Virtual Reality & Physically-Based Simulation
Stereopsis, VR Display Technologies, Stereo Rendering

G. Zachmann
University of Bremen, Germany
cgvr.cs.uni-bremen.de

Which Depth Cues do You Know?

https://www.menti.com/hht2ei5jxt

Depth Cues (*Not* Sorted by Importance)

- Motion parallax: apparent motion of objects relative to each other, when observer moves
- Occlusion (see CG1)
- Stereopsis (binocular/stereo vision)
- Accommodation & convergence
- Defocus blur (a.k.a. blur gradient)
- Perspective foreshortening (see CG1)
- Lighting & shading (see CG1)
- Relative size / familiar size
- Texture gradient

Binocular/Stereoscopic Vision (aka. Stereopsis)

- Stereopsis = "vision with two eyes"
  - The mechanism in human vision for *sensing* depth
- Convergence (a.k.a. vergence) = counter-rotating eye movement (around the vertical axis), so that the optical axes of the eyes intersect at some point (fixation point)
  - So that the fixated object appears on the center of the retina (has highest resolution)
- Focus (a.k.a. accomodation) = adjustment of the eyes' lenses to adapt to different distances
  - So that the fixated object appears sharp on the retina

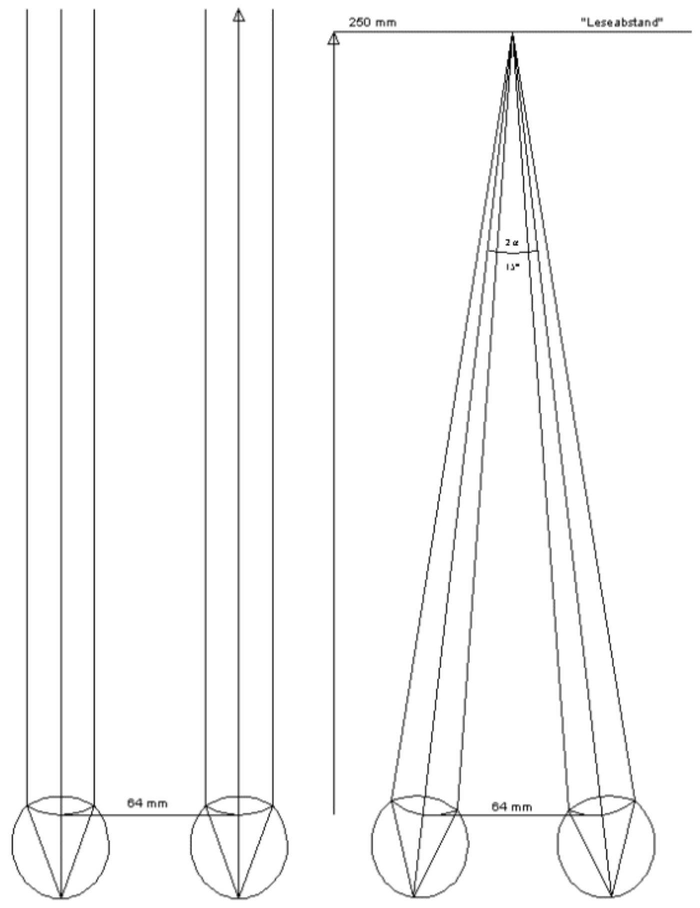

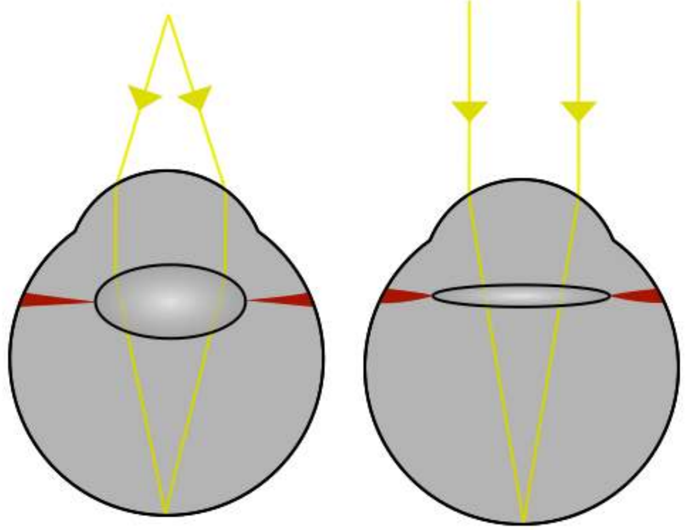

(Fun) Factoids about Stereopsis

- Stereo blindness: affects ~10% of general population
- Some people can actually turn their eyes to *divergence*:

Disparity and Horopter

- Convergence on one object causes disparity δ between corresponding points on the retinas for other objects:

$$\delta = \delta_2 - \delta_1 \;= \gamma - \alpha$$

- Horopter = locus of points in space with same apparent depth as the fixated object = points with 0-disparity

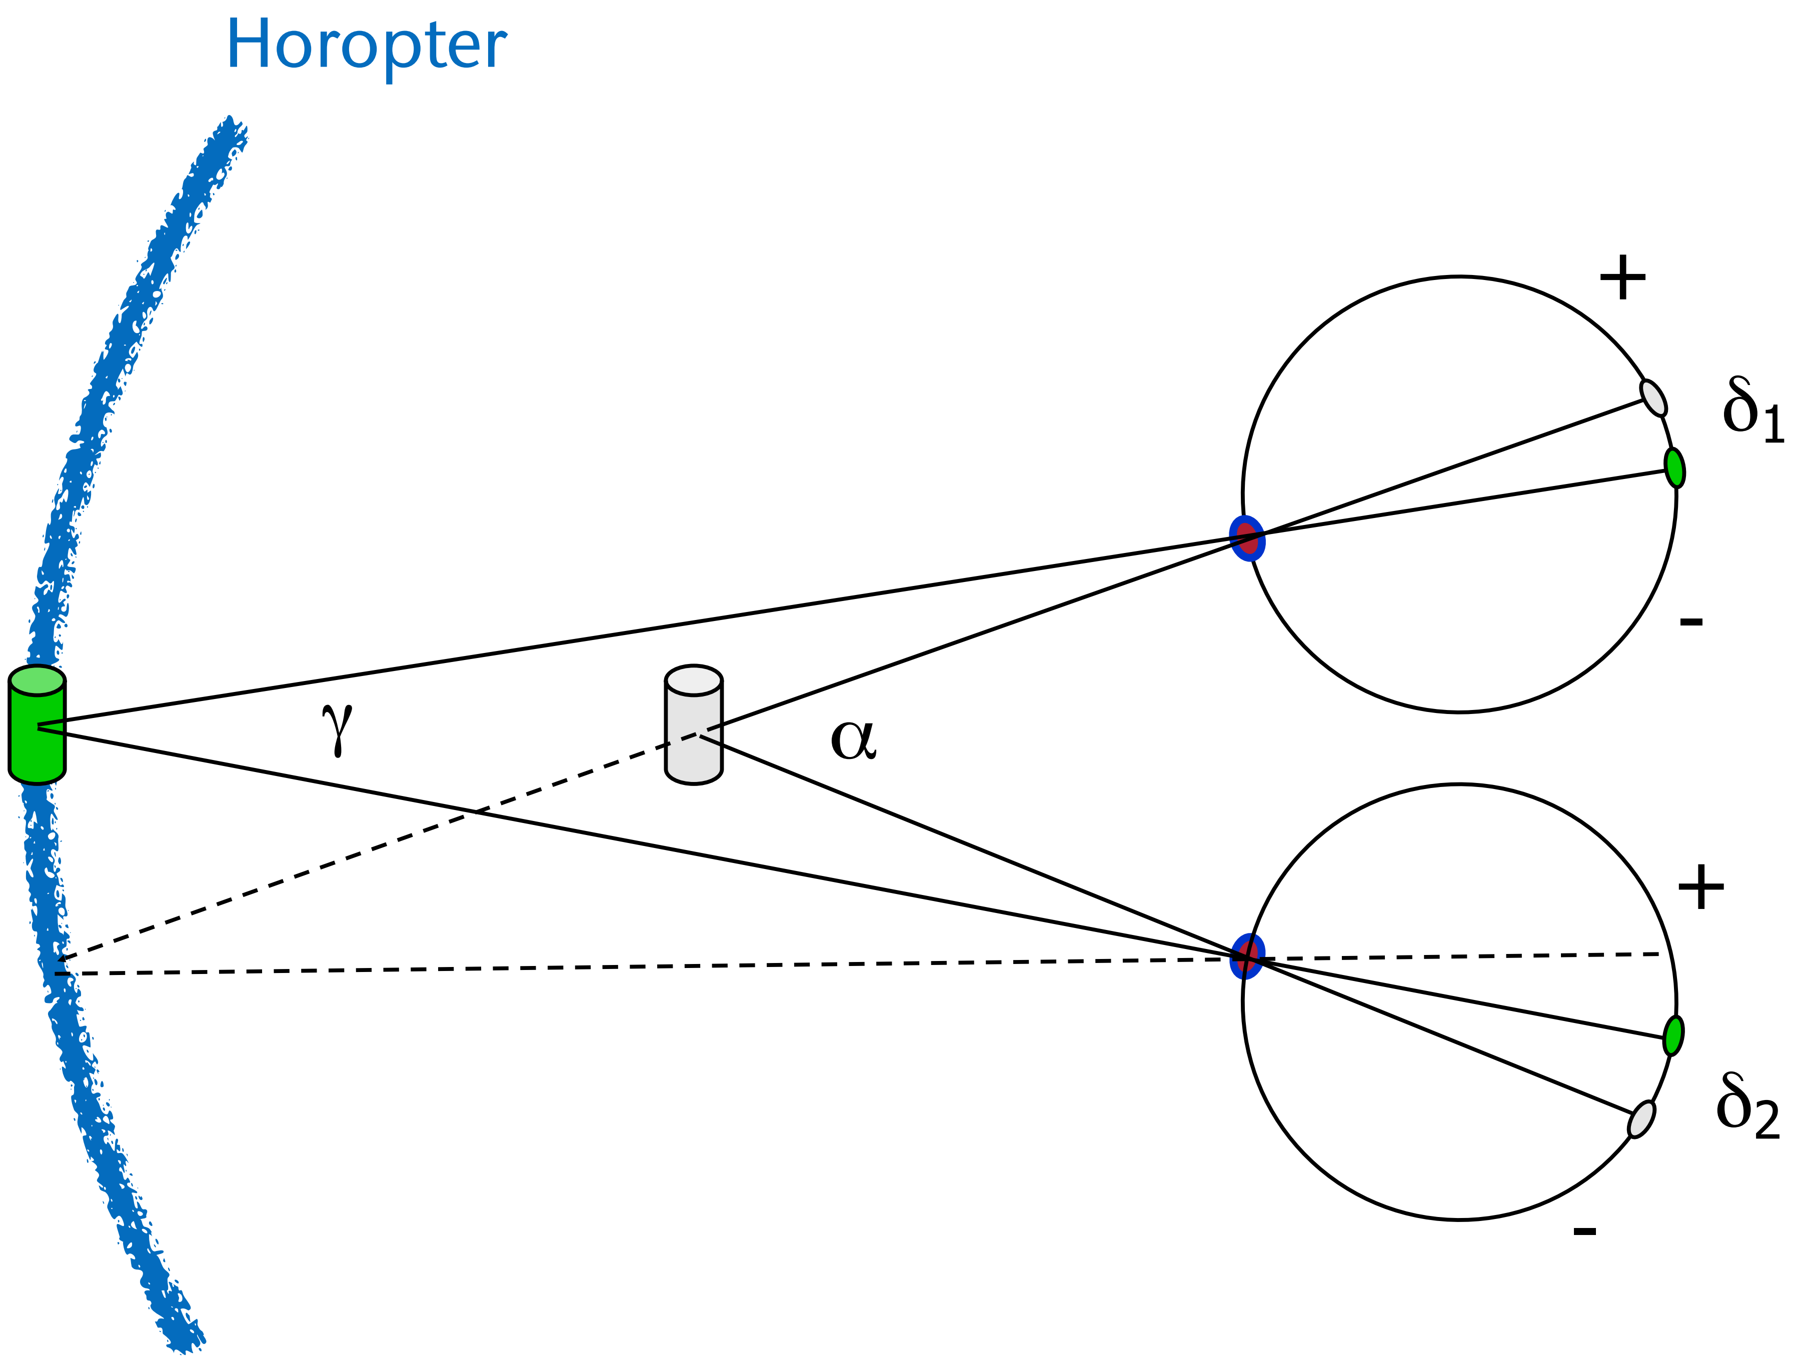

The Shape of the Horopter

- Mathematical construction
  $\longrightarrow$ Vieth-Müller Circle
  = theoretical locus of points in space that stimulate corresponding retinal points

