_1 + z_t) \\ (x_2, y_2, z_2) &\rightarrow (x_2 + x_t, y_2 + y_t, z_2 + z_t) \\ (x_3, y_3, z_3) &\rightarrow (x_3 + x_t, y_3 + y_t, z_3 + z_t) \end{aligned}$$



We need to get to 3D rotation. Start with 2D linear transformations:

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$



$$\begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} M_{11} \\ M_{21} \end{bmatrix}$$

$$\begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} = \begin{bmatrix} M_{12} \\ M_{22} \end{bmatrix}$$

Exaples

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

no effect

$$\begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}$$

scale

$$\begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix}$$

mirror

$$\begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix}$$

rotate 180

$$\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$$

x-shear



$$\begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}$$

y-shear



To rotate (x,y):

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} x' \\ y' \end{bmatrix}$$

A circle of radius r  
 $M_x = \cos\theta$   
 $M_y = \sin\theta$



For rotations: 1) No stretching of axes - preserve scale  
 2) No shearing - orthogonal axes must result  
 3) No mirror images - preserve orientation

4 DOFs [!]

-1 1: Columns have length one

$$m_{11}^2 + m_{21}^2 = 1 \quad m_{12}^2 + m_{22}^2 = 1$$

-1 2: Columns are "perpendicular", Inner-products = 0


$$m_{11}m_{12} + m_{21}m_{22} = 0$$

-0 3:  $\det([: :]) = 1$  (not -1)  
 ↘ mirror image

DOFs = 1

Which matrices remain as rotations?

$$\begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix} \text{ for all } \theta \in [0, 2\pi)$$



To rotate  $(x, y)$ :


$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

A circle of rotations

$$m_{11} = \cos\theta \\ m_{21} = \sin\theta$$



Now try 3D rotations. Which 3x3 matrices are rotations?



- 1)  $\|v_1\| = \|v_2\| = \|v_3\| = 1$
- 2)  $v_1 \cdot v_2 = 0 \quad v_2 \cdot v_3 = 0 \quad v_1 \cdot v_3 = 0$
- 3)  $\det\begin{pmatrix} | & | & | \\ v_1 & v_2 & v_3 \\ | & | & | \end{pmatrix} = 1$

9 DOFs

-3

-3

-0

→ 3 DOFs

A 3D "surface" (manifold)

3 Canonical Rotations: Yaw, Pitch, Roll

Yaw:  $R(\alpha) = \begin{bmatrix} \cos \alpha & 0 & \sin \alpha \\ 0 & 1 & 0 \\ -\sin \alpha & 0 & \cos \alpha \end{bmatrix}$  leave y coordinate untouched

Pitch:  $R(\beta) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \beta & -\sin \beta \\ 0 & \sin \beta & \cos \beta \end{bmatrix}$  leave x untouched  
→ exactly 2D rotation

Roll:  $R(\gamma) = \begin{bmatrix} \cos \gamma & -\sin \gamma & 0 \\ \sin \gamma & \cos \gamma & 0 \\ 0 & 0 & 1 \end{bmatrix}$  leave z untouched

matrix mit

Every 3D rotation  $R$  can be achieved by some  $\alpha, \beta, \gamma$  with  $R = R(\alpha)R(\beta)R(\gamma)$   
with  $\alpha \in [-\pi, \pi], \beta \in [-\pi, \pi], \gamma \in [-\pi/2, \pi/2]$  → 7) error?

Example

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & \alpha & -s\alpha \\ 0 & s\alpha & \alpha \end{bmatrix} \begin{bmatrix} c\beta & 0 & s\beta \\ 0 & 1 & 0 \\ -s\beta & 0 & c\beta \end{bmatrix} \begin{bmatrix} c\gamma & -s\gamma & 0 \\ s\gamma & c\gamma & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Path  $\beta = \pi/2 \rightarrow \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ -1 & 0 & 0 \end{bmatrix}$

Result:

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & s\alpha & -\alpha \\ 0 & -\alpha & s\alpha \end{bmatrix} \begin{bmatrix} c\gamma & -s\gamma & 0 \\ s\gamma & c\gamma & 0 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 1 \\ s(\alpha\gamma) & c(\alpha\gamma) & 0 \\ -c(\alpha\gamma) & s(\alpha\gamma) & 0 \end{bmatrix}$$

apply trig identity



(Roll+Yaw) yields same orientation, regardless of Roll at yaw angles.

Problems


- Order matters: 3D rotations not commutative!!!
- Kinematic singularities and nonuniform representation
  - Vary  $\alpha, \beta, \gamma$  at constant speed, and rotation varies wildly
  - Gimbal lock

(Illustrate with prop) - Problem: Rotation axis keeps changing on body


Euler's rotation theorem? All 3D rotations have an axis-angle representation  
(1776) for orientations



Note



and



are the same rotation!

It would be nice to have some algebra

$$(v, \theta) \circ (v', \theta') = ?$$

so that we always know the resulting axis and angle. keeps changing! (Not in 2D case)

(In OpenGL, exponential coordinates used:  $(\theta v_1, \theta v_2, \theta v_3)$ )  
scale is the angle!

Nice! 3 parameters for 3DOFs  
 $(0, 0) = \text{identity}$

Note: 3DOFs, as needed for 3D rotation

$$\begin{bmatrix}
 1 & 0 & 0 & 0 \\
 0 & \cos(\theta) & -\sin(\theta) & 0 \\
 0 & \sin(\theta) & \cos(\theta) & 0 \\
 0 & 0 & 0 & 1
 \end{bmatrix}$$

$$\vec{v}' = \mathbf{A}(\theta) \cdot \vec{v}$$

(To obtain, extended coordinates,  $(\theta, \dot{\theta}, \ddot{\theta})$ )

state vectors!

Use equations for 3DOF  
 $\dot{x}(0) = \dot{x}_0$

Unit quaternions as a better representation of 3D rotation.

$q = (a, b, c, d) \in \mathbb{R}^4$  and  $a^2 + b^2 + c^2 + d^2 = 1 \Rightarrow$  "normalized"

Set of all  $q$  is a hypersphere ( $S^3$ )

Unity3D/games representation:  $(x, y, z, w) = (b, c, d, a)$

careful!

Most common representation:


$a + bi + cj + dk$

"real" part

"imaginary" parts like complex

numbers  $a + bi$

Encoding a 3D rotation using unit quaternions



$v = (v_x, v_y, v_z) \quad \|v\| = 1$

$\longleftrightarrow (\underbrace{\cos \theta/2}_a, \underbrace{v_x \sin \theta/2}_b, \underbrace{v_y \sin \theta/2}_c, \underbrace{v_z \sin \theta/2}_d)$

arccos to recover angle

Just normalize to recover axis

$\frac{b, c, d}{\sqrt{b^2 + c^2 + d^2}}$

Can't be commutative!!!

Useful examples:

$(1, 0, 0, 0)$	identity rotation
$(0, 1, 0, 0)$	pitch by $\pi$ X
$(0, 0, 1, 0)$	yaw by $\pi$ Y
$(0, 0, 0, 1)$	roll by $\pi$ Z
$(\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}, 0, 0)$	pitch by $\pi/2$
$(\frac{1}{\sqrt{2}}, 0, \frac{1}{\sqrt{2}}, 0)$	yaw by $\pi/2$
$(\frac{1}{\sqrt{2}}, 0, 0, \frac{1}{\sqrt{2}})$	roll by $\pi/2$

↑  
the bigger the axis part, the bigger the rotation

Inverses and multiple representations

$$(a, b, c, d) \xleftrightarrow{\text{equivalent}} (-a, -b, -c, -d)$$

$$(a, b, c, d) \xleftrightarrow{\text{inverse}} (-a, -b, -c, -d)$$

$$(-a, -b, -c, -d) \xleftrightarrow{\text{equivalent}} (-a, b, c, d)$$

Multiplication

$$\text{let } p = (b, c, d)$$

$$q_1 \circ q_2 = q_3$$

$$q_1 \circ q_2 = (a_1 q_2 - p_1 \cdot p_2, \underbrace{p_1 \times p_2 + a_1 p_2 + a_2 p_1}_{a_3, b_3, c_3, d_3})$$

Can't be commutative!!!

R - useful because "native"; rotate directly

q - no singularities

$\pi$  - small change in q = small change in orientation (just like  $\theta$  in 2D rotation)

$\pi$  - normal rotation, hemispheres

YPR - intuitive

- bad singularities!

rotation of rigid body

inverse of multiple rotations

$$(b, c, d, a) \leftrightarrow (a, b, c, d)$$

$$R_3^{-1} R_2^{-1} R_1^{-1} R_1 R_2 R_3 = R_3^{-1} R_2^{-1} R_3 = R_3^{-1} R_3 = I$$

$$(b, c, d) = q + t$$

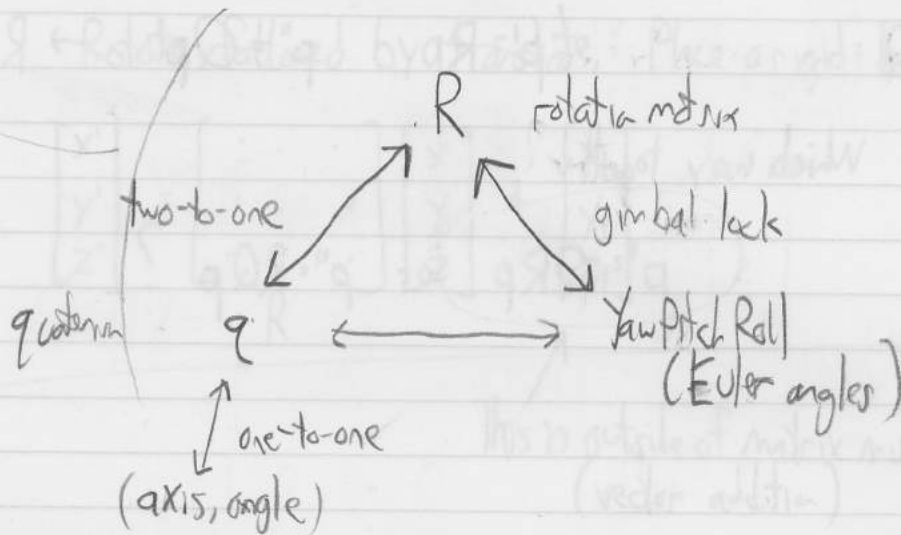
$$p = s p^0 p$$

$$(q \cdot p + q \cdot p + p \times q, q \cdot q - p \cdot p) = p \cdot p$$

$$p_1^2 + p_2^2$$

Can't be commutative!!



Conversions

$$R_1 \cdot R_2 = R_3$$

$$\downarrow \quad \downarrow \quad \uparrow$$

$$q_1 \circ q_2 = q_3$$

Same result either way: Find a proof!

Note: Can apply  $q$  directly to rotate  $(x, y, z)$ .

Operator order: Which is applied first?  $R_1 R_2$

$$R_1 (R_2 \begin{bmatrix} x \\ y \\ z \end{bmatrix}) \rightarrow R_2 \text{ is!}$$

Beware of inverses!

$$(R_1 R_2 R_3)^{-1} = R_3^{-1} R_2^{-1} R_1^{-1} \quad (\text{linear algebra group theory})$$

(85)

Which is first?

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

R Q

$$p' = Rp$$

$$p'' = Qp'$$

→ R before Q

Which way together?

$$p'' = QRp \text{ or } p'' = RQp ?$$

First make

$$\begin{bmatrix} \checkmark R & & & 0 \\ & & & 0 \\ & & & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

Rotate using quaternions  $q$  directly

$$p = (x, y, z, 1) \quad p' = qpq^{-1}$$

$$q^{-1} = (a, -b, -c, -d)$$

Homogeneous transformation matrices - A math hack

Rotate followed by translate: Place a rigid body anywhere

$$\begin{bmatrix} x' \\ y' \\ z' \end{bmatrix} = \underbrace{\begin{bmatrix} \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot \end{bmatrix}}_R \begin{bmatrix} x \\ y \\ z \end{bmatrix} + \begin{bmatrix} x_t \\ y_t \\ z_t \end{bmatrix}$$

↑  
This is outside of matrix multiplication  
(vector addition)

Make a  $4 \times 4$  matrix that is algebraically equivalent  
(and 4D vector)

$$\begin{bmatrix} x' \\ y' \\ z' \\ 1 \end{bmatrix} = \begin{bmatrix} \boxed{\begin{matrix} R \\ \hline 0 & 0 & 0 \end{matrix}} & \begin{bmatrix} x_t \\ y_t \\ z_t \end{bmatrix} \\ \hline & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}$$

Same result, but extra "1" in 4th position

Now you get one  $4 \times 4$  matrix  $T = \begin{bmatrix} R & \begin{matrix} x_t \\ y_t \\ z_t \end{matrix} \\ \hline 0 & 0 & 0 & 1 \end{bmatrix}$  that places rigid bodies

Important question: What is  $T^{-1}$ ? (How to undo the placement?)

What is  $R^{-1}$ ?

$$R^{-1} = R^T \begin{matrix} \swarrow \text{transpose} \\ \uparrow \\ \text{only for "orthogonal matrices"} \\ \text{orthonormal} \end{matrix}$$

Opposite of  $x_1, y_1, z_1$  is  $-x_1, -y_1, -z_1$

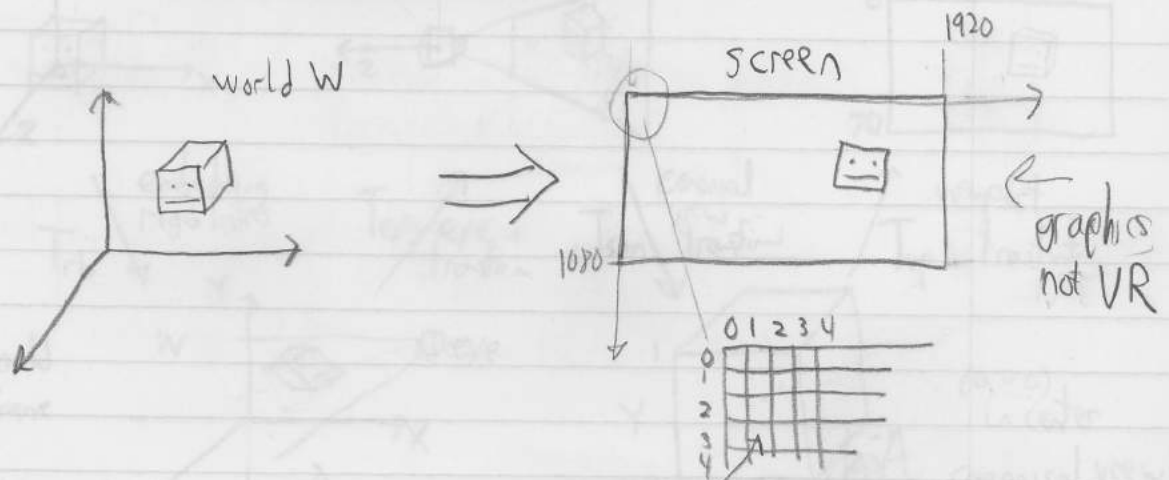
Is  $T^{-1} = \begin{bmatrix} \text{RT} & -x_1 \\ & -y_1 \\ & -z_1 \\ 0 & 0 & 0 & 1 \end{bmatrix}$  ? No!!