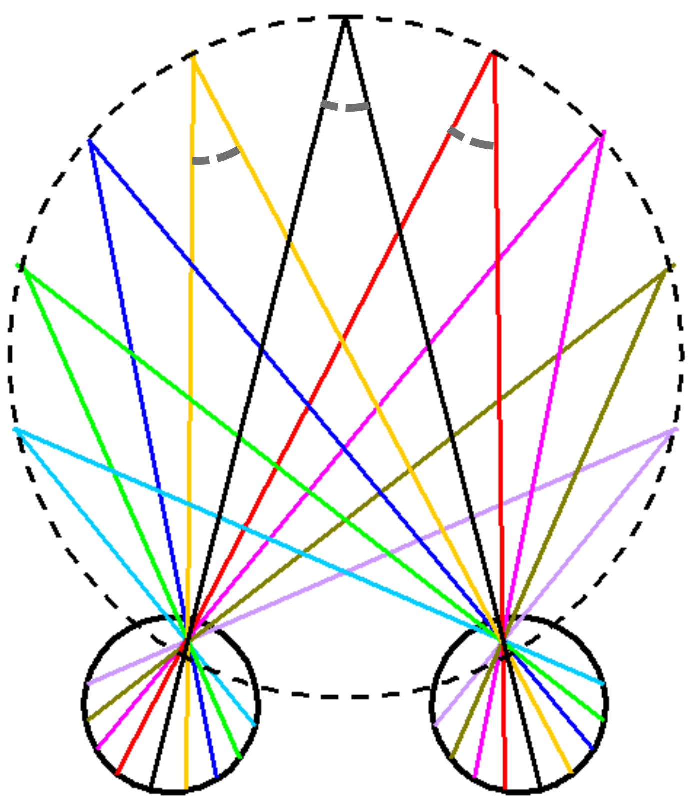

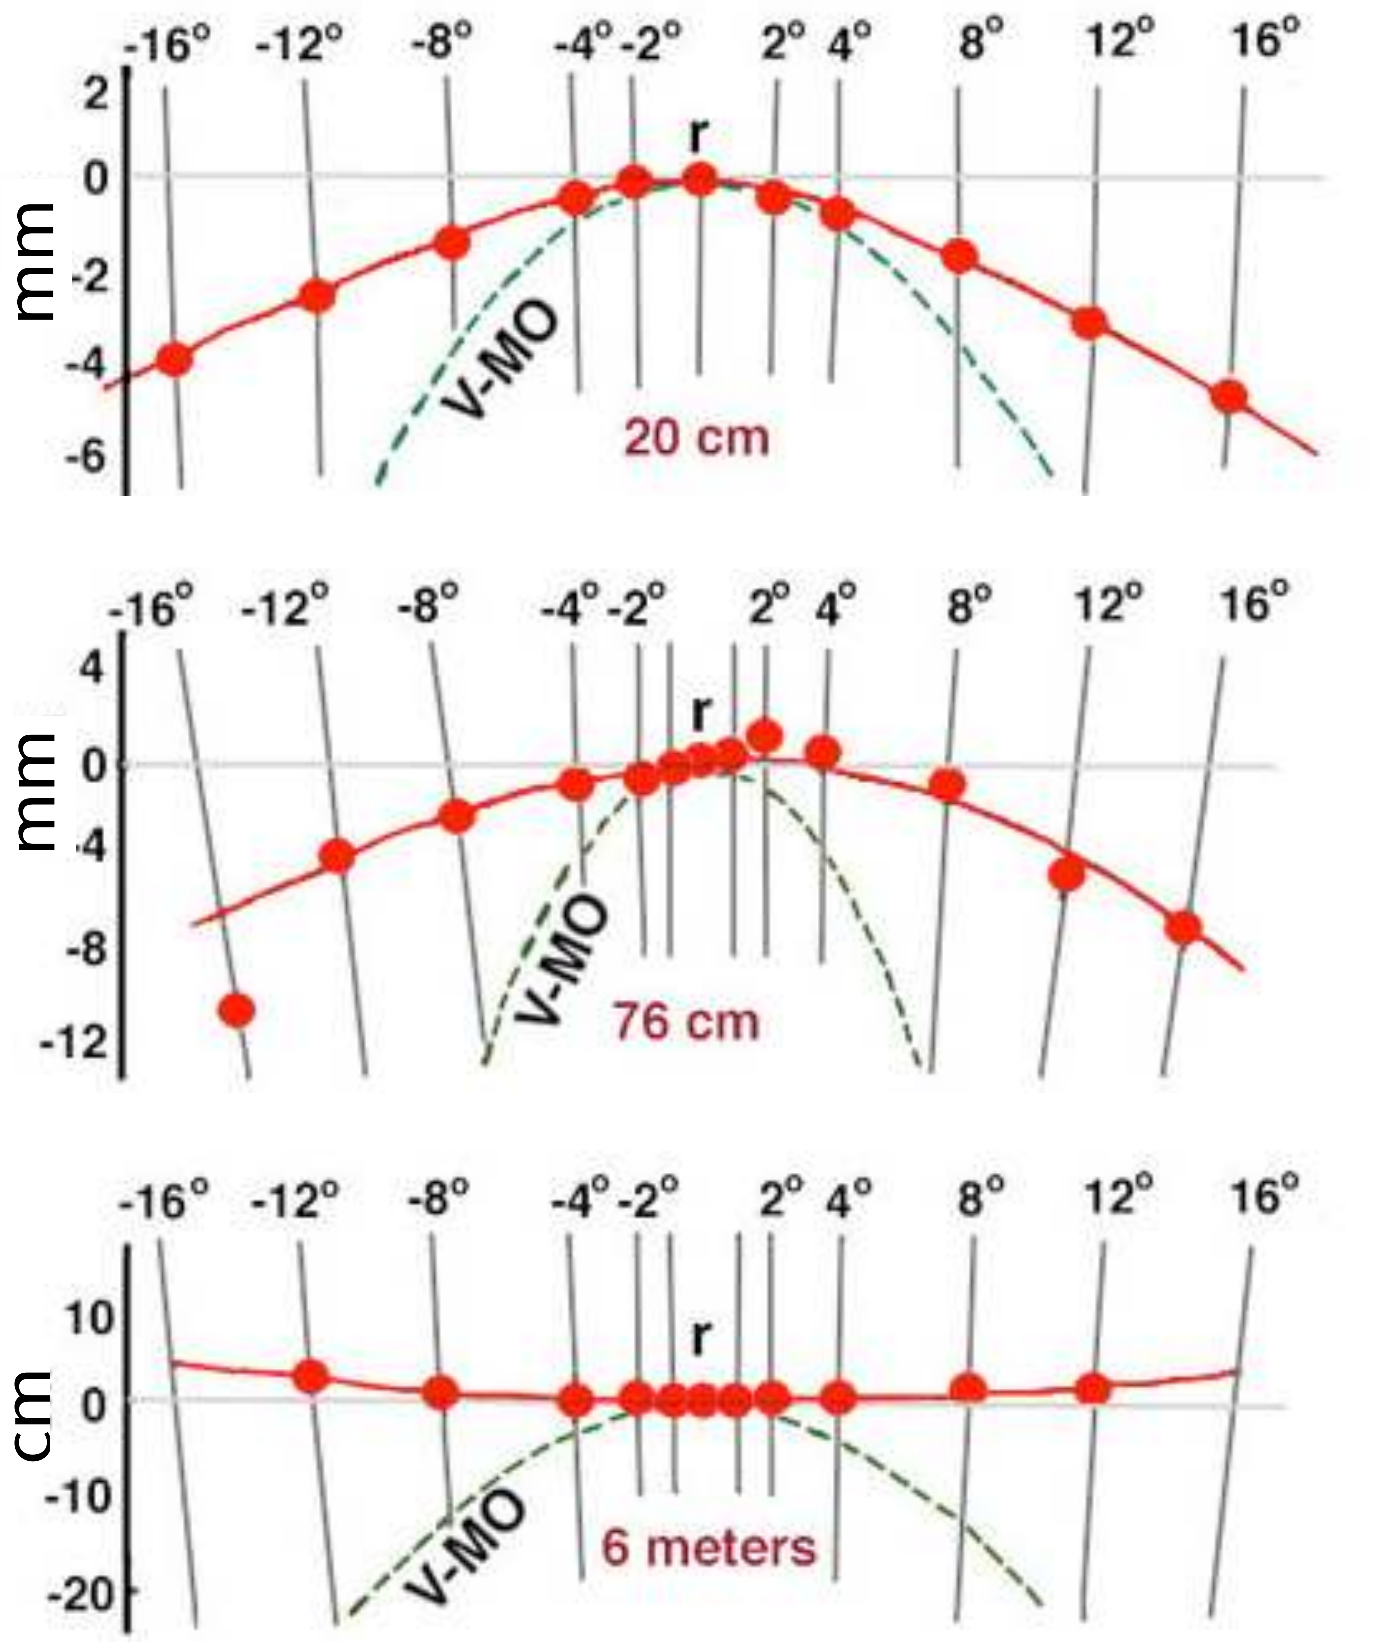

- Measuring the horopter with the "Apparent Fronto-Parallel Plane" method:
  - Subject is asked to arrange a series of objects so that there appears to be no depth difference between them

Panum's Fusional Area

- Disparity limit ≲ 2 deg
- There is a zone/range of depth around the horopter, where the brain is able to fuse the double image of an object
  → Panum's Area of Fusion
- Note: the comfort zone when viewing stereo images is only 1° !

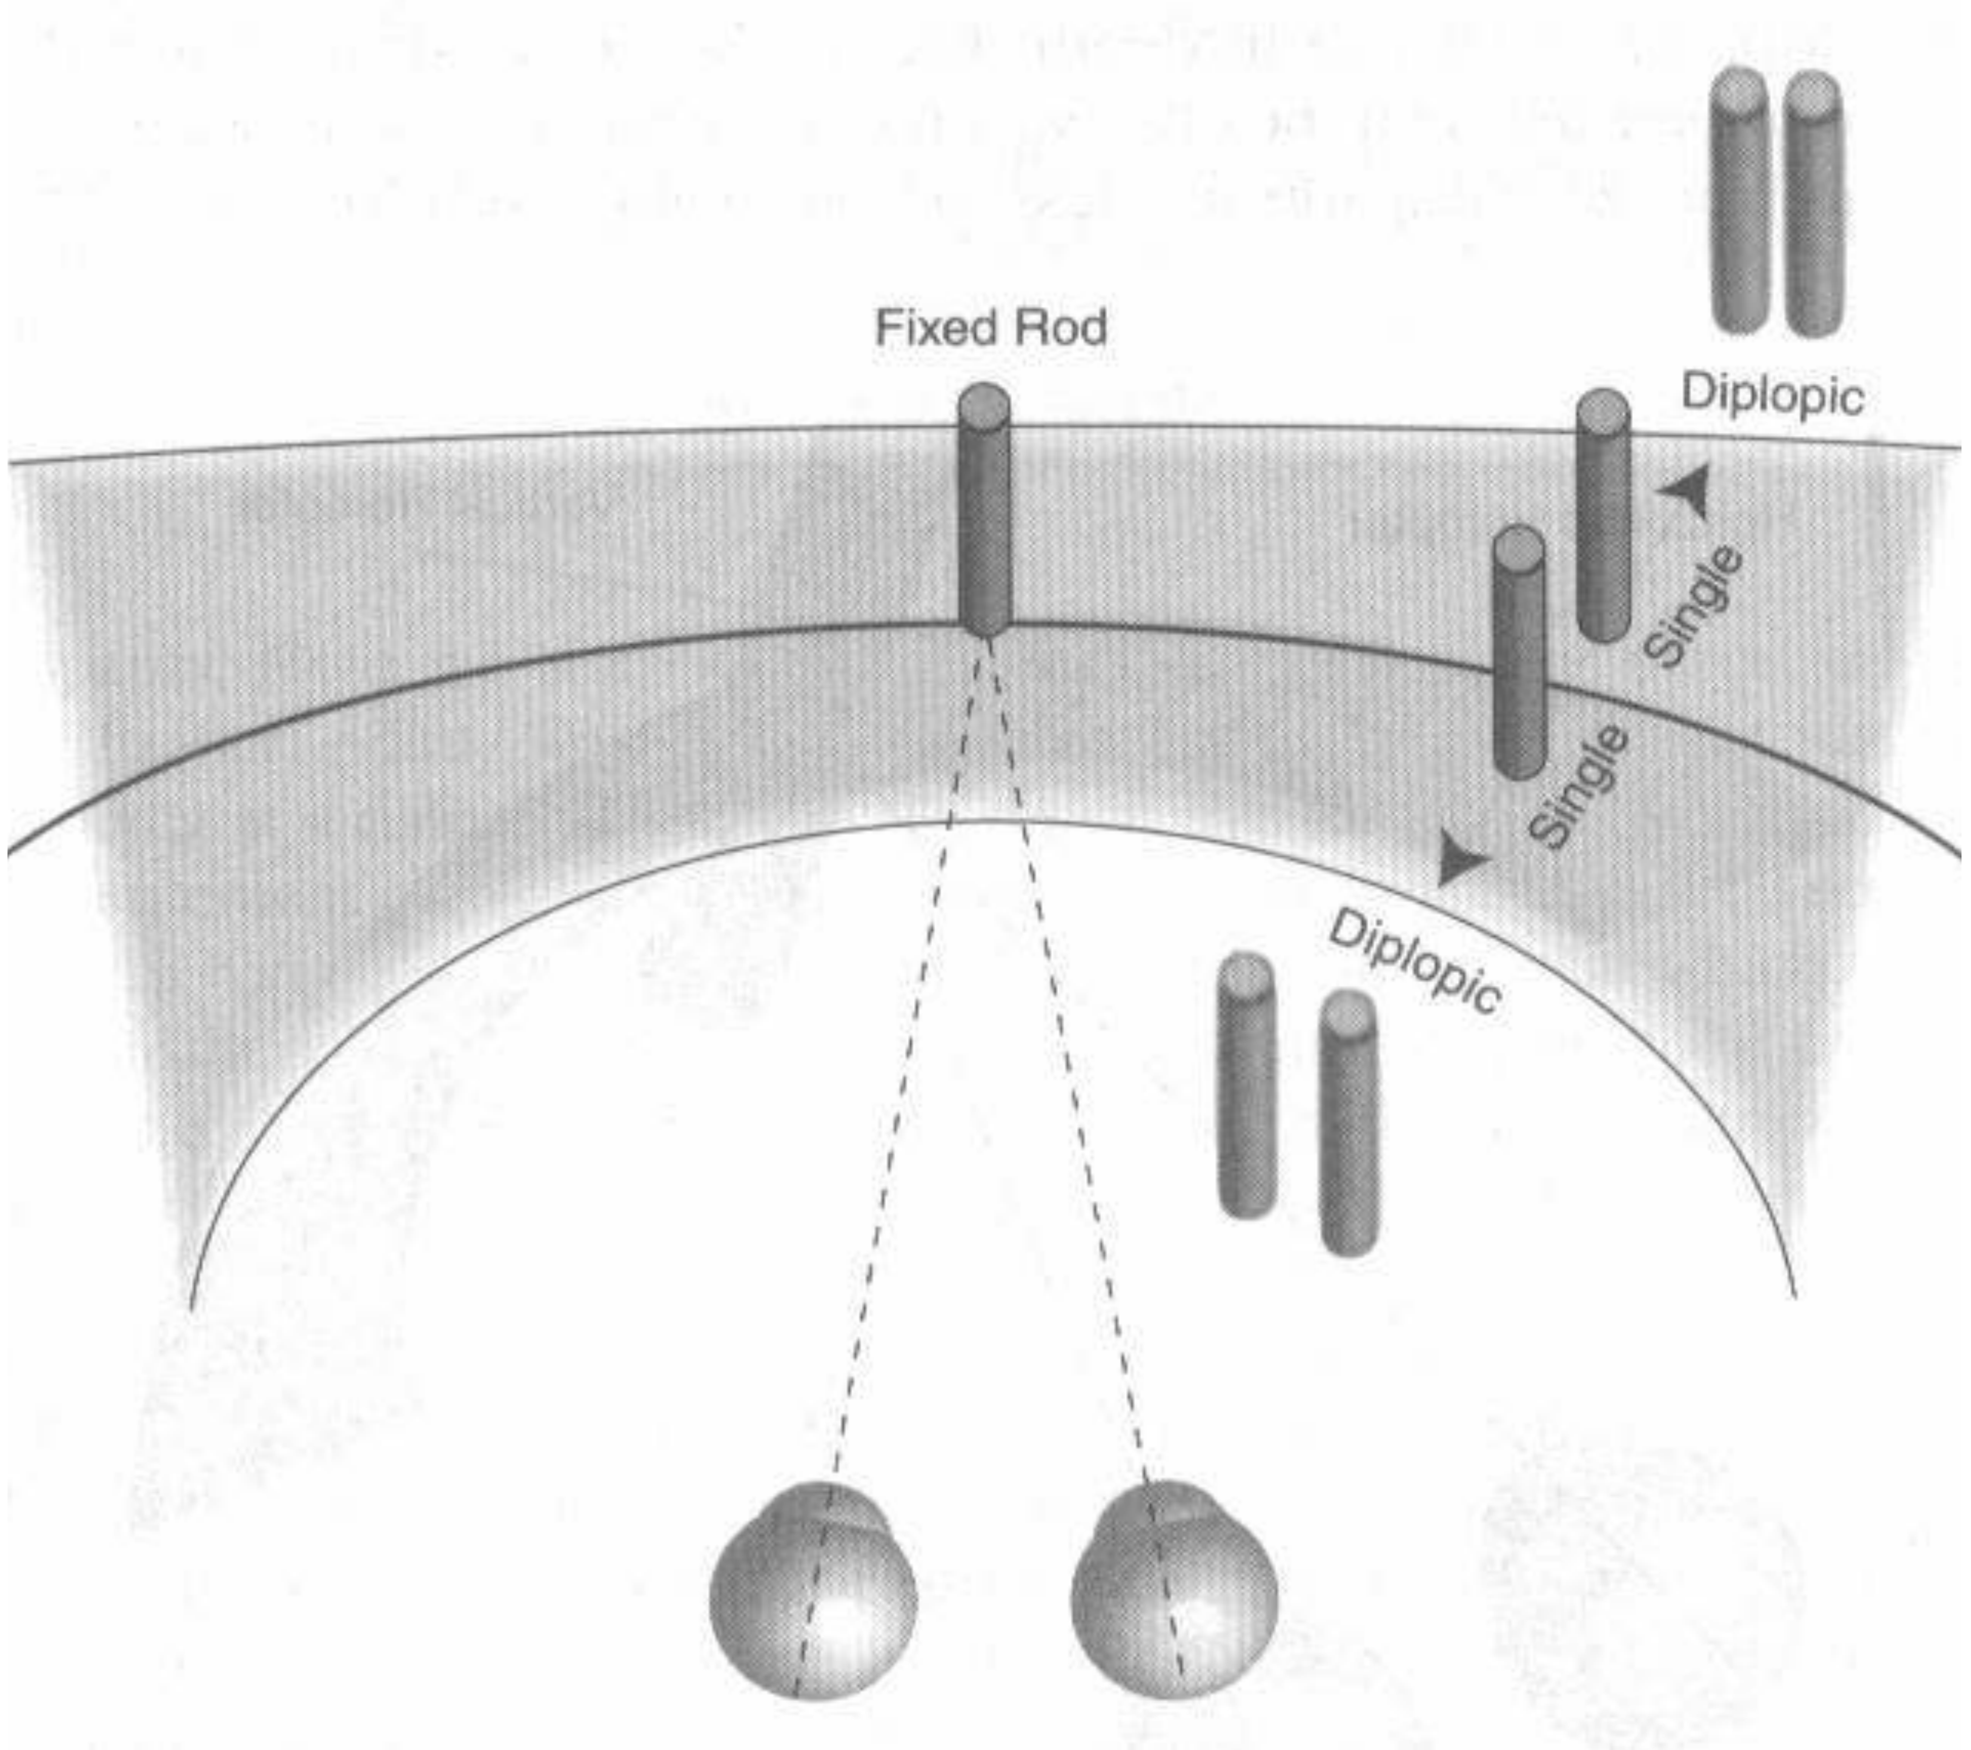

Limitations of Human Stereopsis

- Stereoscopic vision works just up to a few meters (< 6 m, ca.)
- Does not work in the left & right periphery:

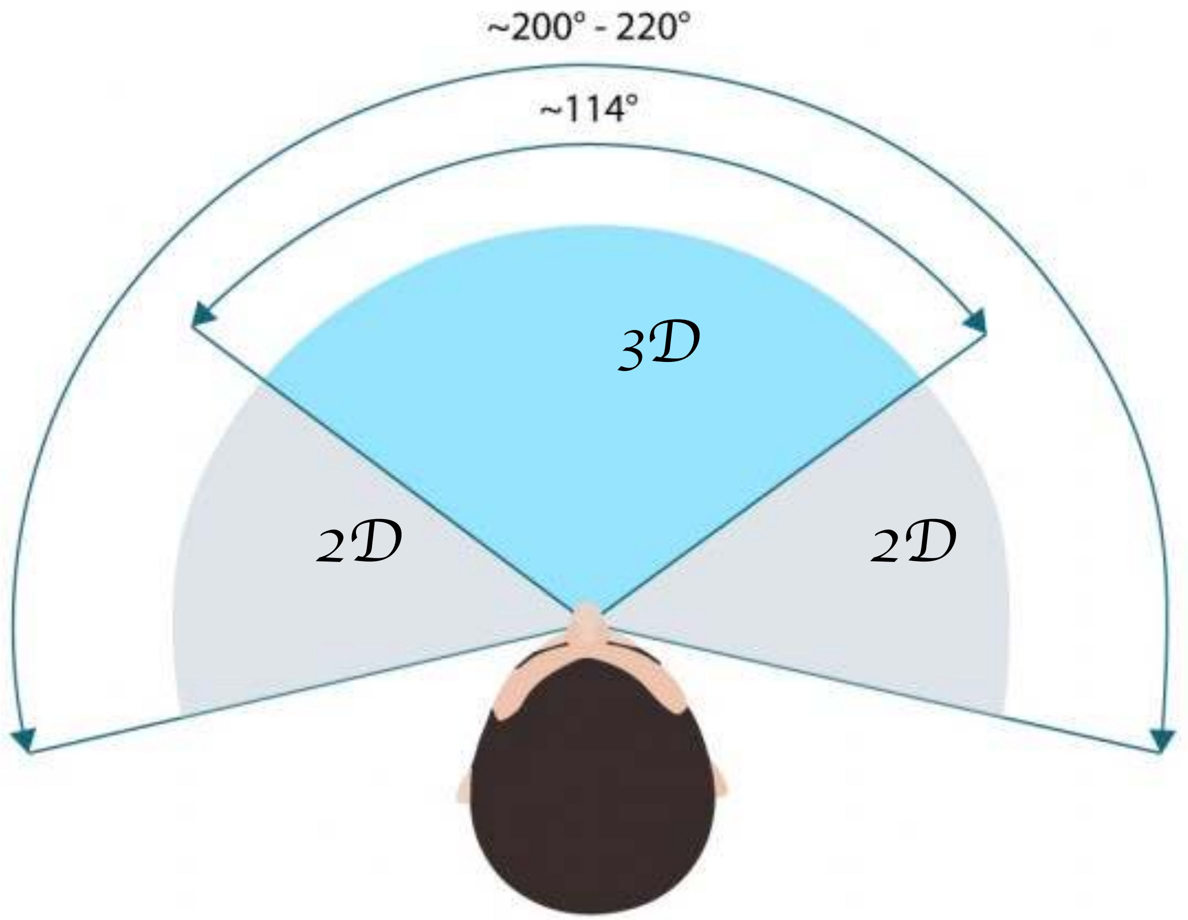

Human

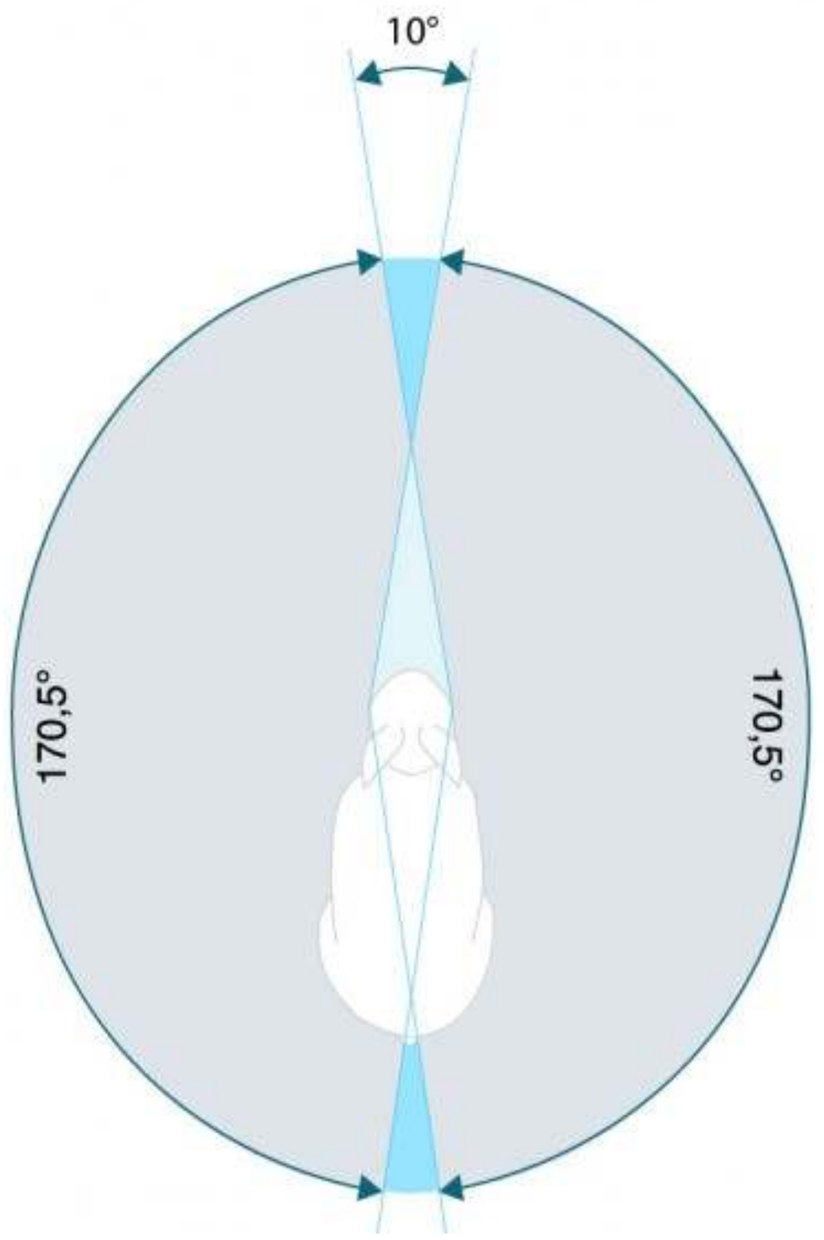

Bunny

Other Depth Cues (Not Exhaustive)