Wrong order of operations

$$T^{-1} = \begin{bmatrix} \text{RT} & 0 \\ & 0 \\ & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & -x_1 \\ 0 & 1 & 0 & -y_1 \\ 0 & 0 & 1 & -z_1 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

Now let's worry only about transforming points from W to screen  
 (points, line segments, triangles)  
 The remainder is the rendering pipeline

## Viewing Transformations



Questions!

How should artificial light propagate?

How to discretize geometric model onto screen?

How to determine what is visible or blocked?

How to make correct perspective?


Efficiency??

Take a shortcut if flaws are not perceptible

Now let's worry only about transforming points from  $W$  to screen  
(points, line segments, triangles, ...)

The remainder is the rendering pipeline.

### Viewing Transformations



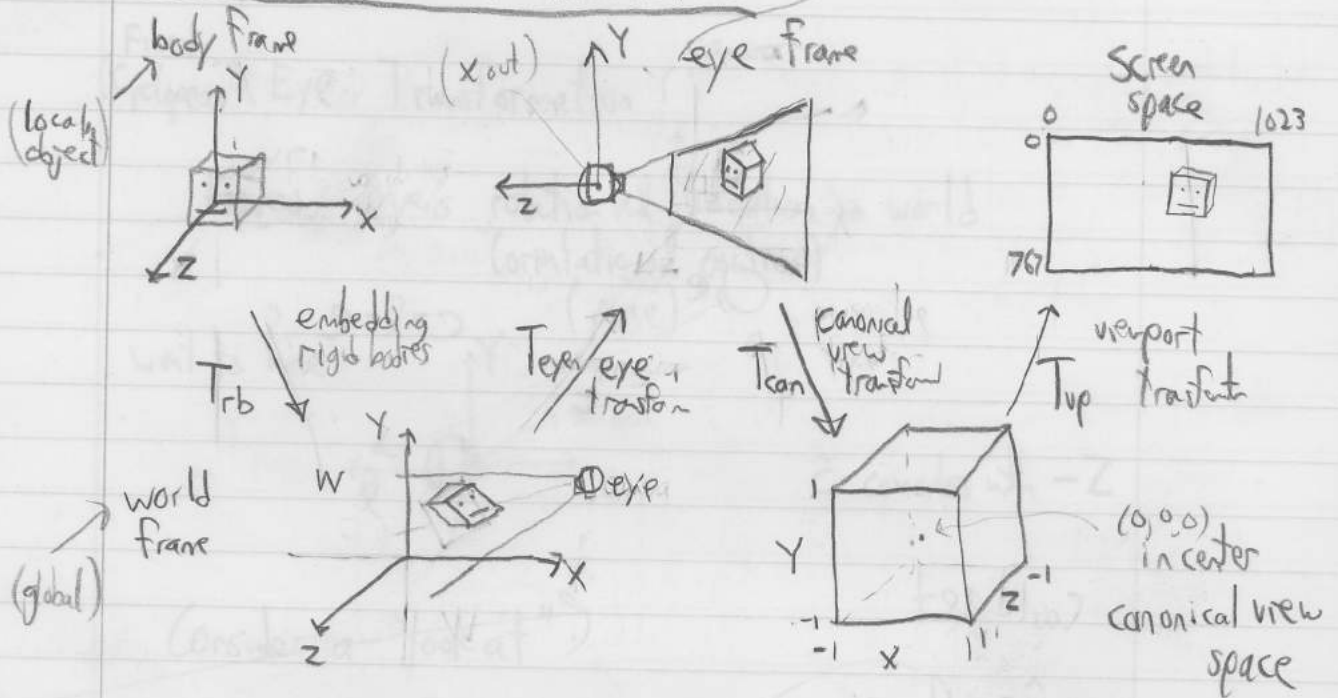
Apply

$$T \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} i \\ j \end{bmatrix}$$

### Units

- Body frame, world frame, eye frame: meters (or feet)
- Canonical view frame: none
- Screen frame: pixel index

# The Chain of Transformation → cyclopean!



Let  $T$  denote a  $4 \times 4$  homogenous transform matrix for viewport from body coordinates

$$T = T_{vp} T_{can} T_{eye} T_{rb}$$

Insert  $T_e, T_r$  for stereo rendering

Insert  $T_{dist}$  to compensate for lens distortion → for VR!

Look at each transformation:

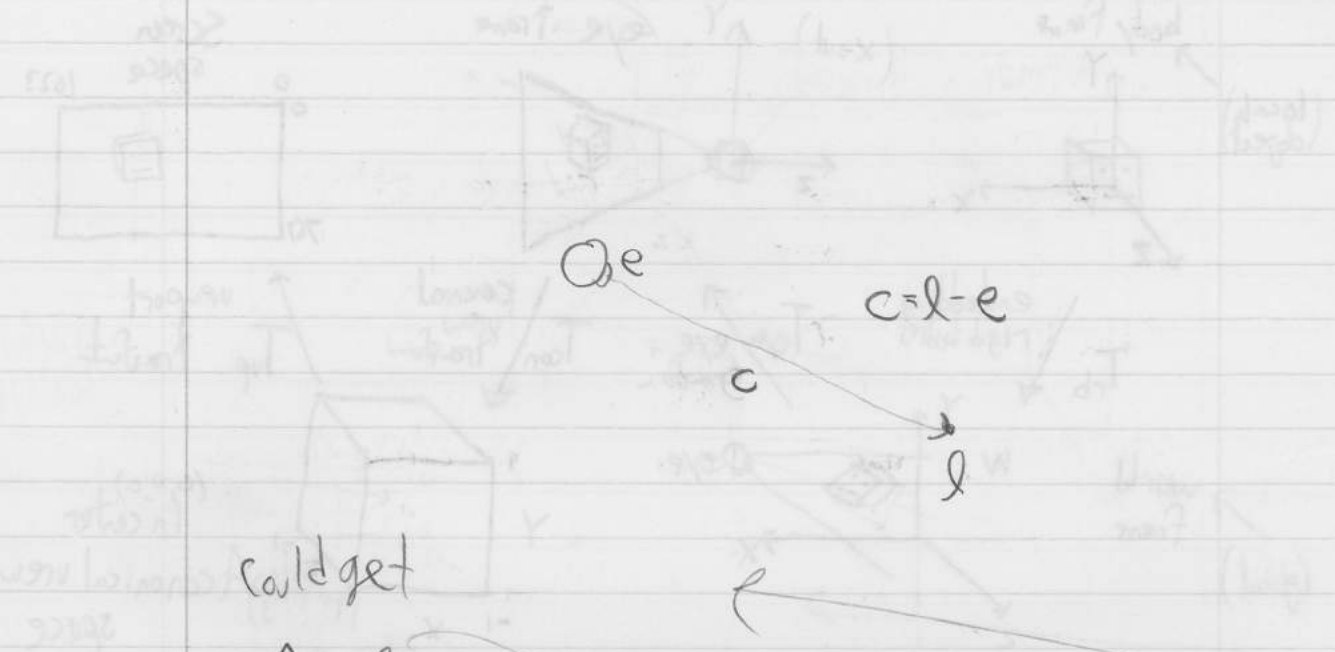
$T_{rb}$  - already covered (rotation + translation)

$T_{eye}$  - becomes an inverse of a rigid body placement

$T_{can}$  - apply perspective transform and rescale, translate

$T_{vp}$  - units conversion and more rescale, translate

# The Chain of Transformation



could get

$$\hat{c} = \frac{l - e}{\|l - e\|}$$

$l = \text{look-at point } (l_x, l_y, l_z)$

Why is  $\hat{y}$  not simply  $\hat{c}$ ?


This allows sloppy input!  $\hat{c}$  need not be perpendicular to image plane.



## Fixed (Cyclopean) Eye Transformation

Express eye's rotation and translation in world  
(orientation and position)  
(pose)

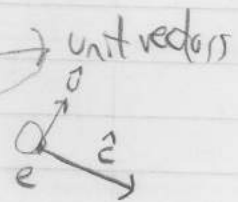
want to have:



$\hat{c}$  coincides with  $-Z$

Consider a "look at":

1. Position of eye:  $e$
2. Looking direction in eye, center:  $\hat{c}$
3. Up direction:  $\hat{u}$




Make a transform to place eye in W

Coordinate axes for eye:

$$\begin{aligned}\hat{x} &= \hat{u} \times \hat{z} \\ \hat{y} &= \hat{z} \times \hat{x} \\ \hat{z} &= -\hat{c}\end{aligned}$$

$$R_{eye} = \begin{bmatrix} \hat{x}_1 & \hat{y}_1 & \hat{z}_1 \\ \hat{x}_2 & \hat{y}_2 & \hat{z}_2 \\ \hat{x}_3 & \hat{y}_3 & \hat{z}_3 \end{bmatrix}$$

To transform world into eye view, need inverse transform!



$$T_{\text{lefteye}} = \begin{bmatrix} 1 & 0 & 0 & +t/2 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} T_{\text{eye}}$$

$$T_{\text{righteye}} = \begin{bmatrix} 1 & 0 & 0 & -t/2 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} T_{\text{eye}}$$

$t$  = inter eye distance  
in virtual world



( $t = 0.064\text{m}$  on  
average in  
real world)



Interpretations: 1) Dragging the world as if it were a rigid body  
 2) Tilt eye at origin ( $e = (0,0,0)$ )  
 Compare: - Yaw world by  $\theta$  (without moving eye)  
 - Yaw eye by  $-\theta$  (" " world)

Thus,  $R_{eye}^T = R_{eye}^{-1}$

$$T_{eye} = \begin{bmatrix} \hat{x}_1 & \hat{x}_2 & \hat{x}_3 & 0 \\ \hat{y}_1 & \hat{y}_2 & \hat{y}_3 & 0 \\ \hat{z}_1 & \hat{z}_2 & \hat{z}_3 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} -e_x \\ -e_y \\ -e_z \\ 1 \end{bmatrix}$$

→ Inverse translation


Stereo  
Canonical View Transform



- projective geometry  
 (space of lines through origin)

Perspective transform:  $h = \frac{d}{z_p} y_p$  → dividing by  $z_p$  is not linear

(PE)



The transformed corners of viewing frustum:

$(l, b, n)$   
 $(r, t, f)$



Tricking matrix multiplication to do the algebra

1D example!  
(homogeneous)

$$\begin{bmatrix} y' \\ h' \end{bmatrix} = \begin{bmatrix} d & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} y_p \\ z_p \\ 1 \end{bmatrix} = \begin{bmatrix} dy_p \\ z_p \end{bmatrix}$$

Define  $h := \frac{y'}{h'} = \frac{dy_p}{z_p}$  (same as previous figure)

n, f

perspective  $T_p =$

$$\begin{bmatrix} n & 0 & 0 & 0 \\ 0 & n & 0 & 0 \\ 0 & 0 & n+f & -fn \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

$$x' = nx \rightarrow \left( \frac{x'}{h'}, \frac{y'}{h'} \right)$$


$$y' = ny$$

$$z' = z(n+f) - fn$$

strange: maintains both ordering (used in pipeline)  
 $h' = z$

Scale and Translate  $T_{st} =$

$$\begin{bmatrix} \frac{2}{r-l} & 0 & 0 & \frac{l+r}{2} \\ 0 & \frac{2}{t-b} & 0 & \frac{b+t}{2} \\ 0 & 0 & \frac{2}{n-f} & \frac{n+f}{2} \\ 0 & 0 & 0 & 1 \end{bmatrix}$$



$x$ : l..r (left-right)  
 $y$ : b..t (bottom-top)  
 $z$ : n..f (near-far)

$$T_{can} = T_{st} T_p$$

(Truth in Shurley)

One remaining part:  $T_{up}$  converts  $-1 \dots 1$  range to pixel coordinates

$n_x = \#$  of horizontal pixels (eg. 1920)

$n_y = \#$  of vertical pixels (eg. 1080)

$$T_{up} = \begin{bmatrix} n_x/2 & 0 & 0 & n_x/2 \\ 0 & n_y/2 & 0 & n_y/2 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} i \\ j \end{bmatrix} \leftarrow \text{pixel indices}$$

Point source of light

- 6 rays horizontal

- Rays are orthogonal to wavefront



without mirrors or lenses, rays always diverge



ray  
ray  
ray

(central) ray  
(central) parallel ray

# Light


Three interpretations:

- 1) particles photons


frequency

$$f = \frac{c}{\lambda}$$

- wavelength  
speed of light

- 2) rays 

Computer scientists love this!  
(graphics, computational geometry)


- 3) waves 

Point source of light:

- Generates wavefronts
- Rays are orthogonal to wavefronts

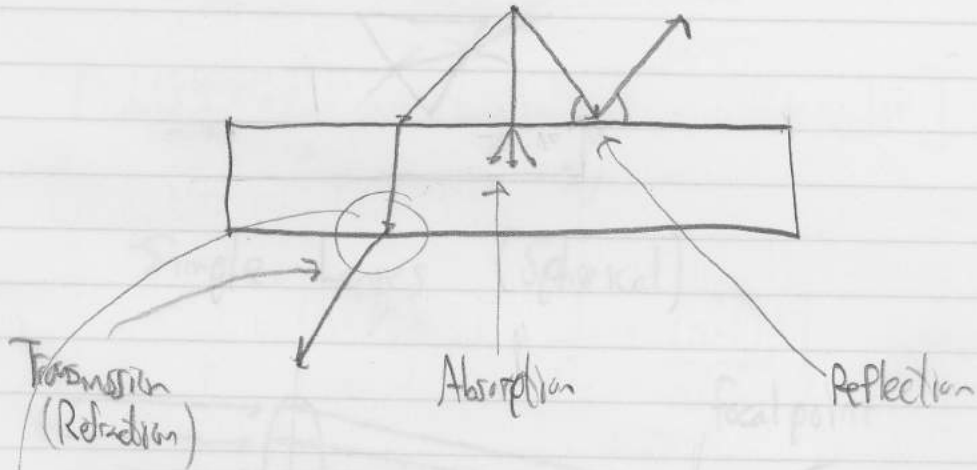




Without mirrors or lenses, rays always diverge




$\rightarrow$  parallel rays  
 $\rightarrow$  collimated rays  
 zero vergence  
 rays to or from infinity