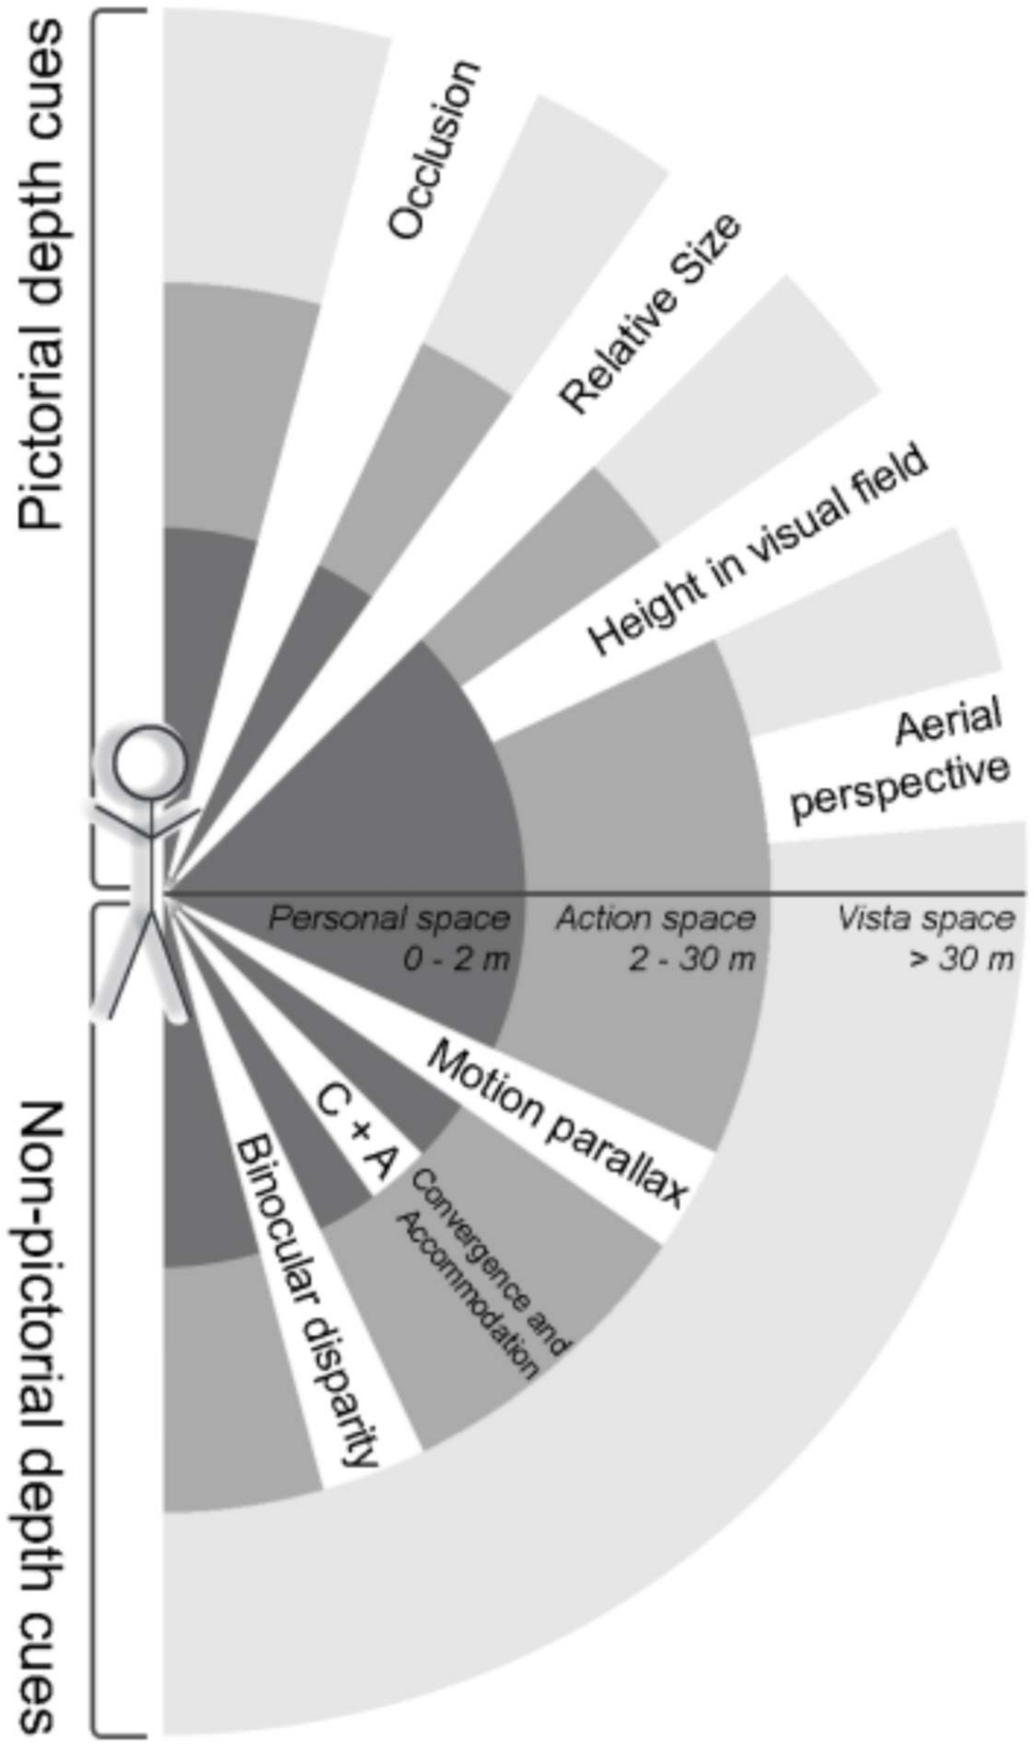

Different Acuities of Different Depth Cues

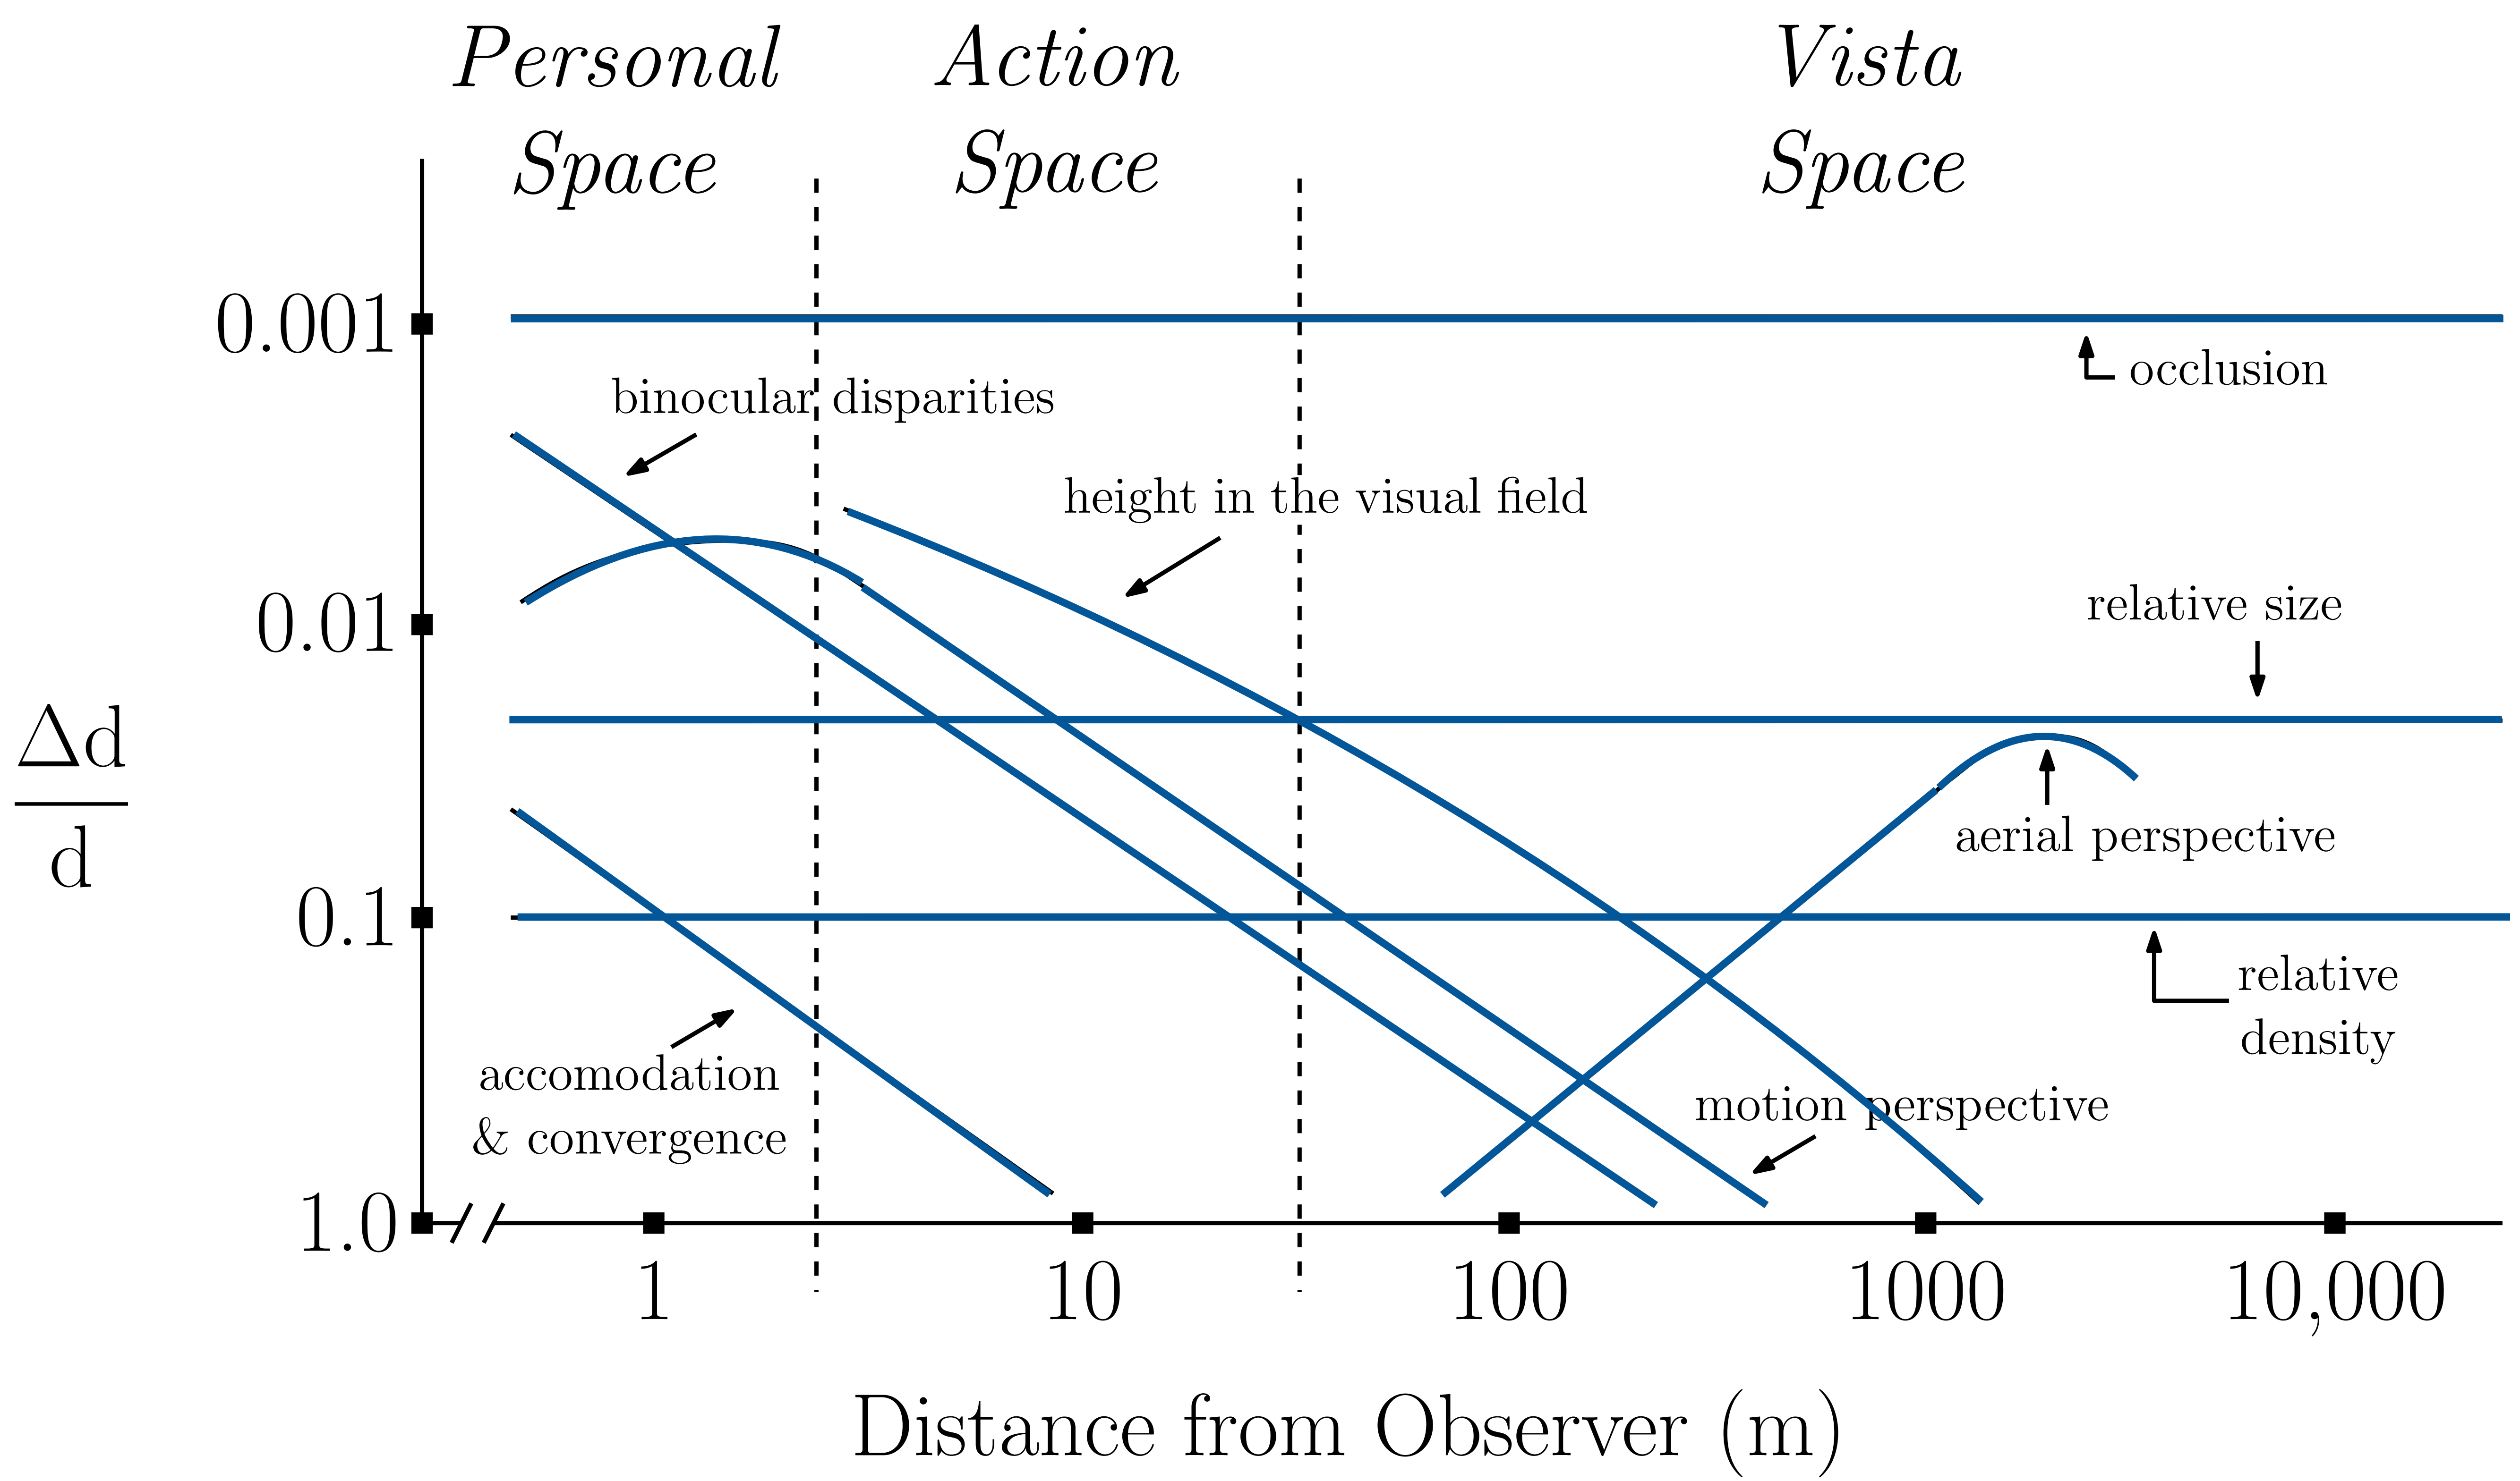

A Short History of Stereo Images/Displays

- Euklid (4th century BC)
- Sir Charles Wheatstone (1838)
- 1860: 1 million stereoscopes sold
- 1950-ies:
- Today (demo):

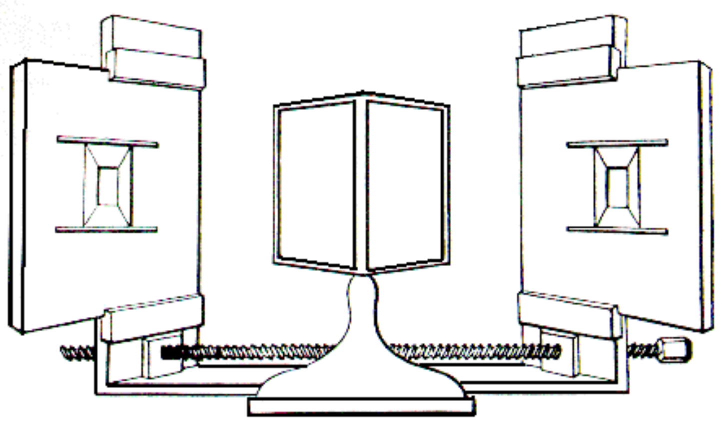

Example Stereogram

- The following image appears to be 3-dimensional, if you can decouple focus (= accomodation) and convergence (you have to scale the slides so that the statues are about 5-7 cm apart, depending on your IPD)

Postcard from 1868

Immersive Displays

- Head-Mounted Displays (HMDs)
- Immersive projection displays (IPDs)
  - Autostereo Monitor
  - Desktop setups
    - E.g. Autostereo monitors, zSpace, or"reach-in"
  - "Powerwall"
  - Workbench
  - Cave

  } A.k.a. World-Fixed Displays
- "Exotic" displays:
  - Retinal displays
  - Holographic displays
  - ...

Stereo Monitor

- Sometimes called "*Fishtank VR*"
- Advantages:
  - Inexpensive
  - Resolution up to 1900 x 1600
  - Well accepted by users (?)
  - No special requirements on the environment
  - Some 3D capabilities
- Disadvantages:
  - Small Field-of-View (FoV)
  - Very little immersion
  - Very limited working volume
  - "*Stereo frame violation*" is very common

52'' Autostereo Display

Stereo monitor with eye tracking (zSpace)

Interesting things you can do with a simple monitor: the "Reach-in idea"

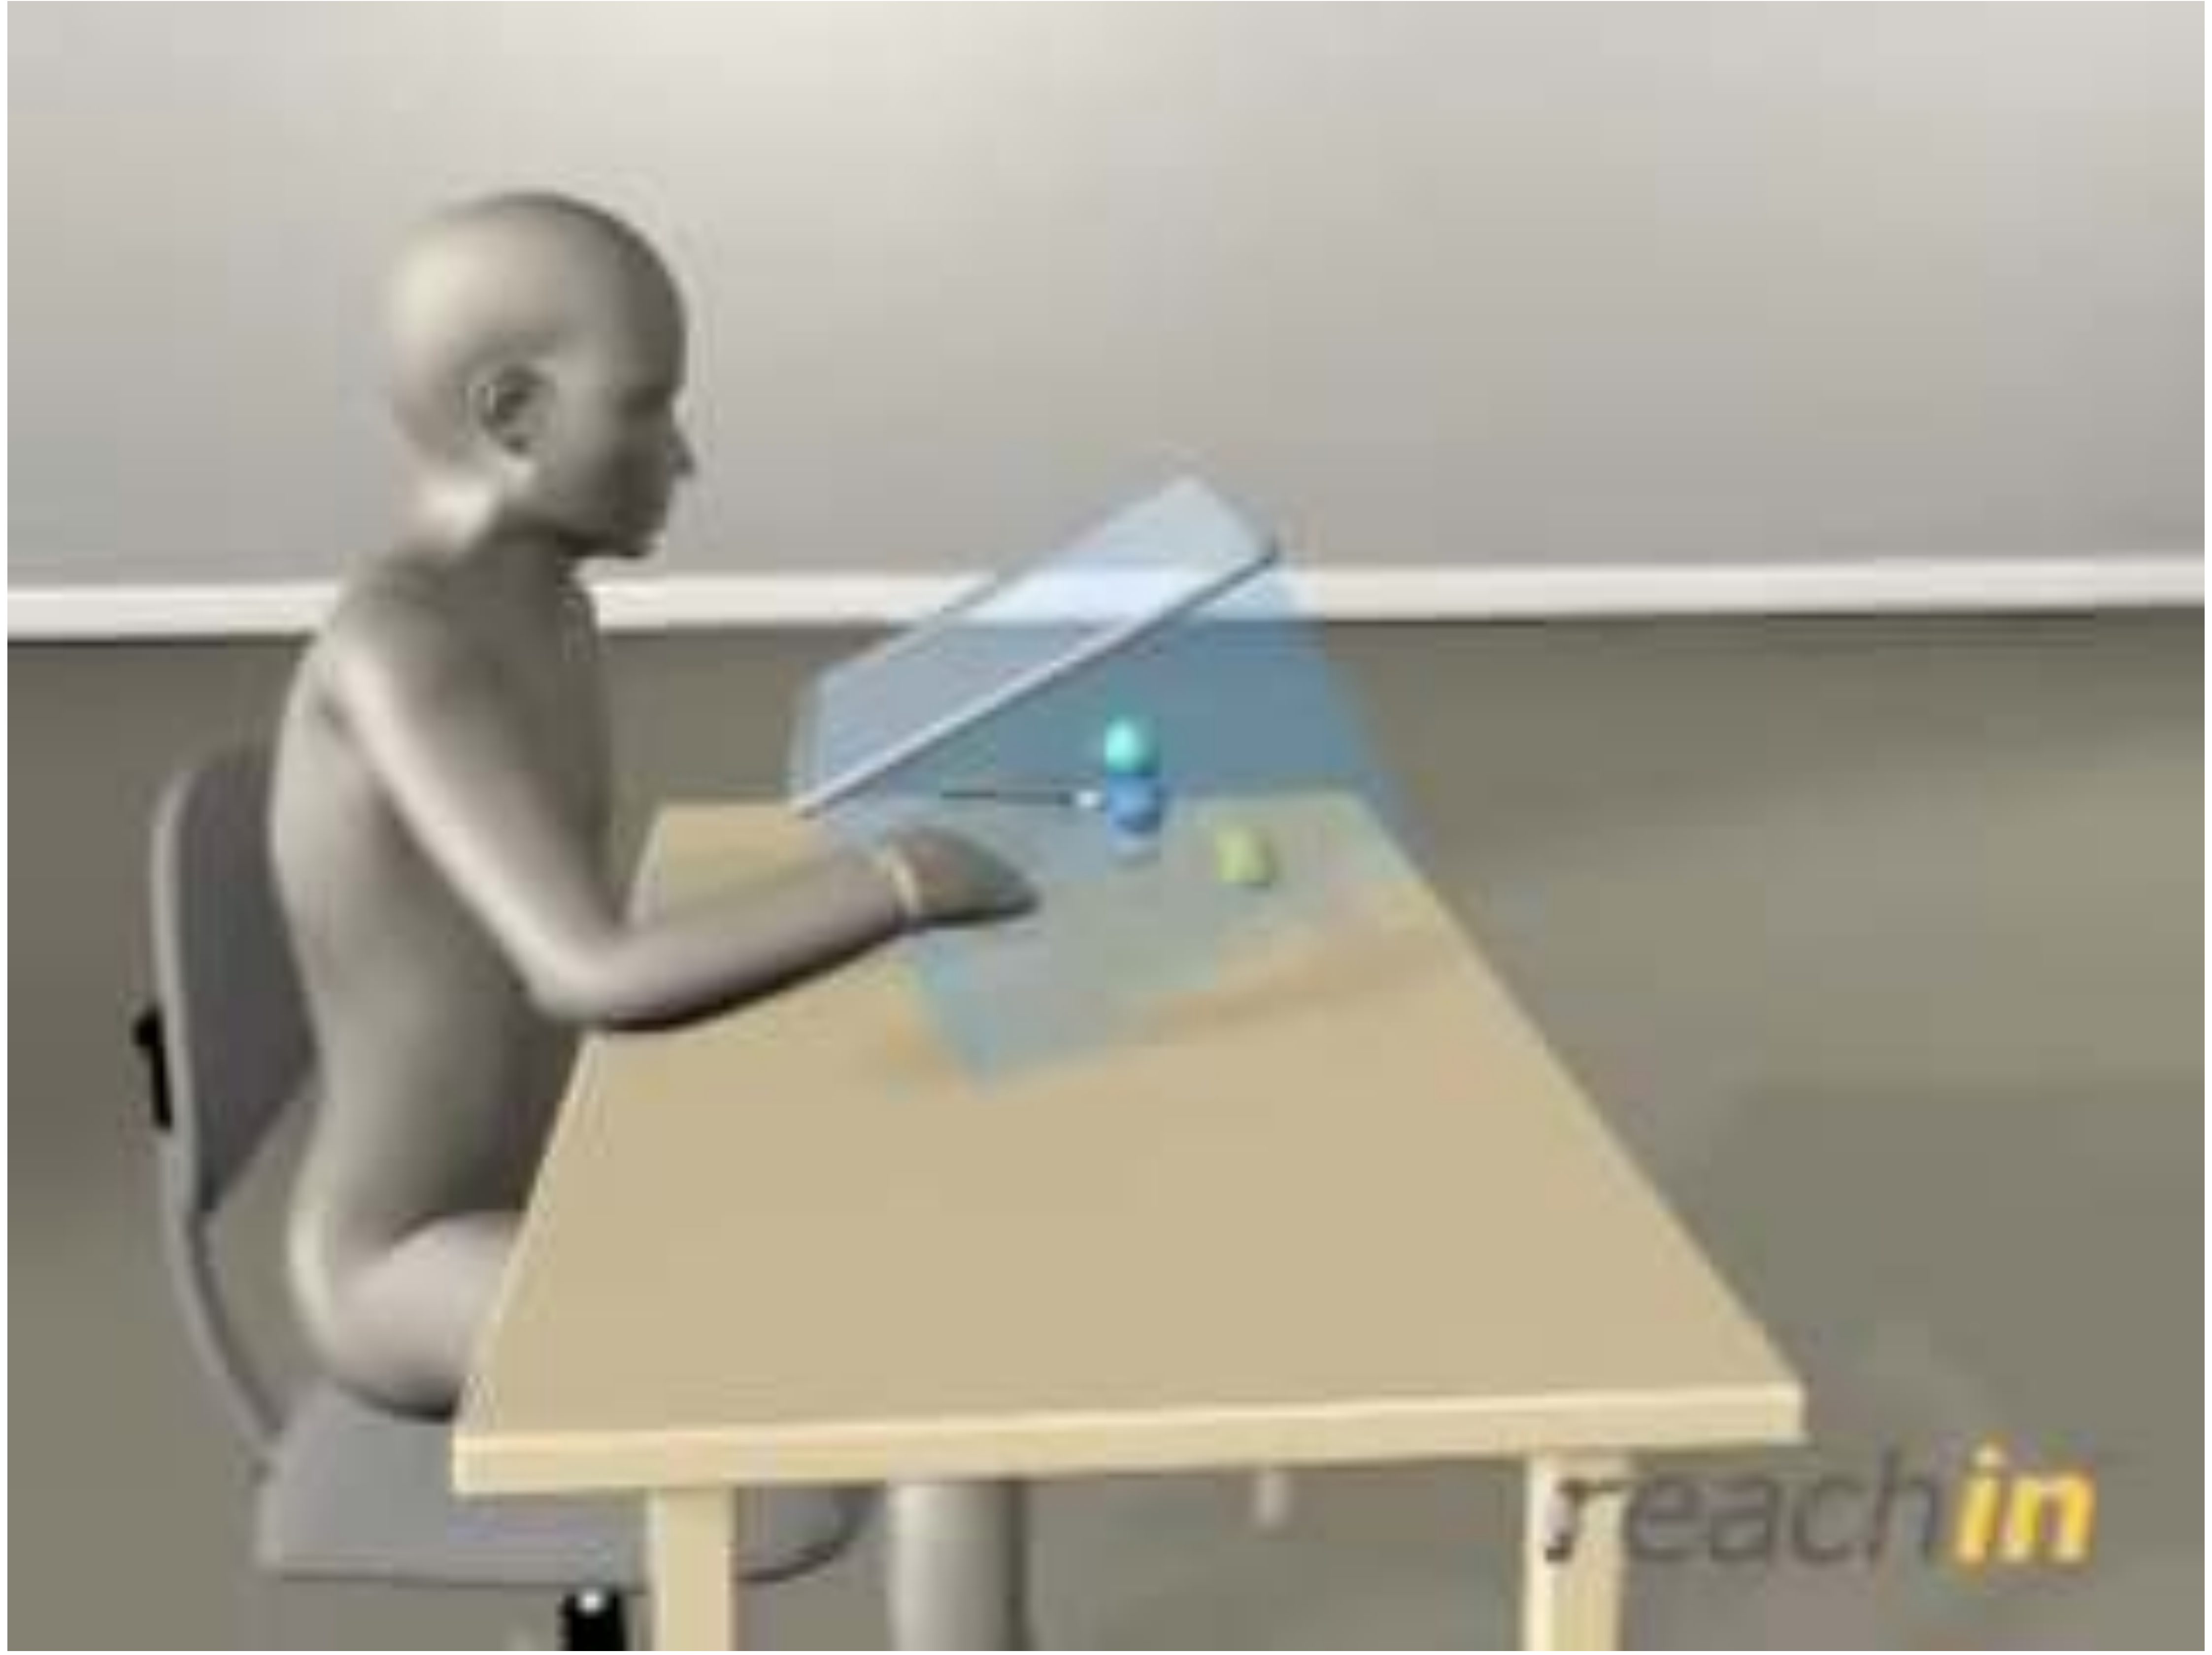

- The problem with a small FoV: there is practically no immersion!