We use materials to bend the rays/waves



Reflection types: Specular   
 Diffuse 

Snell's Law




$$\frac{n_2}{n_1} = \frac{\sin \theta_1}{\sin \theta_2}$$

$$n_i = \frac{c}{\text{speed in medium } i}$$


refractive index




### Refraction

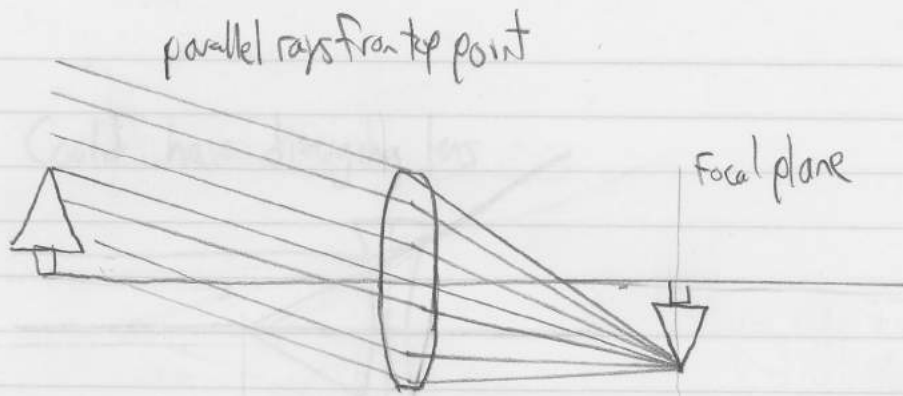


### Simple Lenses (Spherical)



### wave front interpretation






[could do axial point (tree bottom) in another color]

Lensmaker's equation

$$\frac{1}{f} = (n-1) \left[ \frac{1}{R_1} - \frac{1}{R_2} + \frac{(n-1)d}{nR_1R_2} \right]$$


index of refraction  $n$



Thin approximation:  $d \approx 0$

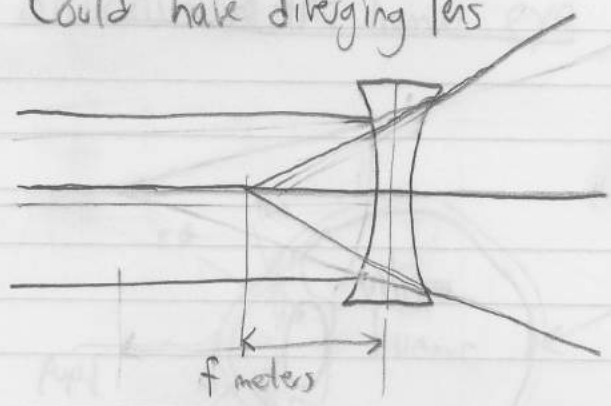
$$\frac{1}{f} = (n-1) \left[ \frac{1}{R_1} - \frac{1}{R_2} \right]$$

Convenient unit: Diopter ( $m^{-1}$ ) converging/diverging power of lens



5D (5 diopters) or  $\left(\frac{1}{f}\right) D$


Could have diverging lens




Lens power  $\left(\frac{-1}{f}\right)D$  → negative because lens causes divergence



$D < 0$



$D = 0$



$D > 0$


Combining lenses or interfaces!

$$D = D_1 + D_2 + D_3 \dots$$

Assumes distance between lenses or surfaces  $\hat{=}$  0  
(thin lens approximation)


↳ (Optometrists must not like algebra!)

(11)



Could have diverging lens

for lens  $(\frac{1}{f}) = D \rightarrow$  Negative power lens gives



Combining lenses in series:

$$D = D_1 + D_2 + D_3 + \dots$$

(thin lens approximation) distance between lenses  $\ll f$


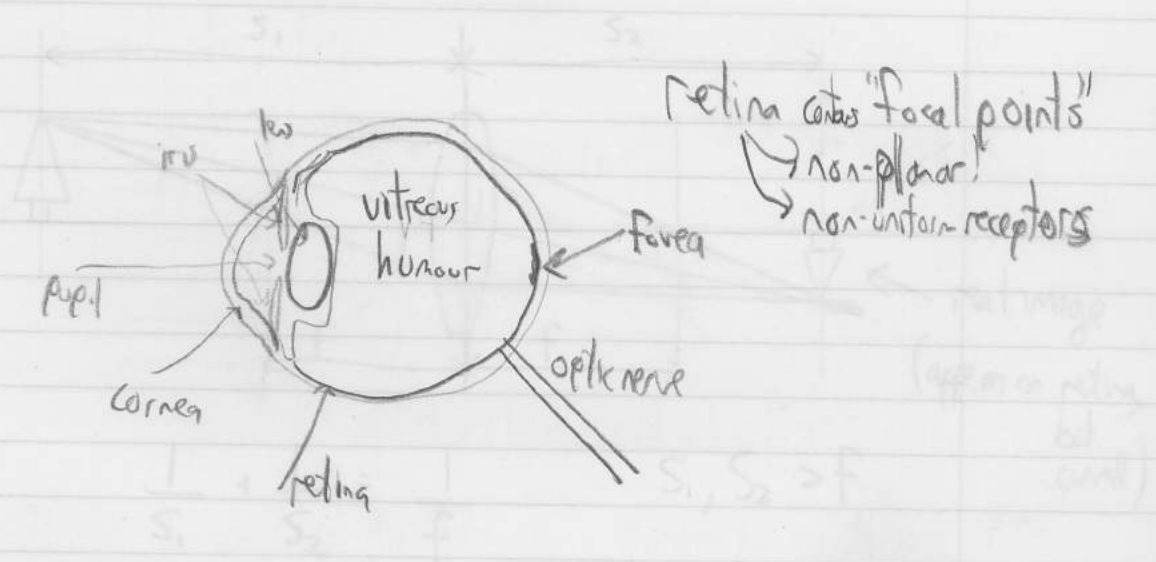
Prin p. 159

W.J. Smith

Modern Optical Engineering

(Optics must not be like algebra)


Structure of the human eye



- $n_1 = 1.009$  air
  - $n_2 = 1.376$  cornea
  - $n_3 = 1.336$  ocular fluid (aqueous)
  - $n_4 = 1.413$  lens
  - $n_5 = 1.337$  vitreous
- (48 is entry to cornea)

Total power of eye's optical system: 59.52D  
(16.8mm)

## Imaging properties of a lens




$$\frac{1}{s_1} + \frac{1}{s_2} = \frac{1}{f} \quad s_1, s_2 > f$$

If  $s_1 = \infty$ , then  $s_2 = f$   
(parallel light)

The human eye lens adjusts its  $f$  (or  $D$ ) to bring closer objects into focus


Could have  $s_2$  be negative: magnifying glass



Think about  
looking into  
RIFT  
(real image still  
produced at  
retina  
by eye's lens)

(13)

Image location of lens




$$\frac{1}{f} = \frac{1}{s} + \frac{1}{s'}$$

At  $s = f$ , the  $s' = \infty$  (parallel light)

The lens eye has object  $f$  (at  $D$ ) for lens eye


Along horizontal axis, this is like watching a stretched movie from  $\square$  to  $\square$



Hand-drawn notes on the left margin, including 'lens', 'object', and 'image'.

Lens Aberrations

1. Spherical aberration




Solution:  $\left( \begin{array}{l} \rightarrow \text{elliptical} \\ \leftarrow \text{hyperbolic} \end{array} \right)$  Harder to manufacture (aspheric lens)

Related: Look into edges of Rift lens and image might not be focused

2. Optical distortion


For high field of view lenses, the image is distorted



Mostly all purple light is a mixture of wavelengths/frequencies  
(Note: blue is not in spectrum)



Inverse of pinushion is barrel distortion; *distortion of wavelengths*



Think old CRT monitors

In RRT to compensate for pinushion distortion, apply a barrel distortion to rendered image.

- Now easy to do in graphics hardware (1990s had warped images)
- Apply in last stage of chain of transform




$$r \mapsto f(r)$$

(only a function of radius, or distance to image center)

### 3. Chromatic aberration

Light is composed of waves, varying between  $\sim 400\text{nm}$  and  $\sim 700\text{nm}$  wavelengths  
( $< 400\text{nm}$  is ultraviolet;  $> 700\text{nm}$  is infrared)

(Each wave has frequency  $f = \frac{c}{\lambda}$ )


The visible spectrum corresponds to pure waves (single sinusoid ):

Red, Orange, Yellow, Green, Blue, Indigo, Violet


Nearly all perceived light is a mixture of wavelengths/frequencies  
(Note: White is not in spectrum)

Eg. Blue Pearl path  $\leftarrow$  red head path

Spectral power - Like a histogram of wavelengths




Spectral reflectance



Perceiving color of an object! Source spectrum AND reflectance spectrum

Dispersion: Speed of light in medium depends on its wavelength!




Something happens in a lens



Eg. Blue focal length < red focal length

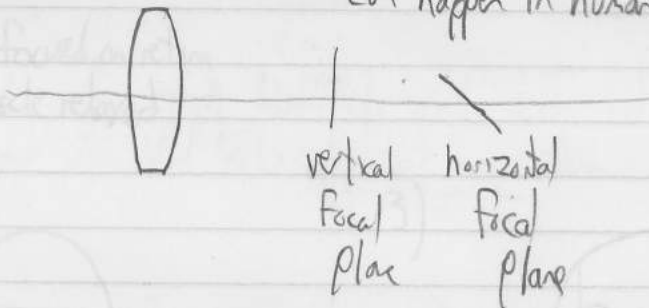
- Solution:
- 1) Find lens with high Abbe number (expensive)
  - 2) Make compound lens



- 3) Compensate in software by shifting based on pixel wavelengths


#### 4. Astigmatism

- Elliptical eccentricity in lens
- Can happen in human eye



Vertical waves vs horizontal waves / no common focus


#### 5. Coma - off axis distortion of image (like "comet")



How does the eye's lens/image system work?


Accommodation: Lens shape changes to increase diopter

Normal eye examples:




- Note that retinal image is not planar!  
- Eye lens is not circular

- Sharply focused on retina
- Eye muscle relaxed




- Blurry on retina
- Eye muscle relaxed



- Eye uses muscles to accommodate (refocus)
- Diopter increases

4) Reading glasses: Make close objects "seem like" infinite distance



- lens makes rays parallel
- No accommodation required

# "The Dress"

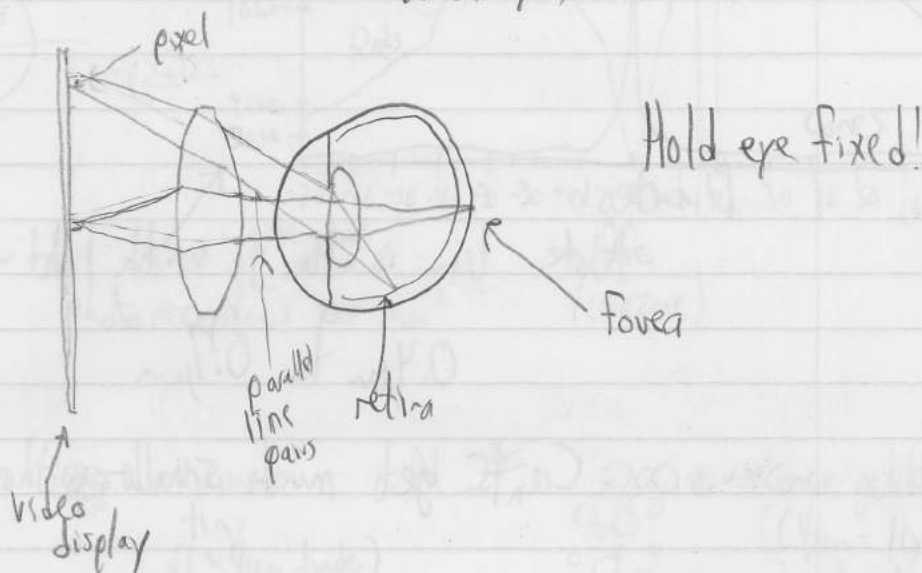
Blue & Black  
White & Gold

Farne! - Red & Yellow

News flash! The dress contained thermochromatic dye!  
When you take a picture, the color shifts on temp



A video rendering system and optical system put real image on retina (left eye)



The image hits a spherical surface of densely packed photoreceptors. Like pixels, but input instead of output

Two kinds! Rods and Cones

	Rods	Cones
Number (per eye)	120,000,000	6,000,000
Function	Low light intensity	Color sensitivity

Ray of light... 120,000,000... 11225  
 wavelength... vol of stimulation...  
 long wave... screen should be 16K x 16K  
 Any more is a waste!

(PP)

interocular space, blue & black, white & gold

Diameter

Note  $1\mu\text{m}$  is close to visible light wavelength;

$0.4\mu\text{m}$  to  $0.7\mu\text{m}$

Can't get much smaller without bleed over, interference


What about varying density?

Area of retina:  $1094\text{ mm}^2$  } → close to 200,000,000 photoreceptors

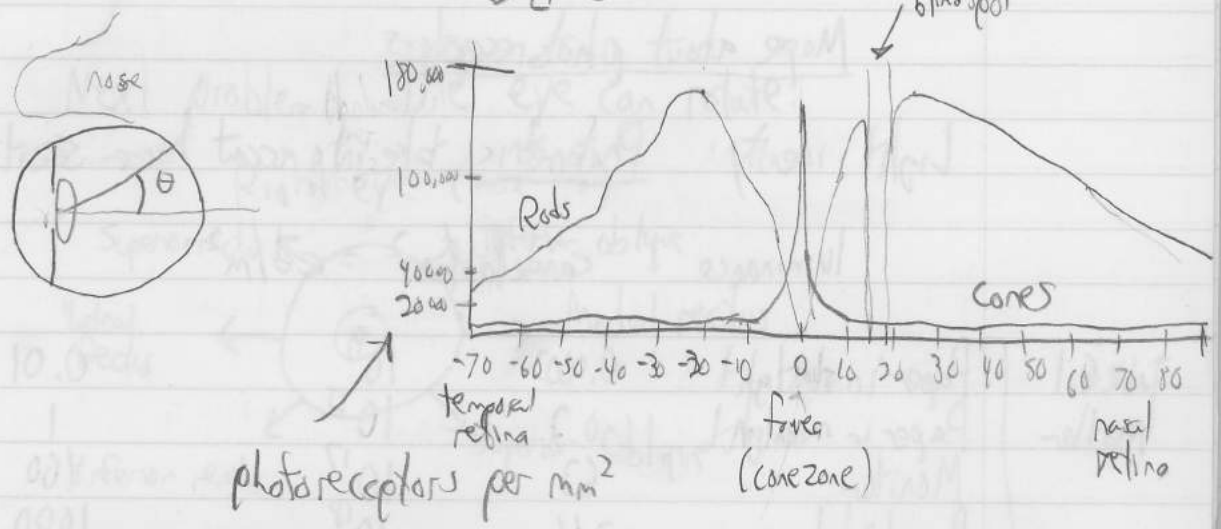
Density at fovea:  $\sim 200,000/\text{mm}^2$

$$\sqrt{200,000,000} \approx 14,142$$

If resolution too low



G. Osterberg  
Acta Ophthalmol. Suppl. 6:1  
1935



At 0°: all cones  
tiny  
(1 to 1µm diameter)


At 2°: cones get bigger  
(4µm - 6µm diameter)

At 50°: mostly rods  
rod ~ 1µm diameter

Implications: - Sharp color at retina; sharp monochrome at periphery  
- Retina/color requires high light intensity

How much resolution is enough for VR?