Head-Mounted Displays (HMD)

- First "true" VR display
- Technologies / characteristics:
  - HMDs using LCDs or OLEDs
  - Weight: Small FoV ⟶ lightweight; large FoV ⟶ heavy
- Advantages:
  - Kind of a "surround display"
  - In theory, very good immersion
  - No *stereo frame violation*
  - Large working volume
  - Almost no special requirements on the working environment
  - No channel separation by multiplexing necessary

Around 1992


Around 1984

Other Models (as of 2022)

Oculus VR / Facebook

HTC Vive

Meta Quest 2

Varjo XR-3

HTC Vive Pro Eye

Playstation VR2

Disadvantages of HMDs

- Uncomfortable when used for a prolonged time ("*invasive interface*")
- Distortions (can be corrected somewhat by pre-distortion)
- Real environment is shut off (good for immersion, bad for collaboration and self-embodiment)
- Manipulation of real controls is difficult (e.g., in mockup of cockpit)
- Every participant needs an HMD (bad: expensive, good: everybody has correct perspective in VE)

The Field-of-View Problem of HMDs

The human visual field

Monocular

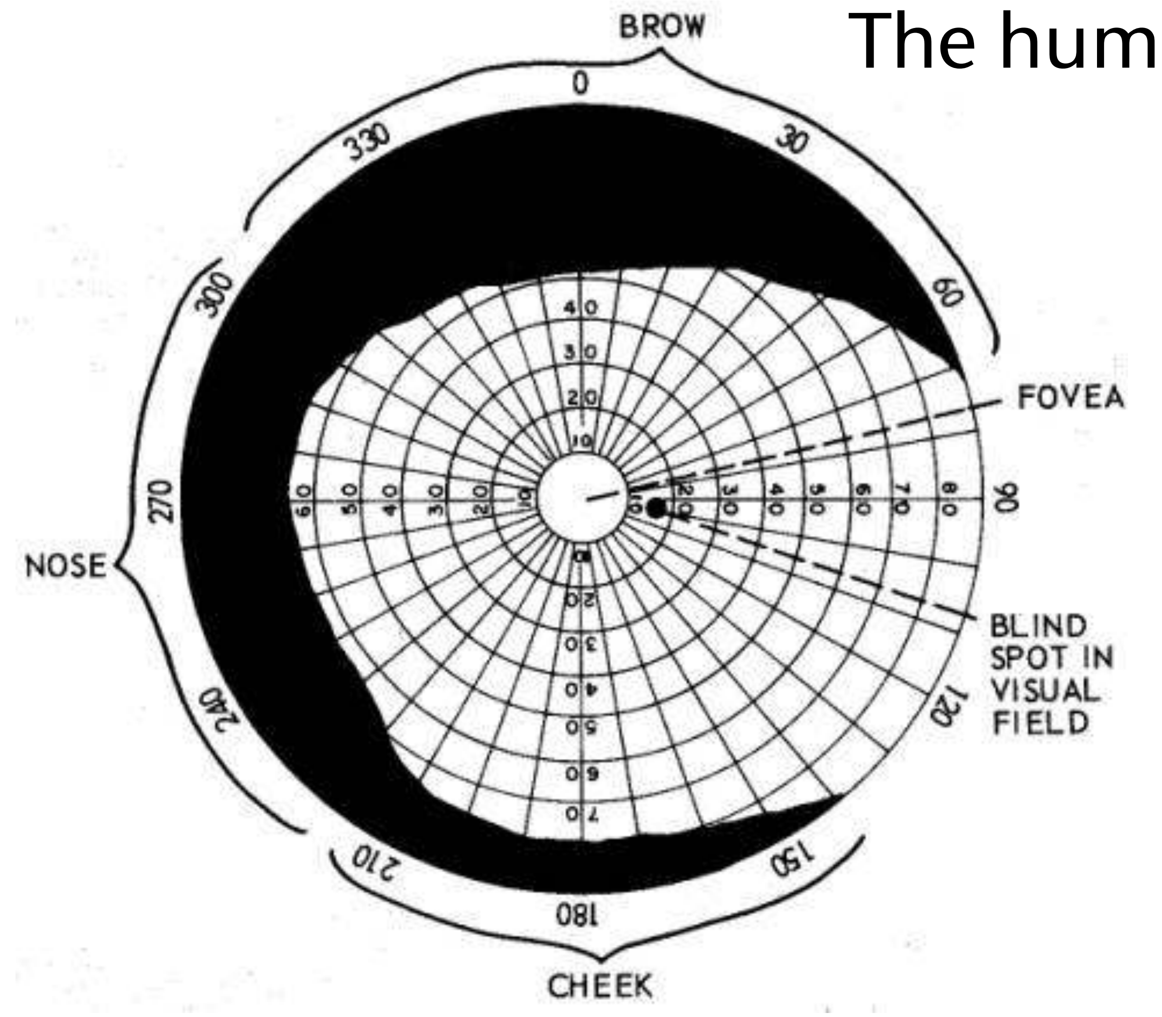

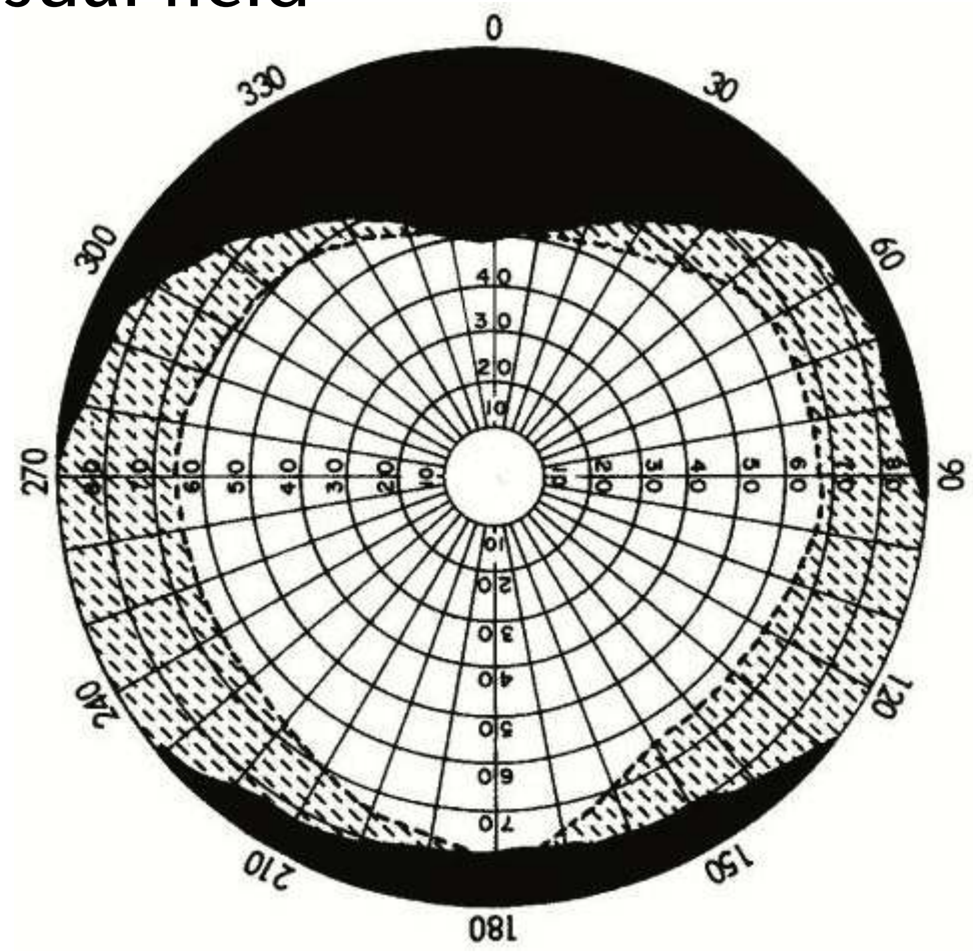

Binocular

Human
Fied of View

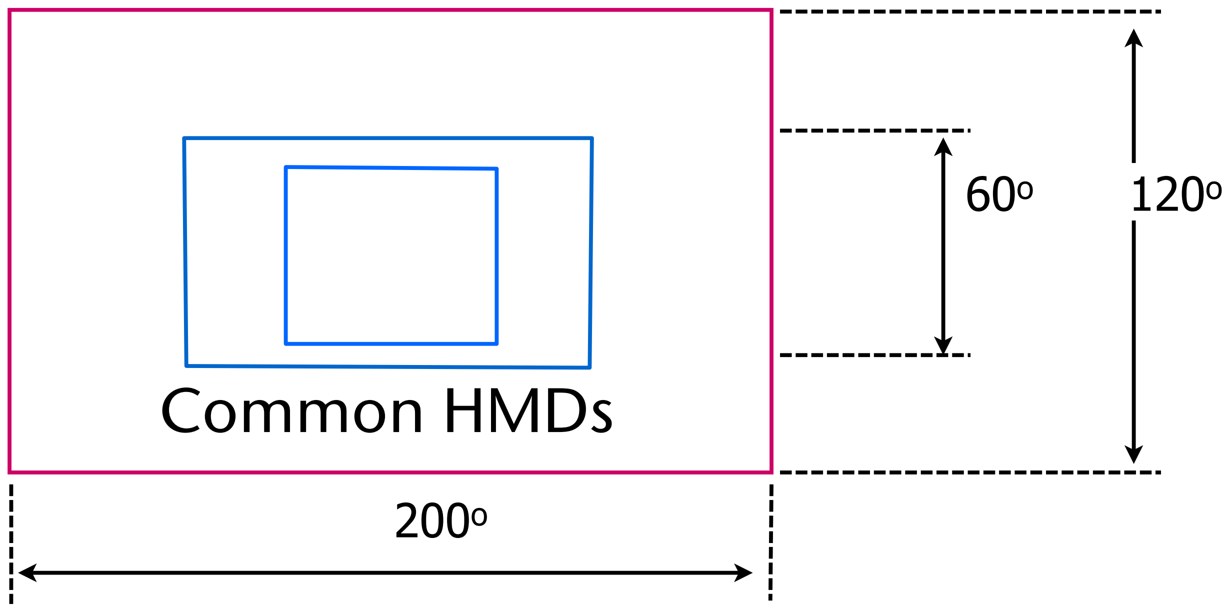

The Resolution Problem of HMDs

- Human visual acuity:
  - 1 photo receptor (cone) = 1 arc min = 1/60 degree
- Display needed for a "retina" HMD:
  - 150° x 135° with 1/60° resolution = 9000 x 8100 pixels per eye
- Challenges:
  - Bandwidth, i.e., moving the data at 60 Hz from GPU to display
  - Miniaturize display panels with 73 Mio pixels

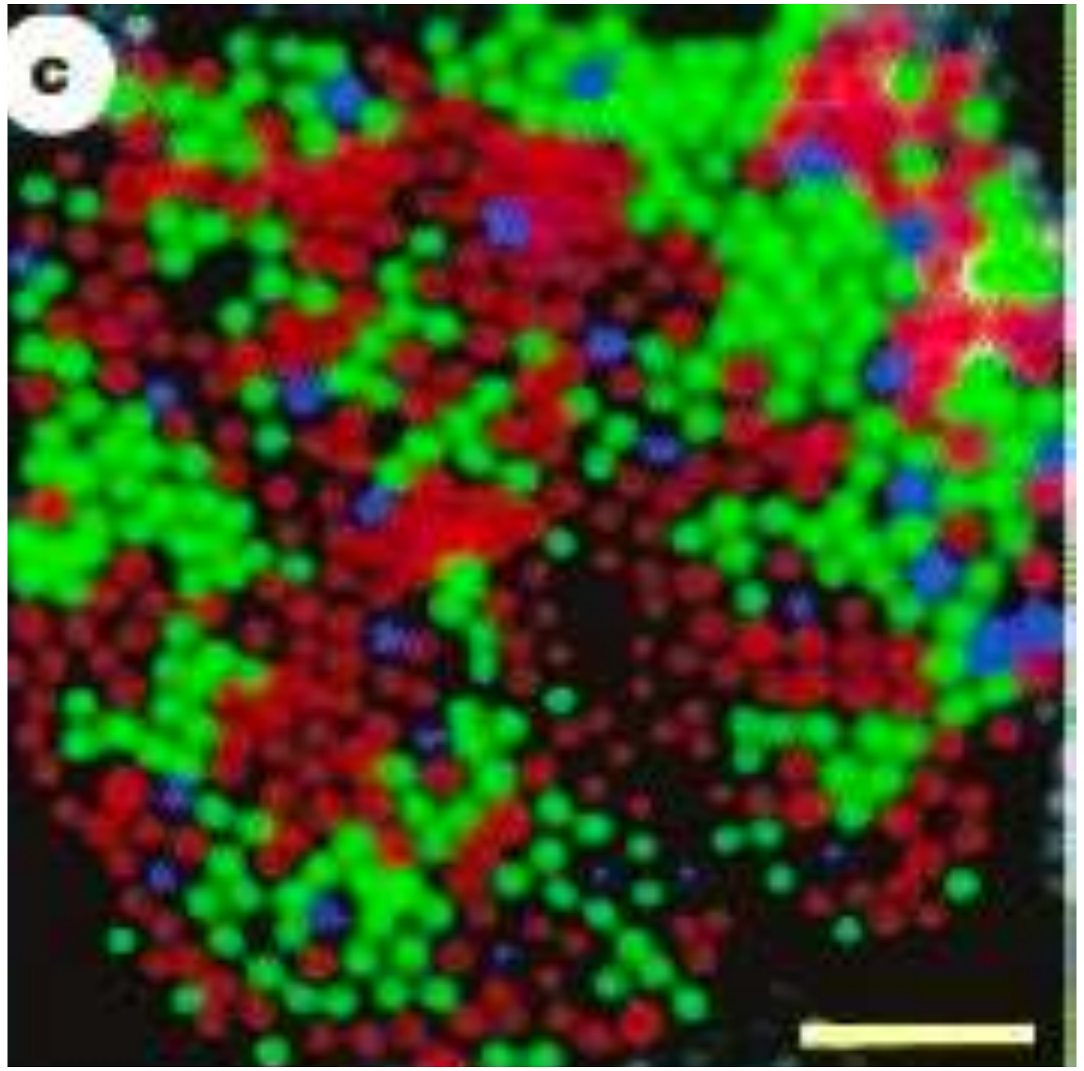

5 arc min

HMDs with Eye Tracking

- Potentials:
  - "Foveated rendering"
    - Requires end-to-end latency of < 10 ms
  - User interaction using eye gaze direction
  - For research
- Dynamically move the zero-parallax plane?
- Control focus depth for depth-of-field rendering?
- Make eye contact with virtual avatars (NPC)?
  - So they "notice" and look back at you
- Shoot enemies in games just by looking at them?

Head Coupled Displays (HCD) – Out-Dated

- HCD = HMD mounted on a "boom"
- Advantage of HCDs over HMDs:
  - Possible to quickly "take the display off" for a moment; or users can just take a "quick peek" into the VE
  - Low weight on the head
  - Extremely good tracking comes built-in
- Disadvantages compared to HMDs:
  - Smaller working volume
  - One hand is always occupied
  - Inertia
- Failed to gain market share

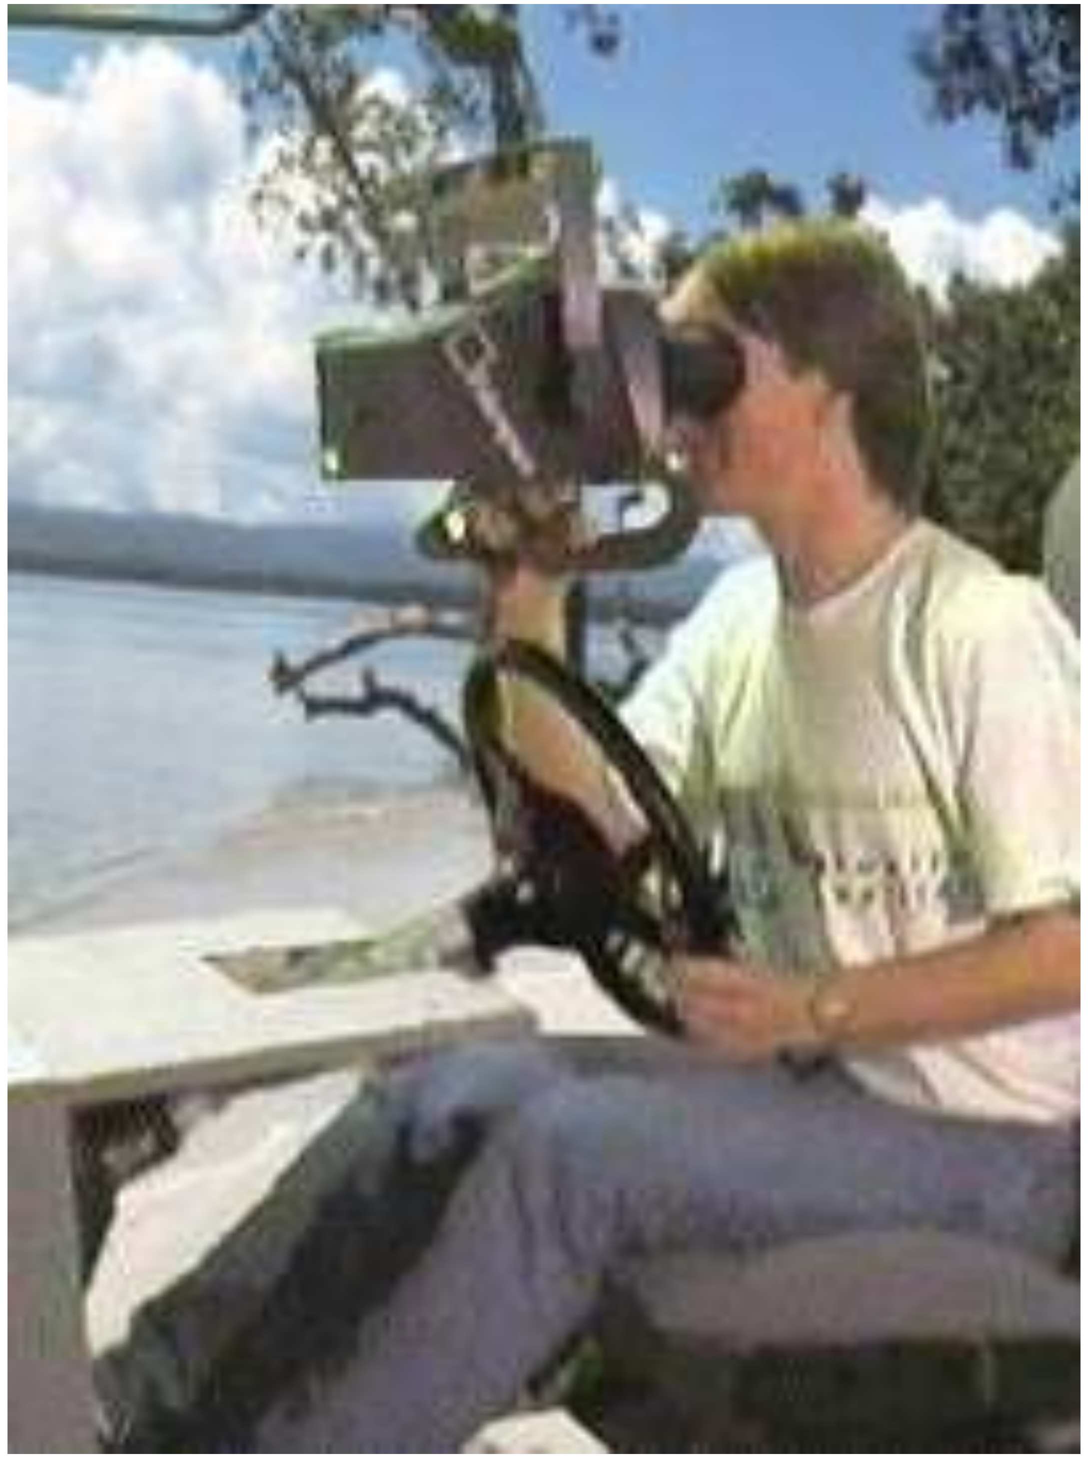

Immersive Projection Displays / Technology (IPD / IPT)

- Idea is (somewhat) similar to cinema theaters
- Setup: 1–6 walls on which VE is projected
- *Powerwall* = 1 wall (e.g., 3x6 meters)
- *Workbench* = 1 horizontal display surface (table)
- *Holobench, L-Shape* = 2 display surfaces, 1 vertical, 1 horizontal
- *Cave* = 3–6 walls

Large-Screen Projection Walls (Powerwalls)

- "*HeyeWall*" (Darmstadt):
  - 24 tiles, 48 PCs
  - Total resolution: 18 Mio pixels (6144 x 3072) in stereo

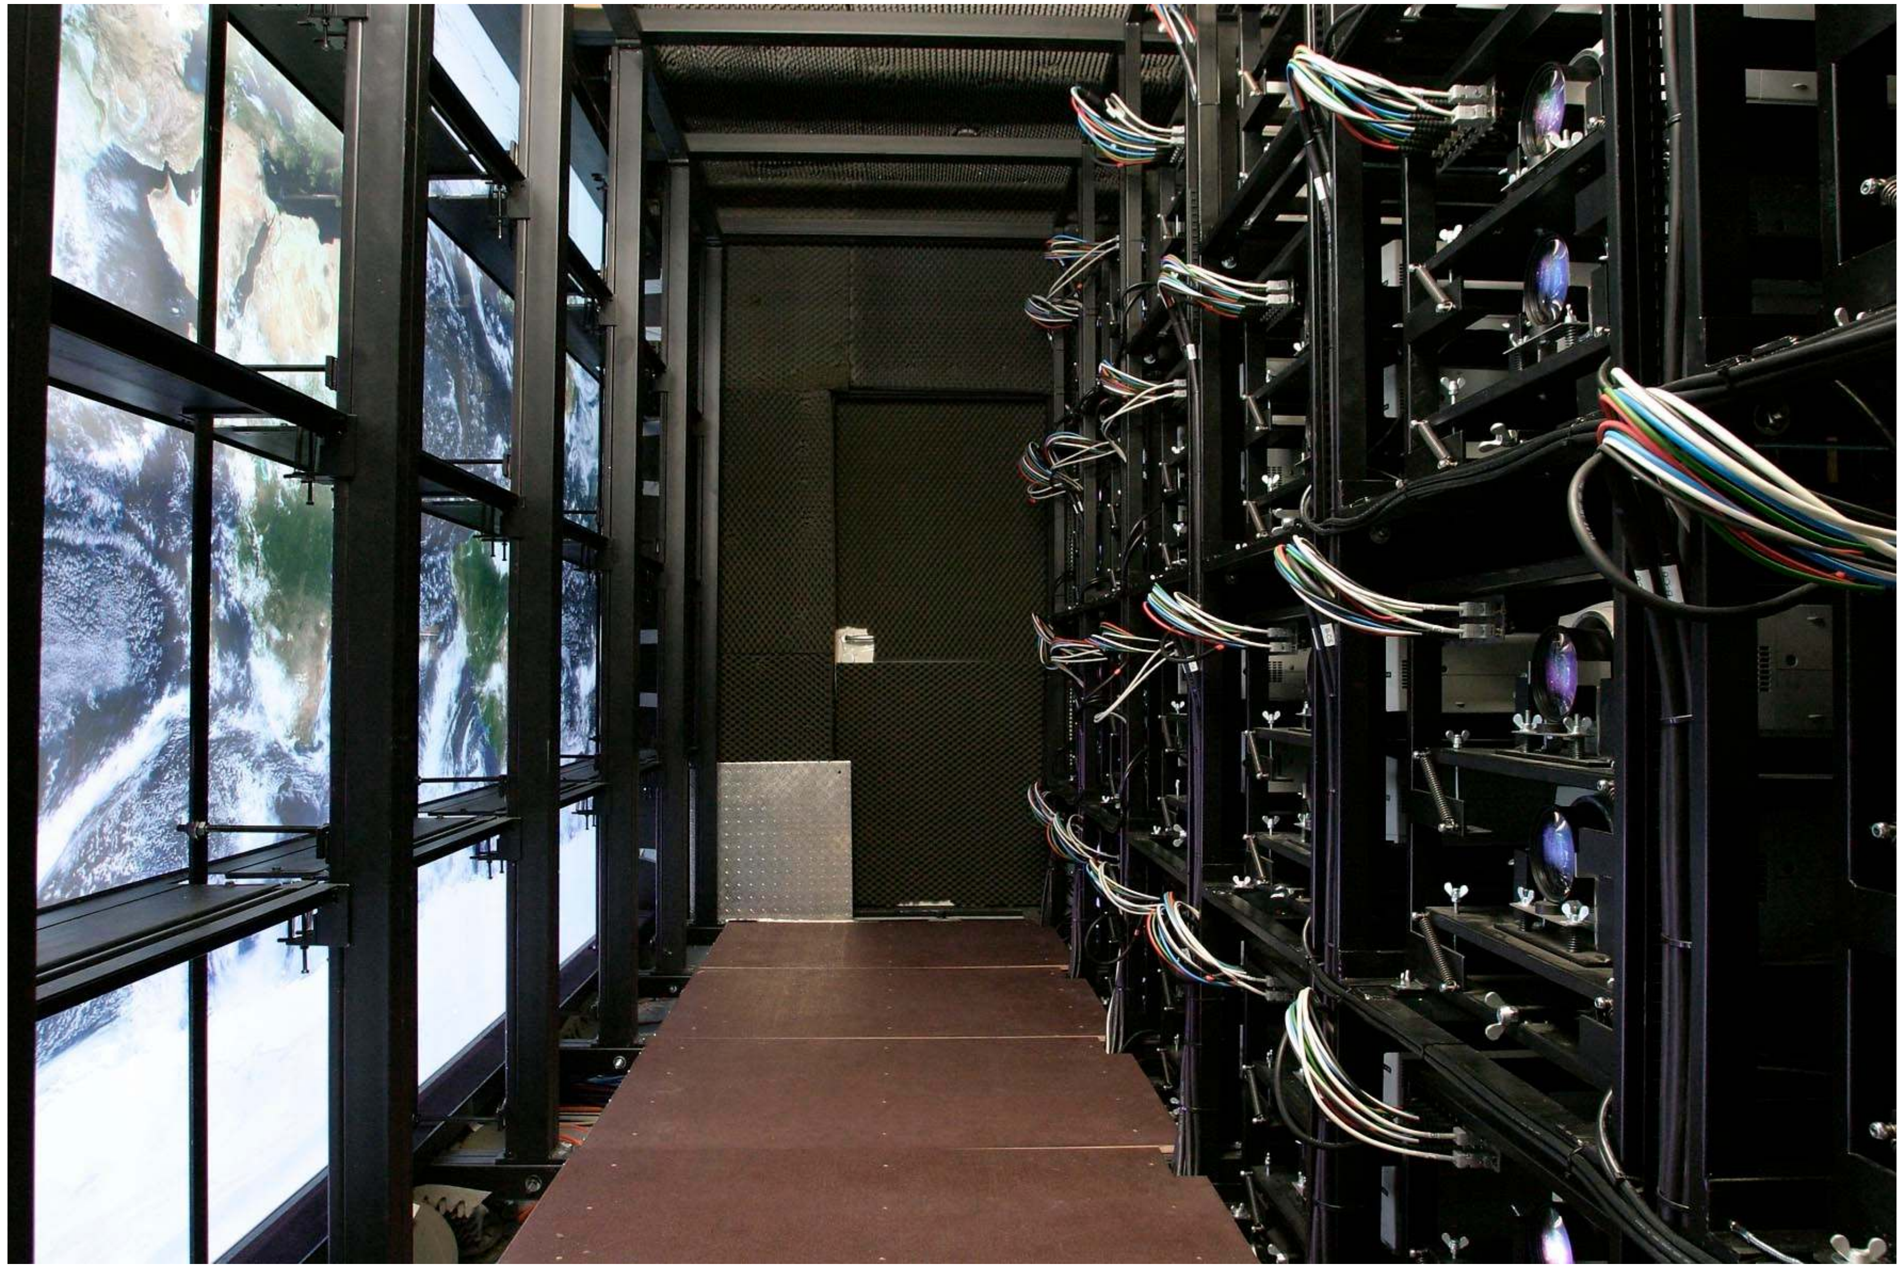

Example Application: Virtual Conference Room

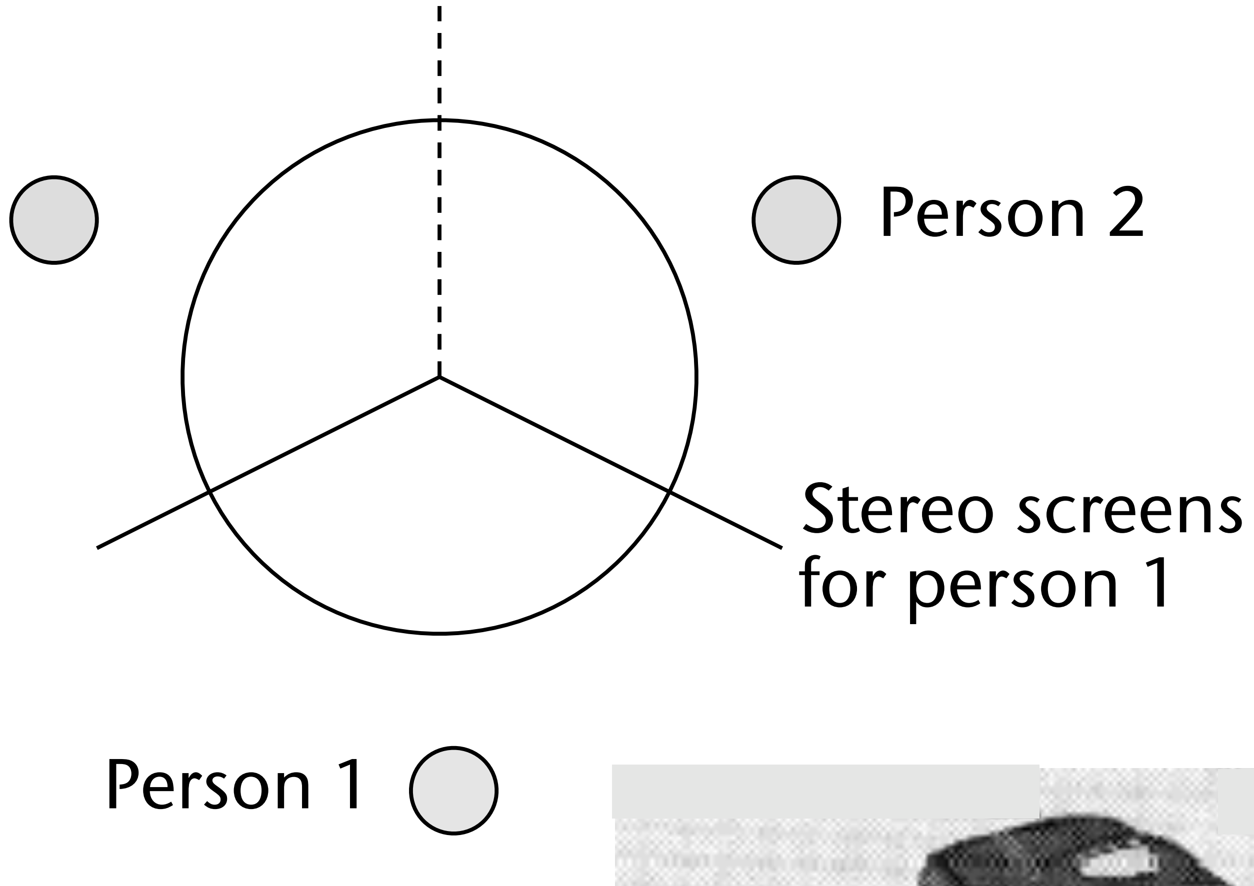

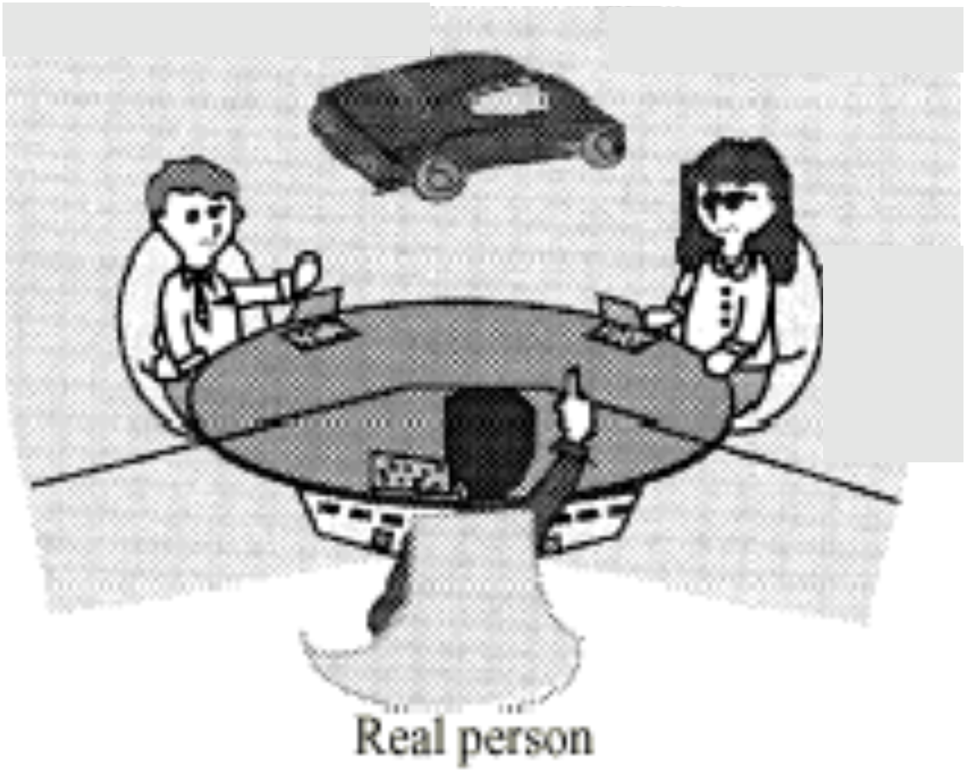

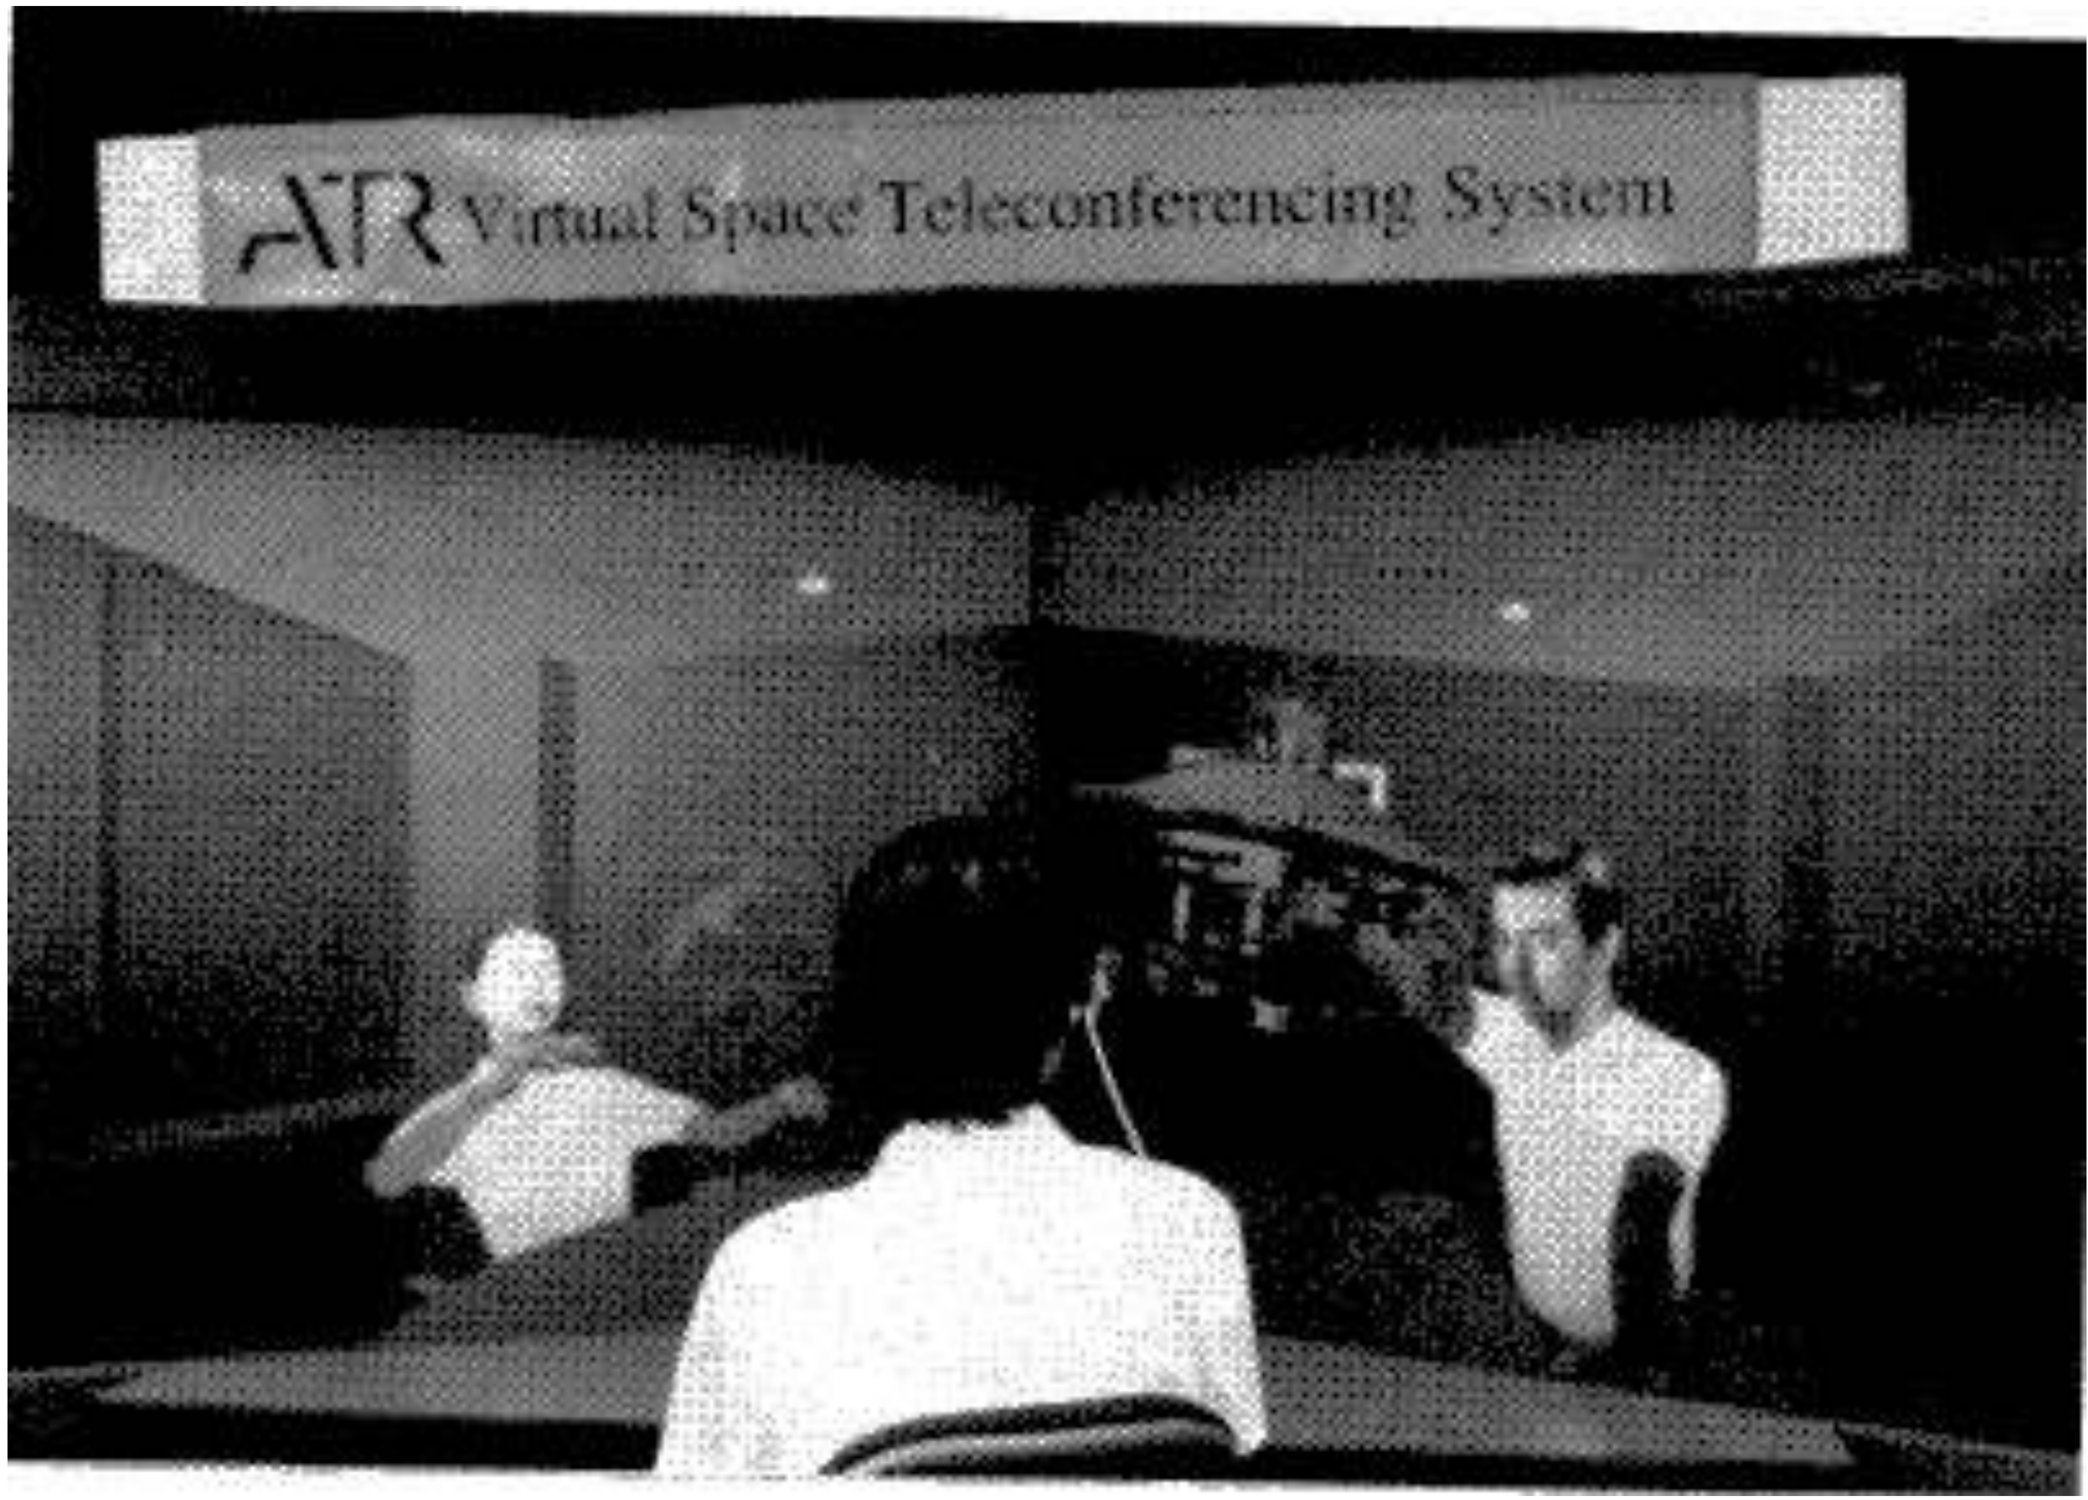

Result: a *single, shared, coherent workspace*, by way of coherently adjoining "desktop IPDs"

Workbench, L-Shape, Holobench, etc.