Rough estimate  $\sqrt{126,000,000} \approx 11225$   
total photoreceptors



Perhaps screen should be 16K x 16K  
Any more is a waste!



50b)

## More about photoreceptors

Light intensity: Photometric take into account human sensitivity

$$\text{luminance} = \text{candelas/m}^2 = \text{cd/m}^2$$

Table 6.1 Material			
Paper in starlight	0.0003	$10^{13}$	0.01
Paper in moonlight	0.3	$10^{15}$	1
Monitor	63	$10^{17}$	100
Room light	316	$10^{18}$	1000
Blue sky	2500	$10^{19}$	10,000
Paper in sunlight	40,000	$10^{20}$	100,000
	Luminous $\text{cd/m}^2$	Photons/ $\text{m}^2 \cdot \text{s} \cdot \text{m}^2$	Photons/receptor

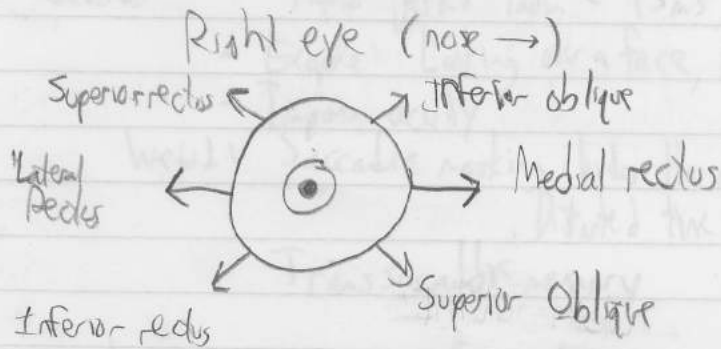
	<u>Scotopic vision</u>	<u>Photopic vision</u>
Photoreceptors	Rods	Cones
Light levels	$< 0.01 \text{ cd/m}^2$	$> 10 \text{ cd/m}^2$
Dark adaptation	35 min	10 min
Color	Monochromatic	Trichromatic

Scotopic: Dark light vision (night)

Photopic: Bright light vision (day)

- 90% of cells devoted to this
- we are "day time" animals

Next problem: The eye can rotate!



LM, MR: Eye yaw (up to 50 degrees), adduction (non-vergence part)

SR, IO, IR, SO: Eye pitch, avoiding roll

Types of eye movements

	Conjugate	Disjunctive
Voluntary	Saccade Pursuit	Convergence Divergence
Involuntary	Vestibulo-Ocular Optokinetic Microsaccades	

Note: Accommodation happens too, but not in VR at present!

(1/2)  
500)

IF VR tracking/rendering data is wrong, the world appears "swimmy"

Or real world can appear swimmy - take off glasses

- 90% of off-axis light  
- 40% by the lens

Movements

- 1. Saccades - rapid "jerks" lastin < 45ms; 900°/s
- Example: Looking over a face, reading
- Improves acuity

Weird: Saccadic masking hides the motion intervals from brain.  
Distorted time perception - second hands on clock

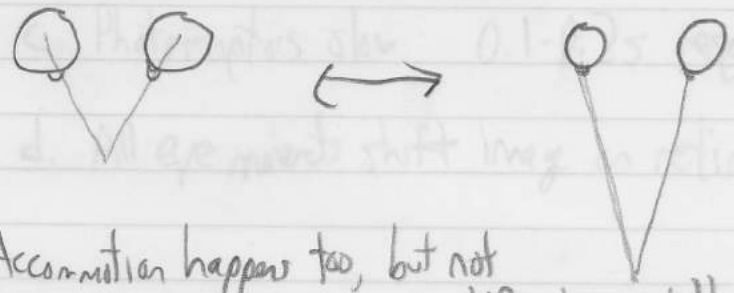
Transsaccadic memory

- 2. Smooth pursuit
  - Track a moving visual signal
  - Example: Moving tennis ball, car
  - < 30°/s, otherwise saccades added
  - Reduces motion blur (slow photoreceptors)

- 3. Vestibulo-ocular reflex
  - Bypasses brain ( $\approx 10ms$ , fastest reflex in body)
  - Keep image stability
  - Counter rotations for 6DOF head motion
  - VOR Gain adaptation

- 4. Optokinetic
  - Alternating between pursuit and saccade
  - Example: Watch a passing train

- 5. Convergence/Divergence - Stereo vision



Note: Accommodation happens too, but not in VR at present!

Process

- 1. Research - early '90s, late '90s, 2000s
- Early: looking on a face, reading
- Japan only

Work: describe working like the other industry for design  
 think the program - read hand  
 Transactable necessary

2. Small groups


- Track a visual signal
- Experimental design: 1000s of trials
- 2000 - others looking at work
- Review with the UI - design groups

3. Vertical - other work

- sample for what program (e.g. 1000 trials)
- low level work
- (e.g. 1000 trials)
- 1000 trials

4. Cite results

Result: very hard to design VR system display that is "equivalent" to



Note: Accurately happen for, but not in VR at all


### 6. Microsaccades

- 1/30 to 2 degrees motion
- Purpose: largely debated, unknown
- Might cause illusory motion

↳ recall slide earlier...

Regarding VR and eye motion:

1. Not looking through optical axis (even if eyeball centered)



Problem: focus, distortion, maybe nausea

2. VOR gain adaptation - swaying effect in VR and real world

3. Complicated interaction between display & photoreceptors

- a. Pixels switch color, intensity at same rate  
(20ms LCD, 100ms OLED)
- b. Frames might be off (black), Frame rates 60, 75, 90 FPS
- c. Photoreceptors slow: 0.1-0.2s response time
- d. All eye movements shift image on retina!

4. Vergence


4. Vergence usually coupled with accommodation  
- Opposites But not in VR → always in focus

- Switch to slides on visual pathways
- Based on Mather Ch 7

Show

- Retina

- Cross section



ON bipolar → cone and rod → connect to amacrine cells  
 OFF bipolar → cone only → connect to ganglion

Horizontal cells: Lateral inhibition

Pathways: Ganglion → LGN/Thalamus part → visual cortex

Make alien circuit joke

Cratellin Tearing

Visual Cortex

End slides

- Receptive field model → funneling, compression
- Opponency → important idea!
- ↳ spectral vs spatial

Ganglia Cells

	Midget 70-80%	Parasol 8-10%	Bistrifurcated <10%
Spectral response	Photopic opponent red-green (center) yellow (peripheral)	P & Scotopic opponent yellow (rod-cone)	Photopic nonopponent blue-yellow
temporal	sustained	transient	

ON-center



OFF-center



Ganglion image

Path to visual cortex

Orientalion Tuning

Visual Cortex

End slides



# Depth Perception Matter Ch10

5. Depth cues:
- 1) Metric - vary continuously with distance
  - 2) Ordinal - ordering: near to far

Multiplicity of depth cues — Not just stereo  
(think about Google StreetView)


1. Retinal image size




2. Height in visual field



10. Stereo cues



3. Texture gradients




(LG explains ones TV)

(52)


Depth Perception

Depth cues - very commonly all distance  
1) Pictorial - original - original  
2) Original - original - original

Show -  
Slides!  
(Dir 15)



Texture gradient

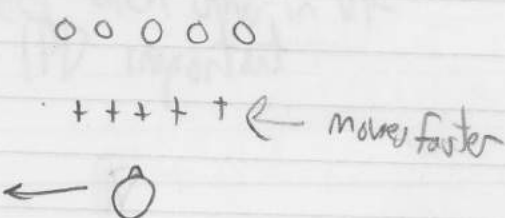


4. Image blur  
Blur parts of image to appear outside of focal plane/depth

5. Atmospheric perspective  
Hazy mountains in distance

6. Accommodation  
For close objects and people under 50

7. Motion parallax  
Optical flow




8. Shadows/shading

9. Interposition

10. Stereo cues

- Vergence angle



- Binocular disparity → left and right eye images are shifted

- Diplopia - multiple images

- Think of IPD in VR!

- Heteropter (LG explains curved TV)

Combination of Cues

Think Bayesian

Scale perception vs depth perception

- Reach and grab something
- They affect each other in VR
- Context, IPD important


Motion Perception

- Purposes:
- Segmentation/Segregation  
Extract moving body from background
  - Extract 3D structure from motion  
(Spin the chair around)
  - Visual guidance of action
    - Manipulation - grab a cup
    - Hand-eye coordination

Object moves, observer and eye fixed.

Notes: Had to do because of south point

## Neural circuitry for motion




The brain uses more information to distinguish  
 - Saccades  
 - Eye movement signals  
 - Eye-scan motion

neuron activated for:  
 $(A \text{ timeshifted}) \times (B \text{ current})$   
 "was at A, now at B"

Possible to confuse with wagon-wheel effect (video)

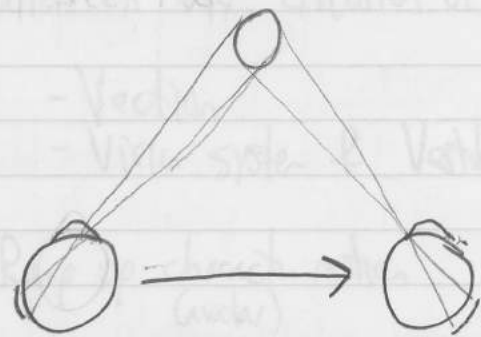
## Object Motion vs Observer Motion

Retinal image moves same way in either case



Object moves, observer and eye fixed

Note: Hard to do because of smooth pursuit



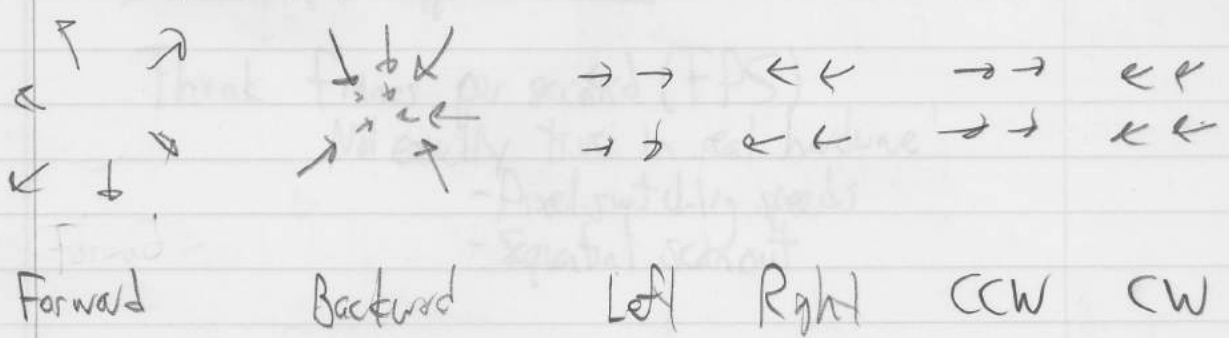
The brain uses more information to distinguish

- Saccadic masking/suppression - suppresses motion detectors
- Eye movement commands
- Large-scale motion - If eye moves, the whole scene moves (Recall big swing illusion)

### Optical Flow

- Track movement of features on retina (or image plane or screen)
- A vector field on image plane (or sphere)

### Self Motions




(Walk around room)

Big problem in VR: Illusion of self motion from optical flow

- vection
- Vision system & Vestibular system in disagreement

Ex: Ramp up character motion (avatar)



Which is more comfortable?

Small acceleration mismatch over long time

vs

Large accel'n mismatch over short time ← better!

Stroboscopic apparent motion

Think frames per second (FPS)

Not exactly true in real hardware!


- Pixel switching speeds
- Sequential scanout

(1)

with loop - if not on floor to ceiling

see Wikipedia - "Frame Rate" -  
"Flicker-Fusion threshold"

Ex: Range up through window  
(window)



What is more comfortable?

2nd window monitor on top line

1st window monitor on top line ←

Stroboscopic effect

That frame rate (FPS)


not really the real picture

- Pixel width

- Pixel height



One ms enough for a frame to register  
Up to 10 FPS, images perceived individually



Early silent films: 16-24 FPS (hand-tuned speed)  
Home movies: 16-18 FPS  
Motion picture industry: 24 FPS standard

- Two-blade projector → 48 FPS effective
  - Three-blade → 72 FPS effective
- Thomas Edison said minimum for comfort is 46  
New motion picture standards proposed: 48 (James Cameron wants)

CRT refresh rate → large monitors, peripheral vision

- 60 Hz - noticeable flicker
- 72 Hz - minimum organic reconstruction
- 85 Hz - comfort for all

Note: Detectable vs not, but headaches, fatigue vs not, and comfortable




↓ very general!

Current talk in VR industry:

- 60 Hz vs 75 Hz vs 90 Hz vs 120 Hz
- Gear VR vs DK2 vs Oculus/Vive vs Sony?


- Google for "blur busters"
  - Abrash blogs
- Everything better than 60Hz

Works because photoreceptors need only 1ms of exposure, and hold for around 100ms



Problems with displays

Most scanout image line by line



Motivated by original CRT - electron gun - analog  
Need to keep phosphors illuminating

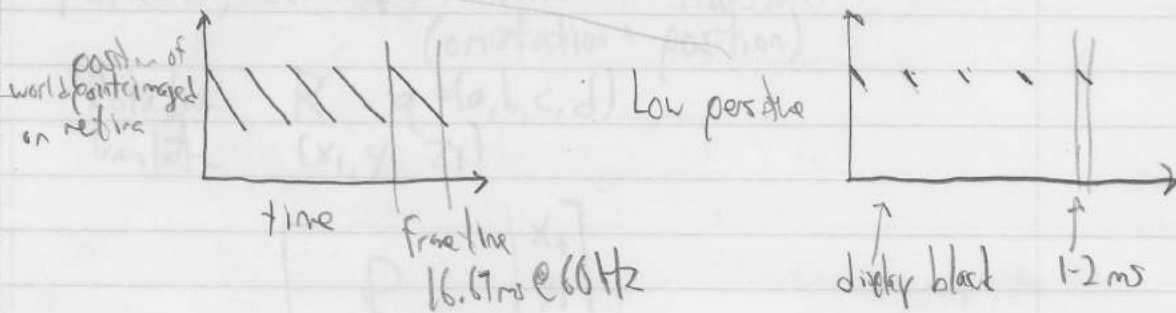
This causes a "waggle" when screen object moves



But not translating across retina in most cases

Recall smooth pursuit and VOR



Both supposed to keep object fixed on retina



Low persistence leads to flicker fusion problem; thus 90Hz or higher

(3)

Google for "the butter" ...  
- About 1/2 ...

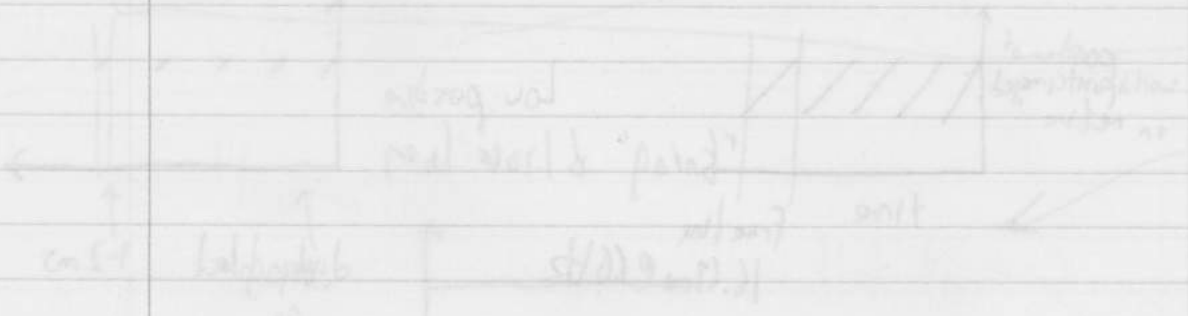



Maintained by ... (RT - electron gun - ...)  
Keep to keep ...

The cross is "paddle" ...



But not ...



Low ... leads to ...

Another problem: LCDs have slow pixel-switching speed  
 result: Blur - see DK1

---

## Tracking Systems

What do we want to track?

Think about rigid bodies

- 1) Head wearing HMD
- 2) Palms of hands
- 3) Eyes
- 4) Fingers
- 5) Entire body
- 6) Movable objects - controller, coffee cup, -
- 7) Other people in the space

For each, estimate rotation + translation  
 (orientation + position)

Rotation  $R$ ,  $q = (a, b, c, d)$


Translate  $(x_t, y_t, z_t)$

$$\begin{bmatrix} R & x_t \\ \hline 0 & 0 & 0 & 1 \end{bmatrix}$$

[Continue in next notebook]

Estimating orientation (Oculus Rift DK1, <sup>Samsung</sup> Gear VR)  
 ↳ Also DK2, but + position

2D case



Merry-go-round

$$\omega = \frac{d\theta}{dt}$$

At  $t=0$ ,  $\theta = \theta_0$

At  $t > 0$ , rotation rate is  $\omega$  rad/s

Thus,  $\theta(t) = \theta_0 + \omega t$

What if  $\omega$  varies over time?

$$\theta(t) = \theta_0 + \int_0^t \omega(s) ds$$

Discrete-time approximation:  $\Delta t = 1 \text{ ms}$   $\omega_i = \omega((i-1)\Delta t)$

$$\theta = \theta_0 + \sum_{i=1}^k \Delta\theta_i \approx \theta_0 + \sum_{i=1}^k \omega_i \Delta t = \hat{\theta} \quad (\text{Euler integration})$$

Now suppose a sensor estimates  $\omega_i$  every  $\Delta t$

$$\hat{\theta} = \theta_0 + \sum_{i=1}^k \hat{\omega}_i \Delta t$$

estimated rotation

sensor reading from 1D gyroscope


[Continue in next notebook]

(82)

Blogger Beau Cronin

Oculus BPG: "O BPG may be the most substantial thing ever written on applied sensorimotor neuroscience"

$$\frac{\partial b}{\partial t} = \omega$$



$\theta = \theta + \omega \Delta t$   
 $\theta = \theta + \omega \Delta t$   
 $\theta = \theta + \omega \Delta t$

What if  $\omega$  varies over time?

$$\theta(t) = \theta_0 + \int \omega(t) dt$$

Discrete-time approximation:  $\Delta t = \text{int}$   $\omega = \omega(t) \Delta t$

$$\theta = \theta_0 + \sum_{i=1}^k \omega_i \Delta t = \theta_0 + \sum_{i=1}^k \omega_i \Delta t$$

Now solve a least squares problem w/ every  $\Delta t$

$$\theta = \theta_0 + \sum_{i=1}^k \omega_i \Delta t = \theta_0 + \sum_{i=1}^k \omega_i \Delta t$$

[Continue in next notebook]   
 sensorimotor for ID dynamics   
 estimator

## Oculus Best Practices Guide

- Goals:
1. Oculomotor comfort - avoid eye strain
  2. Bodily comfort - prevent feelings of disorientation, nausea
  3. Positive user experience - fun, immersive, engaging
  4. Minimal aftereffects - visual-motor functioning after use

### Rendering

- Keep 1:1 correspondence between physical and virtual world head rotation & translation
- Don't make fixed splash screen or images w/o tracking
- R & L eye image differ only by viewpoint
- Supersampling, anti-aliasing - worst at center pixels

### Latency

- Target ~20ms for sensor fusion reading to rendering
- Keep latency constant as possible
- Use predictive tracking

### Tracking & viewpoint

- Never let tracking thread starve
- Don't make large rotating line (esp. horizon)

### Positional tracking

- Mindful of users poking their heads anywhere in camera range
- Better to fade out when limit (e.g. wall) reached



Accelerations

- vection  
problems
- Make mismatches (physical or virtual) as short as possible
  - Accelerations could be constant speeds! Move along curve
  - Have user control acceleration

Movement speed

- Real world 1.4 m/s walking
- Teleporting OK, but try to preserve orientation
- Don't make them look one way while moving another

Cameras

- Don't do zoom effect - bad optical flow/acceleration
- No head bobbing

Managing SimSickness

- Find fresh users, you are not good judge of your own
- Note that people vary in their tolerance
- Make adjustments or options to vary intensity of problems

Degree of Stereoscopic Depth

- Don't make things too close
- Don't make virtual camera R/L distance too large → Fatigue!

User Interfaces

- Make UI sit 2-3 meters in front of user
- Keep in middle 1/3 of screen
- Caution with head movement to select menu
- Embed info into world objects
- Draw crosshairs at object distance

Avatar Control

- Acceleration issues
- Avoid keyboard options if possible
- Head movement for control → Dumpy the elephant

Sound

- Difference between headphones & speakers (6DOF transform)
- NPC (non-player character) speech should not be "center channel"
- Keep position (+ orientation) in mind

Contact


- Units: 1 meter in physical world = 1 meter in Unity
- Most objects between 0.75 - 3.5 m away
- Consider rendering avatar body, but beware of mismatch (eg sitting vs standing)
- Use monocular depth cues
- Avoid strafing, backstepping, spinning

Health & Safety

- Avoid 1-30Hz Flicker → photoepileptic seizures
- Avoid high contrast, high spatial frequency gratings →

Look up:  
Pokemon seizure  
episode


3D orientation case



Axis-angle  $(\vec{v}, \theta)$

3 axis gyroscope measures

$$\vec{\omega} = (\omega_x, \omega_y, \omega_z)$$



If  $\omega$  is constant over  $\Delta t$ , then rotation is

$$\text{by } \Delta\theta = \|\vec{\omega}\| \Delta t = (\omega_x^2 + \omega_y^2 + \omega_z^2)^{1/2} \Delta t$$

$$\text{and axis } \vec{v} = \frac{1}{\|\vec{\omega}\|} (\omega_x, \omega_y, \omega_z)$$

during for  $\Delta t$ .

Let  $q_i$  be quaternion representation of  $(\vec{v}_i, \Delta\theta_i)$

$$q_i = \text{Quat}(\vec{v}_i, \Delta\theta_i)$$

(P2)

Recursively:

$$2D: \hat{\theta}_{\text{current}} = \Delta \hat{\theta}_k + \hat{\theta}_{\text{previous}} = \hat{\omega}_k \Delta t + \hat{\theta}_{\text{previous}}$$

$$3D: \hat{q}_{\text{current}} = \Delta \hat{q}_k \circ \hat{q}_{\text{previous}} = \text{Quat}(\hat{v}_k, \Delta \hat{\theta}_k) \circ \hat{q}_{\text{previous}}$$

Integrating sensor readings to estimate orientation

Recall 2D: 
$$\hat{\theta} = \theta_0 + \sum_{i=1}^k \Delta \hat{\theta}_i = \theta_0 + \Delta \hat{\theta}_1 + \Delta \hat{\theta}_2 + \dots + \Delta \hat{\theta}_k$$

$$= \Delta \hat{\theta}_k + \Delta \hat{\theta}_{k-1} + \dots + \Delta \hat{\theta}_2 + \Delta \hat{\theta}_1 + \theta_0$$

$$\Delta \hat{\theta}_i = \hat{\omega}_i \Delta t$$

Now 3D: 
$$\hat{q} = \Delta \hat{q}_k \circ \Delta \hat{q}_{k-1} \circ \dots \circ \Delta \hat{q}_2 \circ \Delta \hat{q}_1 \circ q_0$$

← match ordering  
← not commutative!


Each  $\Delta \hat{q}_i = \text{Quat}(\hat{v}_i, \Delta \hat{\theta}_i)$

$\hat{v}_i \rightarrow \frac{\hat{\omega}_i}{\|\hat{\omega}_i\|}$   
 $\Delta \hat{\theta}_i \rightarrow \|\hat{\omega}_i\| \Delta t$

Problem: Drift errors

2D:  $d_k = \theta_k - \hat{\theta}_k$

3D:  $d_k = q_k \circ \hat{q}_k^{-1}$



The error is a 3D rotation, with yaw, pitch, roll components

(10)

Integrating error signals to estimate orientation

Recursively:

$$\hat{\theta}_k = \hat{\theta}_{k-1} + \Delta \theta_k$$

$$\hat{\theta}_k = \hat{\theta}_{k-1} + \Delta \theta_k + \Delta \theta_{drift}$$

probabilistic

...

...

not

$$\hat{\theta}_k = \hat{\theta}_{k-1} + \Delta \theta_k + \Delta \theta_{drift}$$


! sensitive

$$\Delta \theta_{drift} = \Delta \theta_{drift} + \Delta \theta_{drift}$$

Note: Drift corrections can come from sensors at very low frequency.



Problem: Drift errors



$$\hat{\theta}_k = \hat{\theta}_{k-1} + \Delta \theta_k + \Delta \theta_{drift}$$



The error is a combination of the drift error and the sensor error

Let pitch + roll error be called tilt error.


Correcting for drift errors:

- 1) Use another sensor to provide world reference
- 2) Gradually apply corrections
  - Fast enough to fully compensate
  - Slow enough to avoid sim. sickness

For tilt error: Need a gravity or "up" sensor

For yaw error: Need a "compass"

Use accelerometer to correct tilt error.



Coordinate frame is estimated orientation

$$\hat{a} = (\hat{a}_x, \hat{a}_y, \hat{a}_z)$$

tilt axis in XZ plane

$$\vec{v}_{\text{tilt}} = (-\hat{a}_z, 0, \hat{a}_x)$$

Need to rotate by  $\phi$  about  $\vec{v}_{\text{tilt}}$  to fix.

Let pitch + roll error be called tilt error

2D case

$$\hat{\theta}_{corrected} = \hat{\theta} + \alpha \phi$$

error measured in  $\theta$

Note: Drift in tilt error of accelerometer is slow enough to avoid tilt drift



Let  $\phi$  be the tilt error



Simple way to correct: Complementary filter

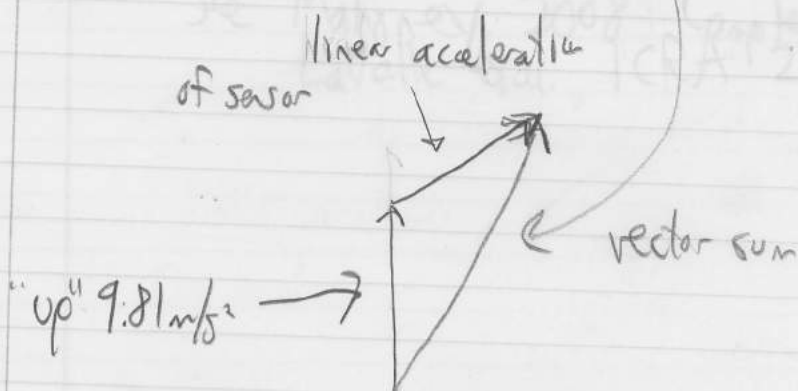
Gain coefficient  $\alpha > 0$   $\alpha \ll 1$  (example  $\alpha = 0.0001$ )

In each  $\Delta t$ ,  $\hat{q}$  changes every  $\Delta t$

$$\hat{q}_{\text{corrected}} = \text{Quat}(\vec{V}_{\text{tilt}}, \alpha \phi) \circ \hat{q}$$

$\alpha$  needs to be large enough to cancel error, but small enough to avoid sim. sickness and perceived jitter.

Problem: Accelerometer measures



Use heuristics to detect when "not moving"

Example:  $\|\hat{a}\| \approx 9.81$

Can make mistakes.  
Accelerate head downward

Use magnetometer to correct yaw error:

Similar to tilt correction:

- Calculate reference error
- Gradually apply using complementary filter

Problems:

- Vector sum of Earth's field, building field, board field
- Calibration hard
- Field might vary over time and position
- Soft iron bias

Accuracy  $\approx$  5 degrees


See Mahoney, 2008: Complementary filters on  $SO(3)$   
 LaValle et al., ICRA 2014

Estimating position + orientation (Rift DK2, Sony Morpheus, Valve Vive)

- IMU (gyro + accel) not enough
  - Drift is too fast from double-integrating  $\ddot{a}$
  - No way to detect drift error

Solutions: 1) Generate non-constant magnetic or EM field


- Razer Hydra, STEM Sixaxe
- UWB Radio



2) Visibility or line of sight methods

Camera arrangements

On headset





inside-out

In world



outside-in


Pinhole camera  
Perspective projection



features in an image

- Features:
- 1) Natural - hard computer vision
    - Extract & maintain from natural scenes
    - Remove moving objects
    - Reliability low
  - 2) Artificial - trivial computer vision
    - Blob detection
    - QR tags, retroreflective markers, LEDs, laser projector
    - Can stay in IR spectrum (invisible to humans)

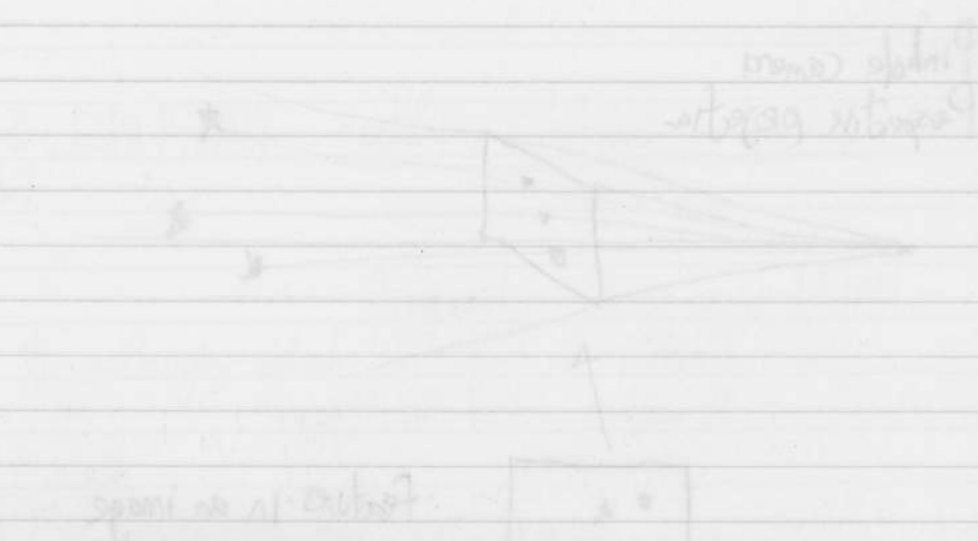
Retroreflective features




MoCap systems  
Vicon  
OptiTrack

power  $\sim \frac{1}{d^2}$  and  $d$  is doubled

12)



Determine position & orientation of triangle from features in image



DOF analysis

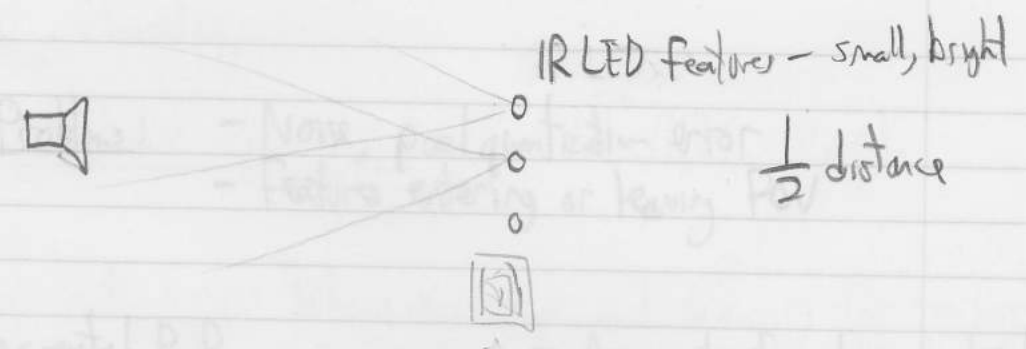
- Start with 6 DOFs (rigid body)
- Each feature subtracts 2 DOFs (xy coords in image)

PIP: Object can rotate freely (3) + near-to-far (1) = 4


P2P: Object can swing around edge between two features (1) and twist (pitch) while changing distance (1) = 2

P3P: No freedom left

Illustrate with real object



PnP problem: Determine rigid-body transform from identified, observed features on a rigid body.  
 perspective points → camera to body




A system of polynomial equations can be obtained.  
 Generally, 8 solutions, but only 4 in front of camera  
 (See Wu, Hu, 2006 PnP Problem Revisited)

P4P, P5P, ... PnP  
 → multiple solutions, coplanarity causes difficulty

(2)

Horizontal velocity

$$\frac{1}{2} g t^2$$



Determine the horizontal distance traveled by the projectile.  
 The horizontal distance is given by  $v_x t$ .  
 The vertical distance is given by  $\frac{1}{2} g t^2$ .  
 The total distance is given by  $\sqrt{(v_x t)^2 + (\frac{1}{2} g t^2)^2}$ .


$$\hat{x}_{cor} = \hat{x} + \alpha_1 \Delta t + \frac{1}{2} \hat{a} \Delta t^2 + \hat{v} \Delta t$$

$$\hat{v}_{cor} = \hat{v} + \alpha_2 \Delta t + \hat{a} \Delta t$$

Illustrate with real object

- Problems:
- Noise, pixel quantization error
  - Features entering or leaving FOV

### Incremental PnP



- Assume transform known already
- Cannot travel far in  $< 20ms$
- Perturb rotation & translation to match image
- Optimization, randomized gradient descent
- Accelerometer helps improve starting point for optimization

### Sensor Fusion:

- Camera can drift-correct orientation
- Accelerometer can help with position estimation
- Use complementary filter (2nd order)
  - ↳ Equivalent to Kalman filter for this problem (Higgins, 1975)
- Double integrator dynamical system model
  - $\ddot{x} = u$       $u = \text{observed linear acceleration}$

↳  $\alpha$  coefficients for both position and velocity




(57)

Problem 1 - Noise and quantization error  
- Feature extraction or learning FOV

- Avoid features from already  
- Can't find features in FOV

(Should redraw to  
make relative  $\ll w$ )



fractional part

- Acceleration takes input starting  
- Part for definition

$$\dot{x}_{acc} = \dot{x} + \alpha \Delta x + \frac{1}{2} \alpha^2 \Delta x^2 + \dots$$
$$\dot{v} = \dot{v} + \alpha \Delta v + \frac{1}{2} \alpha^2 \Delta v^2 + \dots$$

- Cannot do drift control operation  
- Acceleration can help with position estimation  
- UK complementary filter (Shahmoradian)


- Double integrator dynamical system model  
-  $\ddot{x} = u$  - control input

is constant for both position and velocity

Pharology

Lighthouse approach (Valve/HTC Vive)

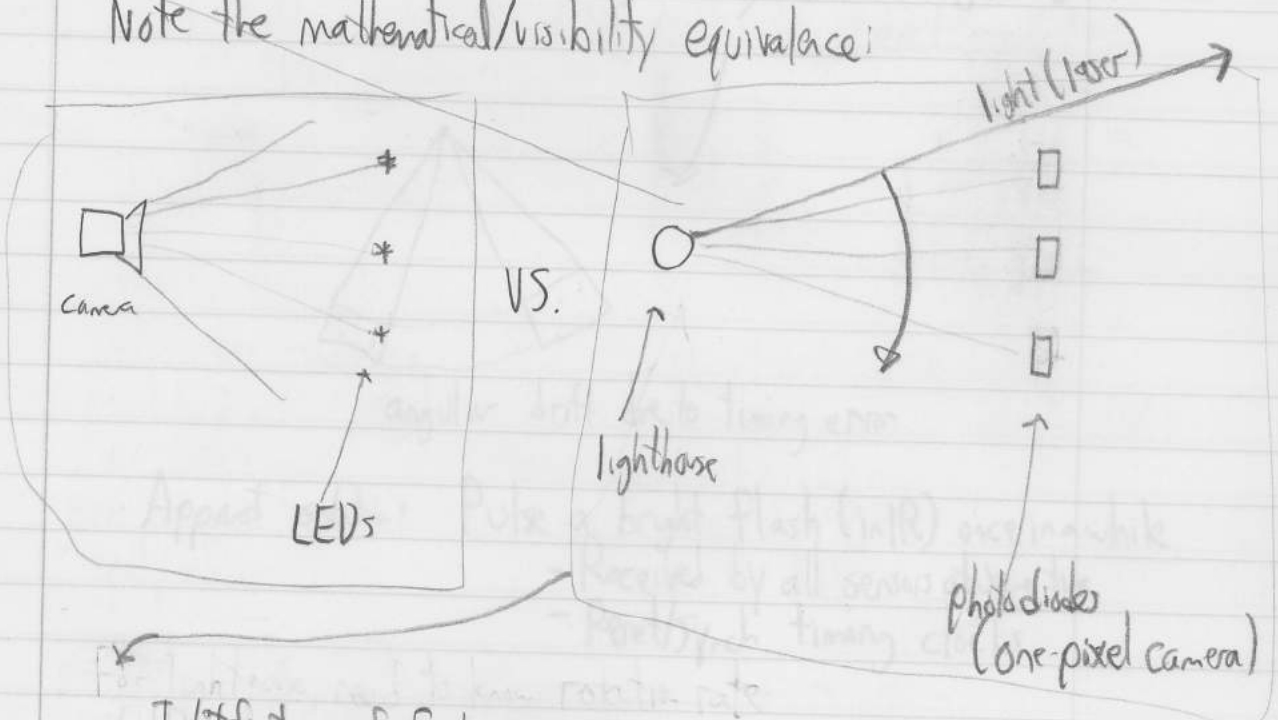
Why not bypass taking an image and scanning for features?  
(blob detection)



- Computationally costly (could do in hardware)
- Limited light per pixel


More resolution not necessarily better

Note the mathematical/visibility equivalence:




- Identification of features is easy
- Need to know rotation rate (100Hz?)
- No image processing
- Little power dissipation because laser
- Does have moving parts

Need horizontal and vertical beams



Problems:

- In Two beams hit same sensor at same time
- Maintaining timing to avoid angular drift



angular drift due to timing error

Apparent solution: Pulse a bright flash (in IR) once in a while

- Received by all sensors at same time
- Reset/synch timing clocks

Rendering

Shirley et al., Ch 4, 8

Recall chain of transformations:

Transform point in body frame to world to screenspace  
(virtual) (pixel location)

$$T = T_{vp} T_{can} T_{eye} T_{rb}$$

↑  
T<sub>e</sub>, T<sub>r</sub>

Next problem: How to set the RGB values of each pixel?


Two approaches:

image-order rendering  
pixel-by-pixel  
ray tracing  
easy to implement from scratch  
usually slower

object-order rendering  
triangle-by-triangle (or primitive)  
rasterization  
used by GPUs, fast  
usually faster


Simplest: Lambertian - illumination proportional to cosine of

$$L = k_d I \cos(\theta, n)$$



### Ray tracing stages


- 1) ray generation
- 2) ray intersection
- 3) shading



#1 is a ray of perspective projection

#2 must find first intersection from focal point

Intersection tests




#3 Shading

Simplest: Lambertian - illumination proportional to cosine of

\*

Unit vectors!




$$L = k_d I \max(0, n \cdot l)$$

↑ pixel color  
 ↑ surface color  
 ↑ light intensity

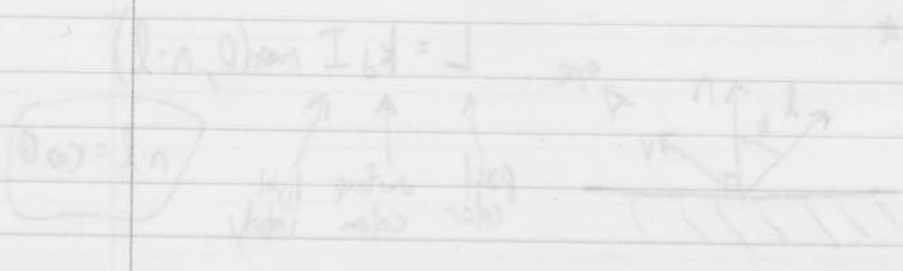

$$n \cdot l = \cos \theta$$

18/18

When is light visible in mirror?




Normal = normal  
(Law of Reflection)



diffuse, flat, matte surface BRDF

Handling highlights: Blinn-Phong shading

Unit  
vectors



bisector:

$$h = \frac{v+l}{\|v+l\|}$$

$$L = k_d I \max(0, n \cdot l) + k_s I \max(0, n \cdot h)^p$$

diffuse                      specular

(p. 83, Shirley)  
 $p = 10$  eggshell  
 $100$  mildly shiny  
 $1000$  really glossy  
 $10000$  nearly mirror like

What about ambient light?

Add  $k_a I_a$  to  $L$

What about multiple light sources?

Add diffuse and specular for each one.


What about multiple reflections?

Need to Global illumination — very expensive  
 (visibility with mirrors)

Transparency? Recall absorption, refraction, reflection picture  
 from optics

More general model for shading: BRDF

Bidirectional Reflectance Distribution Function



$$f(\theta_i, \phi_i, \theta_r, \phi_r) = \frac{\text{radiance} \leftarrow \text{reflected energy}}{\text{irradiance} \leftarrow \text{received energy}}$$

Relative symmetry  $\rightarrow f(\theta_i, \theta_r, \phi_i - \phi_r)$  - for physics consistency

Lambertian: Special case with  $f$  constant

Problem with VR: Stereo eye views



Need to do twice the work for shiny surfaces



More general model for shading: BRDF  
Bidirectional Reflectance Distribution Function



radiance  $L(x, y, z, \omega)$   
radiance  $L(x, y, z, \omega')$

Useful adaptations

Z-buffer from light source perspective can generate shadow map (show parts of object obscured by light)

Problems with VR: Stereo eye view




Need to be able to track the work for each eye

Rasterization

Most common - GPUs

Object-order rendering - triangle-by-triangle



Fill in pixels by sampling midpoint

Problem: Aliasing

Depth ordering problem: Which triangle is in front?

Painter's algorithm: Sort by furthest to closest; paint far to near order

Problem:  depth cycle - segment triangles

Problem: Sorting too costly

Z-buffer: Store depth value ( $z$ ) at each pixel.  
 Render triangles in arbitrary order.  
 Render pixel only if new  $z$  value is less than old one at that pixel.

Clipping: Remove triangles (or parts of triangles) that are behind the eye, or too close (vertical clipping plane)

~~VR~~ problem, where?


Efficiency: Culling - Eliminate unnecessary rasterizations

Backface: Remove triangles on "back" of object  
(not in view)

View frustum: Remove triangles not in viewing frustum

Propagating color, RGB, other attributes efficiently:

Barycentric coordinates



$$p = \alpha_1 p_1 + \alpha_2 p_2 + \alpha_3 p_3 \quad \text{for } 0 \leq \alpha_1, \alpha_2, \alpha_3 \leq 1 \text{ and } \alpha_1 + \alpha_2 + \alpha_3 = 1$$

( $\alpha_i = 0 \rightarrow$  on edge  $\alpha_i, \alpha_j = 0 \rightarrow$  on vertex)

Efficient interpolation

$$R = \alpha_1 R_1 + \alpha_2 R_2 + \alpha_3 R_3$$


$$G = \alpha_1 G_1 + \alpha_2 G_2 + \alpha_3 G_3$$

$$B = \alpha_1 B_1 + \alpha_2 B_2 + \alpha_3 B_3$$

$$\begin{array}{ccc} \uparrow & \uparrow & \uparrow \\ atp_1 & atp_2 & atp_3 \end{array}$$

Efficiency - Cutting - Eliminate unnecessary features  
Backlog: more things on back of queue

Regular Patterns pose a other resolution problem.  
Wavefront to ensure phase issues & boundaries



$p = x, y, \text{ or } z, \text{ for } 0 \leq x, y, z \leq 1$  and  $x+y+z=1$

$(x, y, z) = 0 \rightarrow \text{or } \text{or } \text{or}$

Efficient interpolation

$R = x, y, z, \text{ or } p_1, p_2, p_3$

$G = \alpha, \beta, \gamma, \text{ or } p_1, p_2, p_3$


$B = \alpha, \beta, \gamma, \text{ or } p_1, p_2, p_3$

↑ ↑ ↑

$\alpha, \beta, \gamma$


Other mappings over triangles:

Texture mapping (like painting a picture onto surface)



- Look up mipmapping for scaling textures (a form of anti-aliasing)

Bump mapping!



Map a pattern of normals to surface  
Make it look rougher

Normal mapping!

Use surface normals from high-resolution surface



→ Make plane look spherical (for example)

Use Lambertian shading to make surface look curved.

28

Other mappings are triangles

Texture mapping (the mapping of texture on surface)

Main differences:

- 1) Head motion
- 2) Stereo
- 3) Low resolution

- Lack of mipmapping for scaling texture (a form of antialiasing)

Point mapping

Map a feature of ramp to surface  
Make it look complex



Normal mapping: Use surface normals for high-resolution

→ Make plane look spherical (for example)



Use Laplacian smoothing to make surface look curved


Problems with rendering for VR:

- Shading:
- Highlights (specularities) need stereo perspective → double work
  - Texture maps look like painted cardboard
  - Bump/normal maps look fake


- Aliasing:
- Resolutions look low

Head: Motion turns "staircase" into "escalator"

- Stereo causes mismatched stairs per edge



- Render target:
- Unusual shape
  - Use stencil buffer in GPU



- Correct for optical distortion due to lenses  
pincushion - barrel annihilation

Latency: - GPU pipeline optimized for triangle throughput, not latency.

↑  
triangles/sec

- Latency compensation by post-rendering image warp (time warp)

Visibility problems for 6DOF case

Geometry errors: - Thin objects look fake or implausible (leaves in Tuscany)

- Holes, isolated points more obvious (think about SLAM data)

- Chromatic aberration?

- head in walls

(you can move your head into/through obstacles)



Antialiasing - Get rid of "jaggies"

Supersampling



(low discrepancy)

Average the colors for the final pixel value

Common special case (hack): MSAA - multisample antialiasing

Only supersample depth (z) and stencil values (in or out of triangle)

Compute polynomial barrel distortion


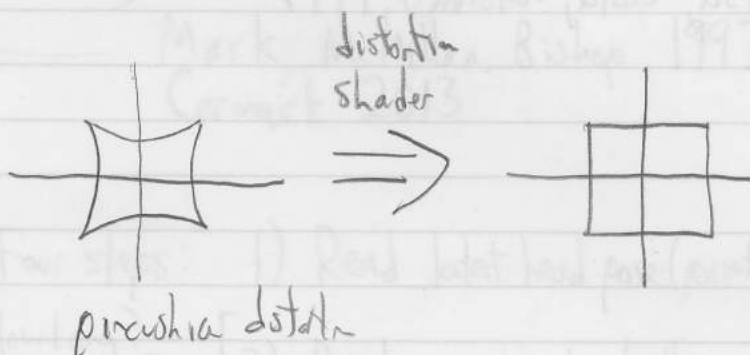
$$f(r) = 1 + c_1 r^2 + c_2 r^4 + c_3 r^6 \quad (1 - c_4 r^8) \dots$$

Odd powers not needed (theory of optics)

$c_1, c_2, c_3$  chosen by hand or laser measured

Apply  $f(r)$  to every pixel to barrel warp image  
 Precision - barrel annihilation

## Distortion Shading for High POV:



- Use polar coordinates from optical center
- Assume rotation invariance for distortion  
( $\theta$  does not matter)

Radially symmetry

$$(r, \theta) \mapsto (f(r), \theta)$$

Compute polynomial barrel distortion

$$f(r) = 1 + c_1 r^2 + c_2 r^4 + c_3 r^6 (+ c_4 r^8 + \dots)$$

Odd powers not needed (theory of optics)

$c_1, c_2, c_3$  chosen by hand or laser measurements




Apply  $f(r)$  to every pixel to barrel warp image  
Pincushion - barrel annihilation!

tune until grid lines straight

(op)

Also, Darsa, Costa, Varshney, 1997



- Use polar coordinates for off-axis  
 - Assume constant variance for distance  
 (theta not matter)

transform spatially

$$(r, \theta) \rightarrow (r', \theta')$$

constant polynomial point distribution

$$f(r) = 1 + c_1 r^2 + c_2 r^4 + c_3 r^6 + \dots$$

Odd powers not needed (theory of optics)

$c_1, c_2, c_3$  chosen by hand or least squares


edges are sharp

Apply  $f(r)$  to every pixel to correct image  
 Bicubic - 8-bit convolution


Post-rendering image warp (time warp)

Mark, McMillan, Bishop 1997  
Carmack 2013

- Four steps: 1) Read latest head pose (orientation & position)  
How long? 10-50ms? [2) Render scene into buffer using pose for viewpoint  
3) Read latest head pose  
4) Adjust rendered output to fake the newer viewpoint



pitch upward  
occurred in  
step 2  $\Rightarrow$



no data!

Image shifting accounts for small pitch and yaws  
(vertical) (horizontal)

Image rotation accounts for small roll.



What about position?


Some filtering possible

Some faking possible:


- Rescaling image to account for depth ( $z$ ) change
- Shift image to account for  $x, y$  change

Problems!

- Perspective projection not quite right as  $z$  changes (small change OK)
- Visibility problems



Another use: Improve frame rate by inserting faked viewpoints




Note! We are not doing lenses, images


(25)

Some factors possible:


- Decreasing number of particles (or change in volume)
- Shift temp. to increase kinetic energy




- Velocity problems
- Pressure problems in physics or 2 change (small change OK)



Another use: change from one state by varying forces




## Audio for VR



Sound waves: Similar to light except:


- Fluctuating air pressure (instead of EM/photons)  $2 \times 10^4$  Hz
- Frequency only 20 Hz to 20000 Hz (instead of  $400-700$  nm wavelength)  
(2000 Hz? in basilar membrane)
- Speed 343 m/s in air (not  $c = 3 \times 10^8$  m/s)  
(wavelengths:)

Interacting with other media: Just like light




Note! We are not doing lenses, images

Frequency spectrum: Just like light




pure tone



ROXGBIV  
like in visible spectrum



complex wave



like mixed colors (white)

Think filtering, Fourier transforms

Human auditory system

(show slides for pictures - dur 24)

Explain components: pinna, tympanic membrane, cochlea, Organ of Corti, hair cells, --



Also, higher above noise threshold




# Perception of Sound Weber's law (JND)

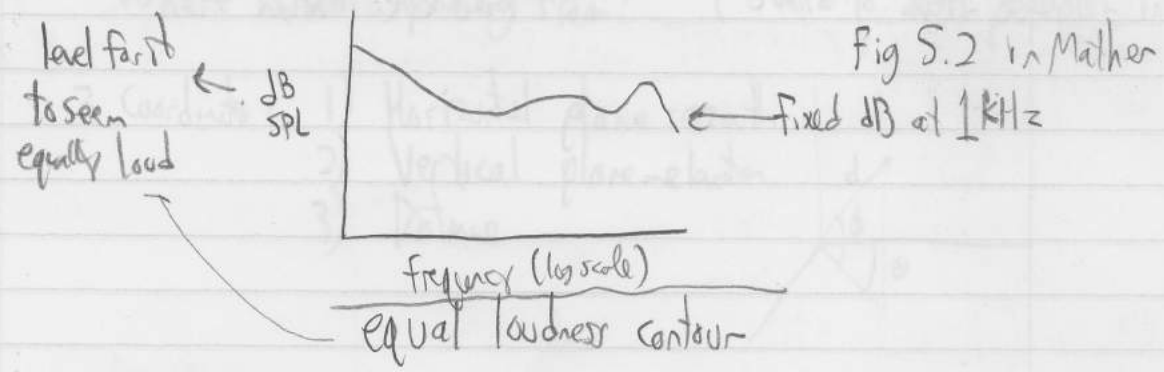
Recall Stevens' power law:  $m \mapsto cm^x$

$m$  = magnitude of stimulus  
 $p$  = perceived magnitude


Loudness at 3000 Hz:  $x \approx 0.67$



## Loudness as function of frequency



## Loudness adaptation



Also, loudness above noise threshold

Perception of sound


Recall store power law:  $m \rightarrow cm^x$

No. of neurons  
to perceive stimulus

Lawson of 30000's in cortex




-  $\phi$  and  $\theta$  are known for vision  
because of position on retina,  
and eye rotation



(sharp) period  
- natural resonance

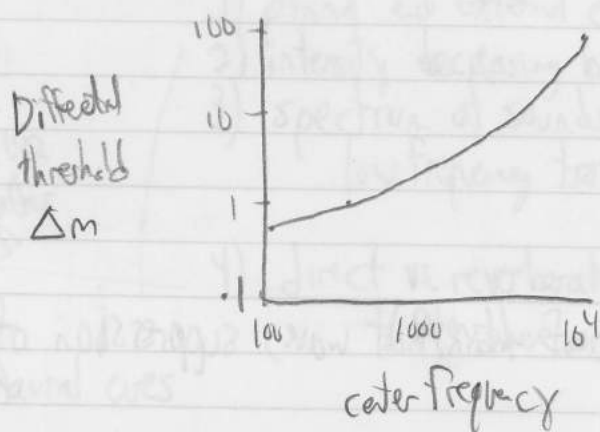
Resonance frequency



Resonance  
of cortex  
low


All neurons are not firing

## Pitch perception (recall Weber law and JND)

Auditory localization

Where is sound coming from? (Similar to depth perception in vision)

- 3 Coordinates:
- 1) Horizontal plane - azimuth
  - 2) Vertical plane - elevation
  - 3) Distance




## Minimum audible angle (MAA) — JND



- $\sim 1^\circ$  below 1000Hz, straight ahead
- $\sim 5^\circ$  to the side
- terrible around 1500-2000Hz on side

(2P)

Pitch perception (real Weber law of JND)




Audio illusion!

Preference of main (first wave) suppression of later wave (reverberated)



vs.



only clear one

both heard

Minimum audible angle (MAA) - JND

- 1° below 1000 Hz  
 - 2° in the ear  
 - terrible sound 100-1000 Hz or ear



Monaural cues

p. 142  
Matter

- 1) pinna and external ear canal shape
- 2) intensity decreasing by inverse square law
- 3) spectrum of sounds

low frequency travels further! Distal thunder is only low frequency rumble

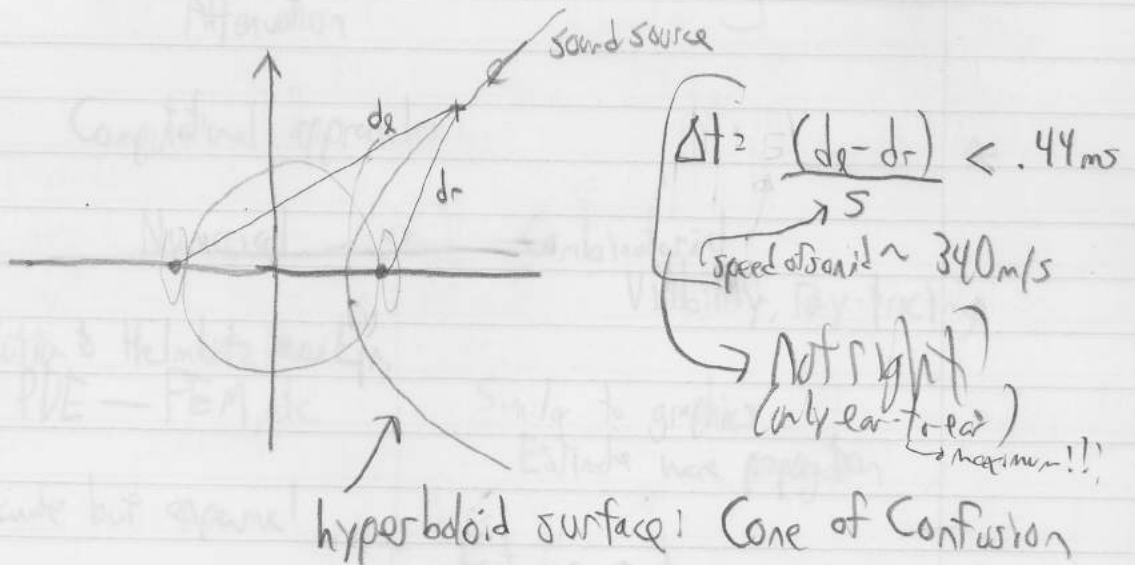
- 4) direct vs. reverberation energy (When there are reflecting surfaces)

Binaural cues

ILD: Inter-aural Level Differences - acoustic shadow of ear ("back face")

ITD: Inter-aural Time Differences - arrival time, distance between ears ~14cm

In sensing, Multilateration from TDOA - time difference of arrival (DECCA system, WWII) modern GPS





## Sound rendering

Four steps

### 1. Modeling

Geometric models - walls, obstacles  
 Acoustic material properties - absorption, reflection, refraction  
 Sound sources - point source, parallel wave source, loudness

Lower resolution than for graphics (not like pixels)  
 Simplified

### 2. Propagation

Reflection - specular, scattering  
 Diffraction  
 Refraction  
 Doppler effect  
 Attenuation

Yes, there is an equivalent to BRDF for audio

### Computational approaches

Numerical	vs	Combinatorial
Solution to Helmholtz wave Eqn. PDE - FEM, etc.		Visibility, ray-tracing
Accurate but expensive		Similar to graphics Estimate wave propagation
		Fast, approximate

Sound rendering  
UNC talk about that  
Modeling

Acoustic material properties - absorption, reflection, refraction  
Sound sources - point source, line source, surface source  
Lower frequencies - the graphics (not the sound)  
Propagation (time)

Kind of like applying barrel distortion  
at the end for visual rendering



Computational approaches

Numerical

Finite Element Method (FEM)

Acoustic boundary conditions

Fast algorithms

Sampling & graphics -  
Fast algorithms

Numerical  
Variability of frequency



### 3. Rendering

Use head positional orientation to determine pressure signal at right and left ears



Problem: Account for scattering due to:

- Outer ears (pinna)
- Head
- Whole body (clothing?)

HRTF (Head-Related Transfer Function)

resolves CONE of confusion!

$$H(\omega) = \frac{\text{Output}(\omega)}{\text{Input}(\omega)}$$

Adjust for scattering using linear transfer function (in frequency  $\omega$  domain)


- Depends on individual ear, head, body shape
- Depends on x, y, z of audio source
- Can estimate by recording in anechoic chamber - various impulse responses

distance

$$H(\omega, \theta, \phi, d)$$

or

$$H(\omega, \theta, \phi) \text{ - farfield approx}$$



- Quantize angles, distance
- Could assure "far field" and distance is "at infinity"

## 4. Display

Especially one "pixel" display per ear

Time is more important than space  
Frequency, phase, amplitudes

Surround sound (fixed in space) vs. headphones vs. earbuds (eliminates pinna altogether)

What about head position & orientation?  
Like CAVE

HRTF should account

## Problems for developers:

1. How much modeling needed? How to accomplish?
  - Need good middleware
2. How to evaluate correctness or accuracy?
  - Ears not as sensitive as eyes (and eyes already play tricks)
  - HRTFs are person-dependent - like eyeglasses
3. Computational cost

# 3D User Interfaces Bowman et al., 2005

←  
2. Depth

Include vehicles, grappling, jumping high

←

## Visualization of oneself:

- Torso
- Arms, legs, hands
- Face in mirror

←  
3. CAVE

How much walking needed for 3D visualization?  
 - Head gear  
 - How to evaluate correctness of 3D  
 - How to evaluate accuracy  
 - HSTP or known objects - the experience  
 - Cognitive cost

## Interfaces for VR

Some categories:


1. Locomotion - Moving oneself from place to place  
 Issues: Speed, vestibular mismatch, getting lost
2. Manipulation - Touching, feeling, grabbing, carrying, discarding objects  
 Issues: Selection, force feedback, grasping realism, depth perception
3. System Control - Menus, windows, text  
 Issues: Readability, comfort, speed
4. Social - chat, texting,  
 Issues: face covered, body language, spatial arrangement, uncanny valley
5. Others?? Text editing? Street fighting? Cinematography?

Simulating ANY interface that exists in the physical world!

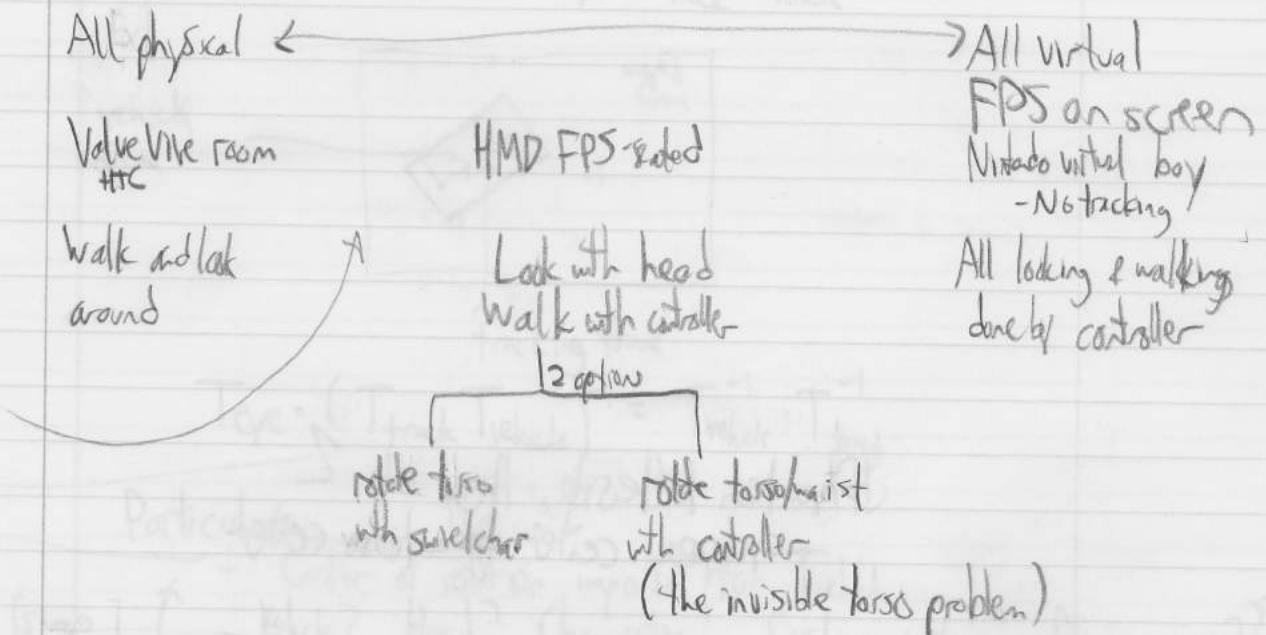
3D User Interfaces  
Bowman et al., 2005

- Stationary bicycle
- Omnidirectional treadmill

Bruder, Lubos, Steinicke, 2015  
Minimum Radius: 22m



Locomotion - Traveling in VR



How to apply it:

Recall chain of transforms:  $T_{up} T_{can} T_{eye} T_{rb}$

Consider how  $T_{eye}$  is formed from "look at" to get  $\hat{x}, \hat{y}, \hat{z}, e$

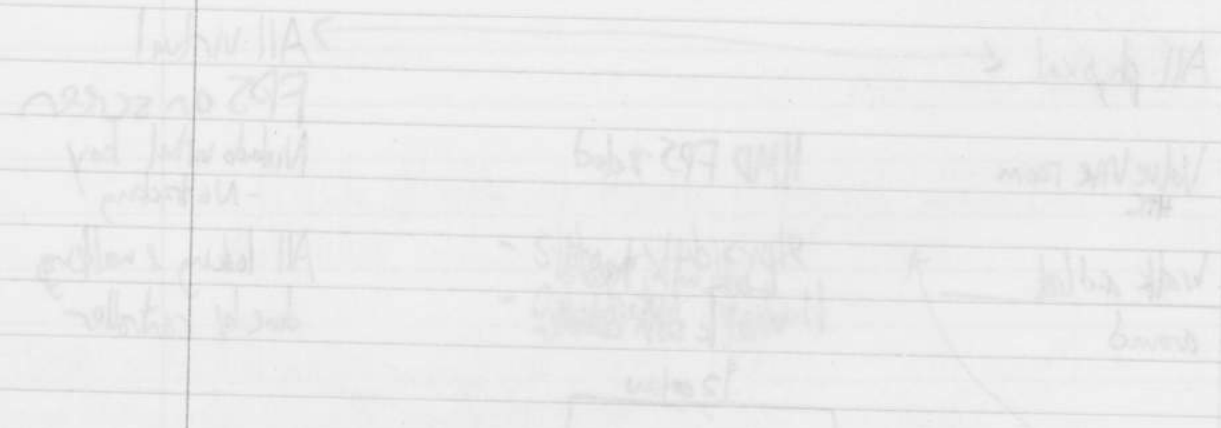
All physical:  $T_{eye}$  entirely determined by head tracking  
 All virtual:  $T_{eye}$  " " " Controller

Cool trick for all physical walking in unbounded virtual worlds:  
 Redirected walking (Razaqre, Kohn, Whittan, 2001)

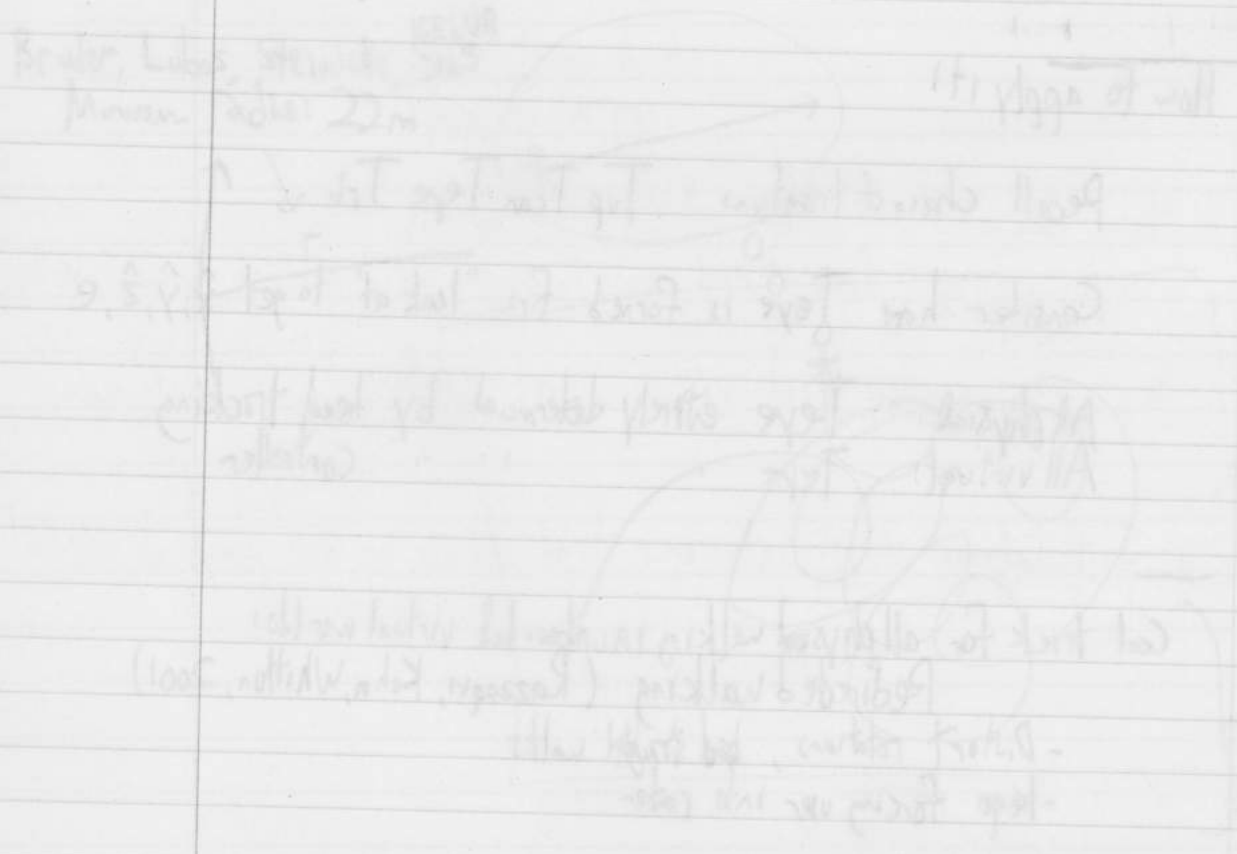
- Distort rotations, but straight walks
- Keep forcing user into center

(05)

Location - location in IR




If you lean to the side, then cyclopedal center  $\neq$  tarsal center



For HMD FPS:

Lookat frame transferred by  $T_{track} T_{vehicle}$



$$T_{eye} = (T_{track} T_{vehicle})^{-1} = T_{vehicle}^{-1} T_{track}^{-1}$$

Particulars:

- Center of rotation important for  $T_{vehicle}$   
Vehicle? Head? Chair center? Cyclopean retina? Player mass??
- For virtual walking in a fixed chair,  $T_{vehicle}$  is equivalent to 3D rigid body transform  
Provides x, z offset, and yaw direction (torso)
- For free-flying ship,  $T_{vehicle}$  is full 3D rigid body transform

#2 is more comfortable because the brain spots the best, large mismatch as either

Idea: Do the same for rotation (windtanks)

Discrete jumps for orientation changes

Too large jumps → Confusion about new direction



Problems with vestibular mismatch (vection)


- Linear acceleration in VR causes perceived acceleration from vision!



Note: Arrows and the accel. vectors

Vestibular organ knows there is no motion

Two strategies for initiating a walk



#1

#2

#2 is more comfortable because the brain rejects the brief, large mismatch as outlier

Idea: Do the same for rotations (windlands)


Discrete jumps for orientation change

Too large jumps → Confusion about new direction


- Angular motion, even when constant, induces vection



Circular motion



$$\vec{a} = \omega \times \vec{v} = \omega \times (\omega \times \vec{r})$$



Beware!

Further problem: Even motion at constant velocity (straight line, constant speed) is uncomfortable. Why?

Is continuous motion important?

Depends on application - usually not important!

- Teleporting - great for large areas; possible location confusion
- VR cinema - pick your seat from seating chart

Use a map or W/M (World In Miniature)


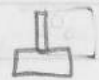
Issue: Wayfinding? - Learning a virtual city or forest  
 - Visual landmarks important - like real life!

ManipulationFactors

- 1) Selection: Distance and direction to object  
Size  
Clutter around object  
Occlusion (maybe partial)
- 2) Positioning: Distance and direction of initial and goal  
Positioning precision requirement  
Does goal have obvious basin of attraction?
- 3) Rotation: Amount and precision of rotation.  
All 3 DOFs?

Input device categories:

- 1) Metric - Motions are tracked through space; consider DOFs  
Desk mouse, air mouse (gyro, 3D), Leap Motion, gloves
- 2) Switch - Motions induced by pressing, pushing  
Joystick, buttons

Even Atari 2600 (1977) had both: Metric was paddle  and  
Switch was joystick 

Which one is artificial steering wheel?



Could be either, but

Manipulation

1) Selection: Nature of order to object  
See  
Clifford and Spill  
Coffman (order control)

2) Positioning: Nature of order to object  
Positioning  
Positioning  
Positioning

See Bowman et al  
book



1) Order - Nature of order to object  
2) Order - Nature of order to object  
3) Order - Nature of order to object

4) Order - Nature of order to object  
5) Order - Nature of order to object



Can be either

Note that both have scaling issues over space and/or time

- How far to move mouse?
- How long to hold button?
- How much to rotate in a tap of joystick?

- Perceptual issue! Learnable motor program

- Precision issue! Can a high precision task be accomplished?

Like press  on window corner w/mouse

- Note! No vestibular mismatch with scaling! For manipulation!

Selection idea! Do ray casting to point at objects - like laser pointer

Manipulation idea! Make a virtual hand on an extensible arm

Final advice! Reduce DOFs as much as possible in task

General issue! Haptic feedback

Rumble

Pressure/force

Gloves

Manipulator robot hand

3D printed objects

Important application! Medical simulation or teleurgery

See Oculus BPG and Bownen et al.



Flow to the next  
- the high level of  
- the way to create a set of objects

- Perception/Action / Learning with action  
- Perception: how can a high-level action be accomplished?  
- The next level is to learn from action  
- Note: No real-time learning with action for the moment

What is the goal? To learn to control the objects - like a robot

Perception: how to learn a model for an object's action

First action: how to learn a model for an object's action

Control: how to learn a model for an object's action  
Example  
Action/force

Go here? Or somewhere else?




Important applications: learning to control an object or robot

## System Control

Sending commands to the application

Most common: Graphical menu  
(Others: Voice command, gesture command)



How to select? Head looking, controller, -- the clicking game

### Guidelines:

- Embed menu in virtual world - like a billboard  
(Do not attach to force in front of eyes)
- Lazily pulling menu into looking direction is OK
- Distance in virtual world should be 2 to 3 meters away
- Fit inside middle 1/3 of viewing area - avoid head movements to read
- Embed info into the world - ammo count appears on gun

### What about text entry?

Keyboard difficult  
 Split keyboard with separation  
 (Glove) Chorded keyboard

Social Interaction

Virtual communities - sociology

↳ fr. book by Howard Rheingold  
 Includes Plato, chatrooms, email lists, Usenet, MUDs  
 multi user dungeon

Special case:

Virtual worlds


Science fiction: Snowcrash, Ready Player One, Neuromancer!

MMORPGs: Massively multiplayer online roleplaying games  
 teamwork!

Second Life (Open Simulator)

An entire economy formed - Linden Dollars  
 GDP \$64 million in 2009 - like Bitcoins  
 No government - damage, that's hard to achieve

Face-to-face interaction:



↑  
 Allows fantasy, altercations

↑  
 Approaches reality




## Achieving realism


- Use omnidirectional video camera
- Is person in real or placed into virtual world? Extracting problem
- What if an HMD occludes their face? Camera inside?
- Effect of delays - at least synchronize audio+video

## With avatars, achieving the impossible (Transformed Social Interaction)

- Change gender, "face", species
- Change height, weight, eye color, hair color
- Empathy can be taught!!  
Police training videos
- Look at multiple people simultaneously (Bailenson et al 04)



Teacher in physical classroom



Teacher in virtual classroom

In VR the teacher can look at every student simultaneously because they each get a unique world on their HMD

Also, if students are not looking at you, then their avatar fades away

Bringing realism back into the avatar!

- Eye tracking (blinks, winks?)  
(More latency allowed than for foveated rendering)
- Face modeling, animating, and tracking  
see Paul Debevec, USC
- Beware the uncanny valley!  
People get creeped out
- Hand and body tracking  
Poor tracking worse than none at all, (eg  
Leap motion, Kinect can be poor  
Can communicate body language)

Final comment on interfaces! Any interface from physical world can be simulated in VR.

Hence, a VR interface can be simulated in VR!  
And on and on...

(Put on a virtual HMD while already in VR)

Like the levels in Inception