Workbench

Holobench

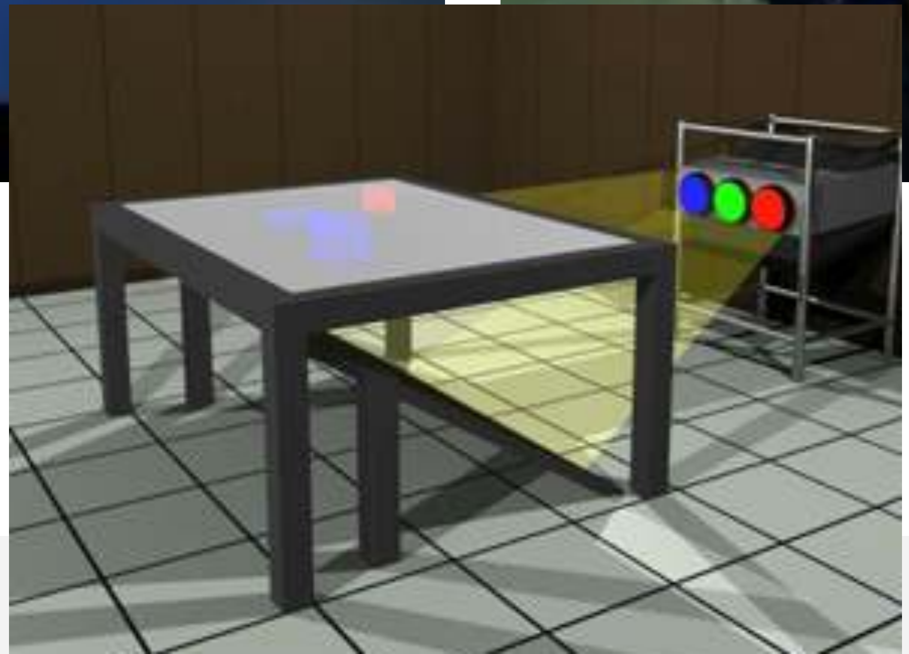

Principle of the workbench

The Cave

3-wall cave

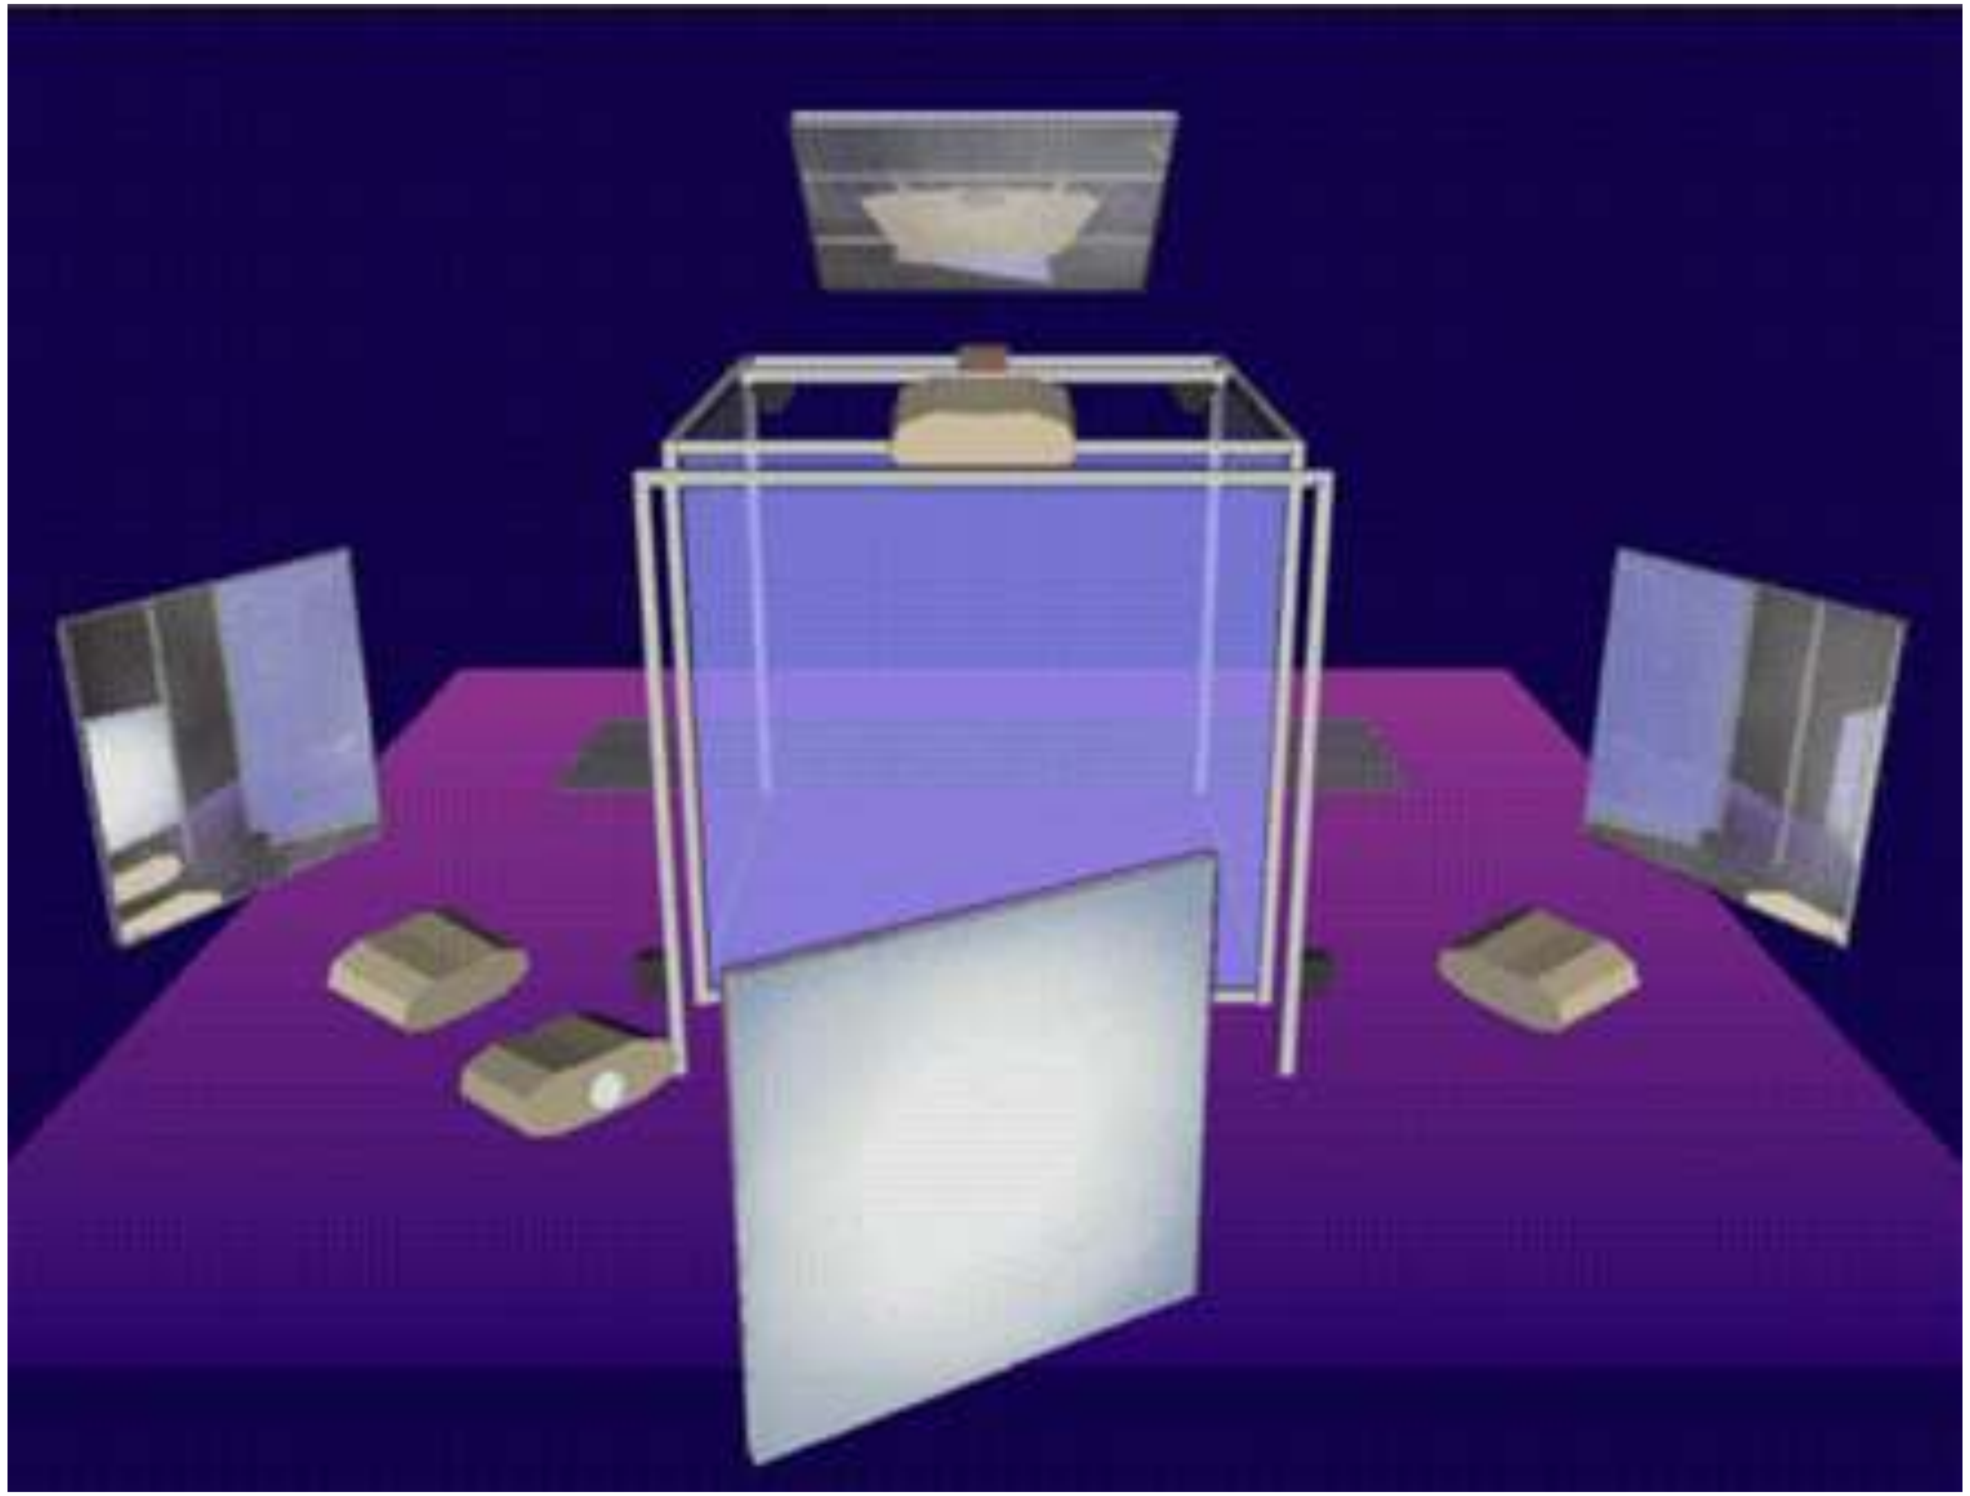

Schematic of the arrangement of the mirrors


6-wall cave, Alborg, DK


5-wall cave, FhG-IGD, Darmstadt

5-sided CAVE at University RWTH Aachen

Disney Imagineering's DISH

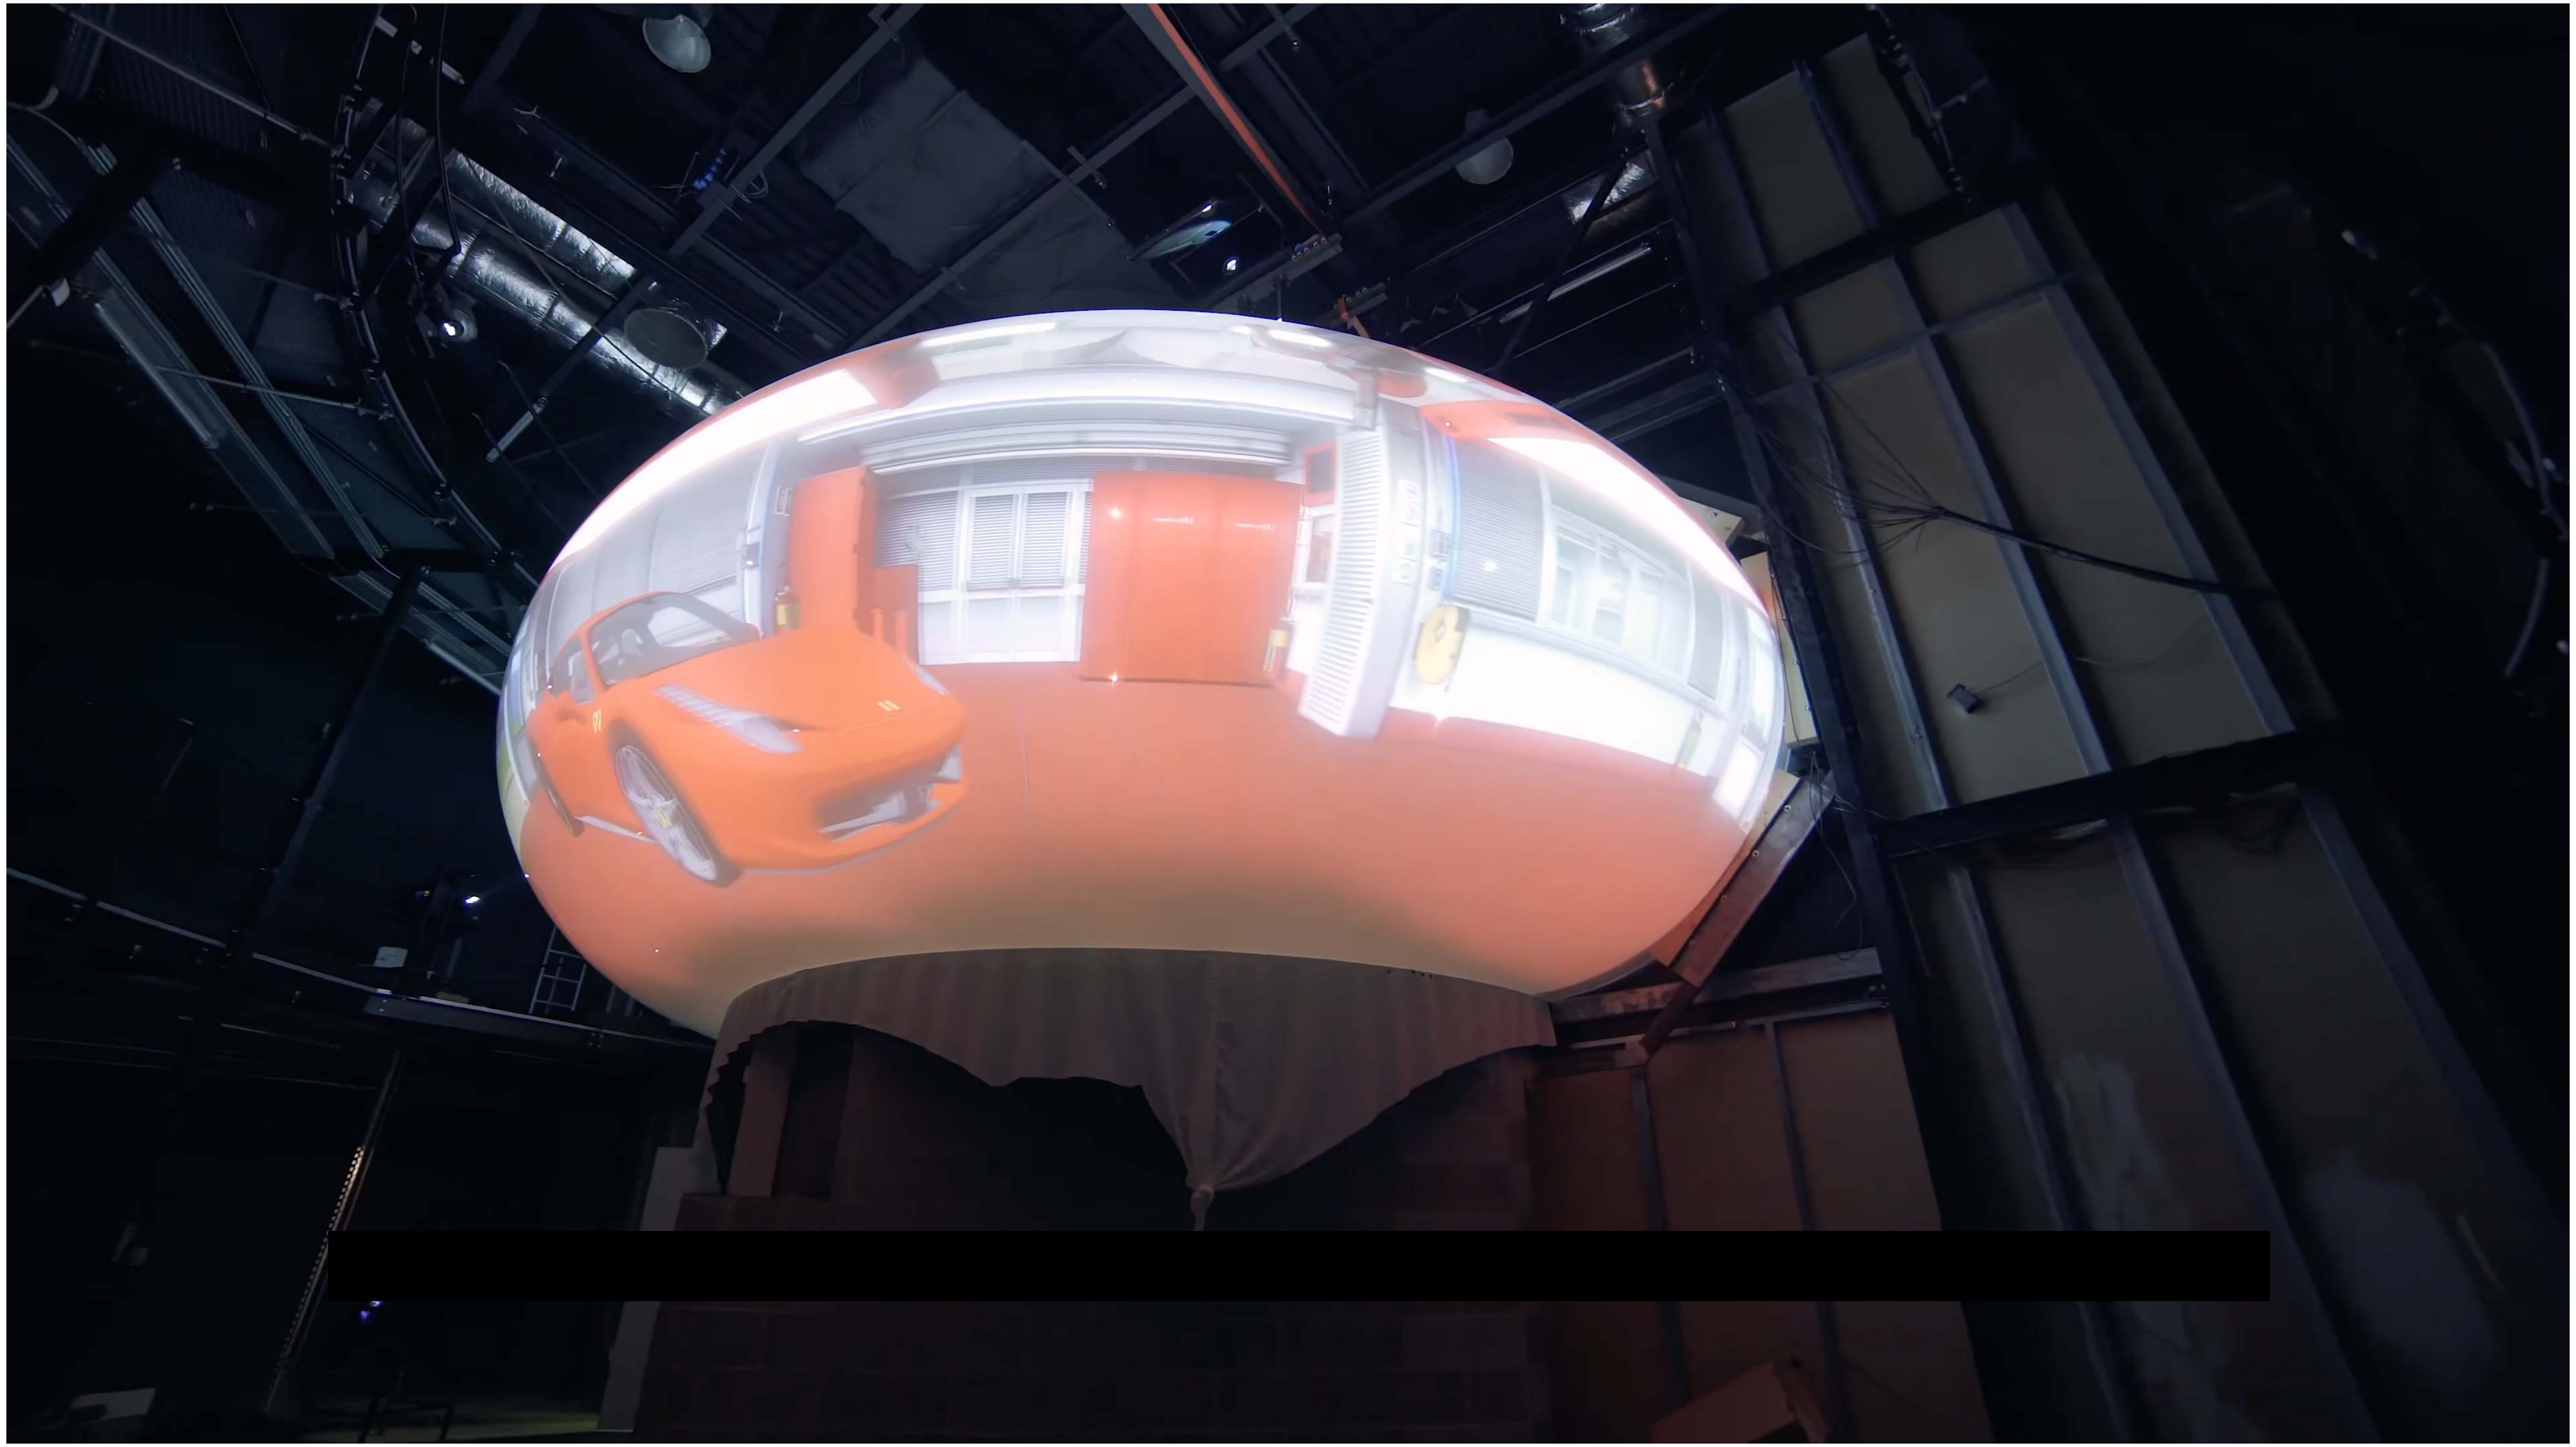

TORE, University of Lille, France

Surface = half of the outside of a torus ⟶ same curvature everywhere and 180° horizontally and vertically

RealityDeck - Immersive Giga-Pixel Display

- 308 x 30" LCD displays
- 2560x1600 resolution per display
- 1.5 Giga pixels of resolution in total
- 40'x30'x11' physical dimensions
- 85 dual quad-core, dual-GPU cluster nodes

http://www.cs.stonybrook.edu/~realitydeck/

Curved Screens

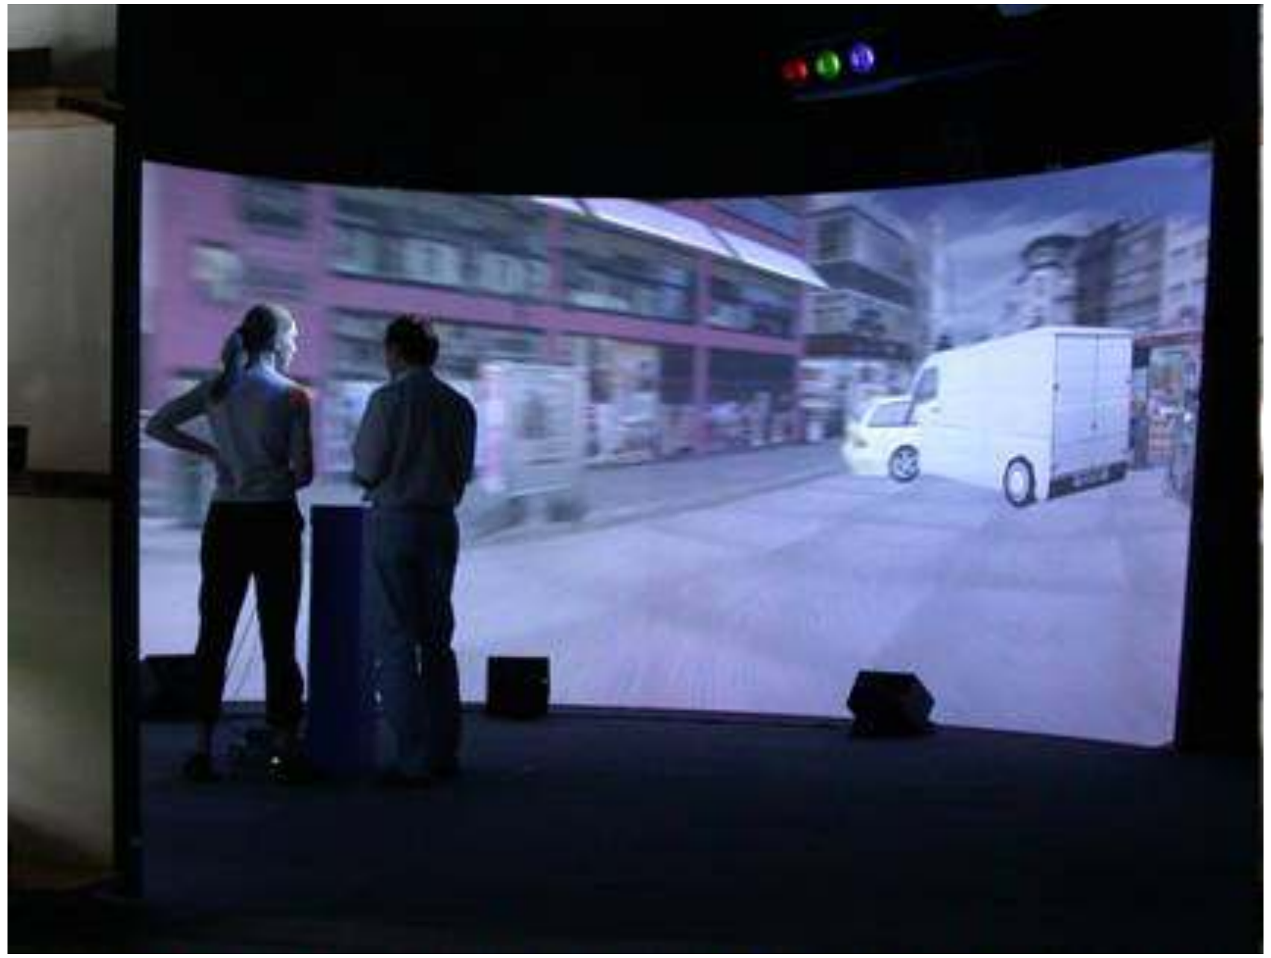

- Usually, with wall-sized screens (curved or not), some kind of edge blending and color correction between projectors is necessary

Curved Screen made out of 3D-TVs

- Idea: construct the walls of a Cave / curved powerwall out of a (small) number of 3D TVs
- Advantage: reconfiguration is relatively easily (just put the walls on wheels)


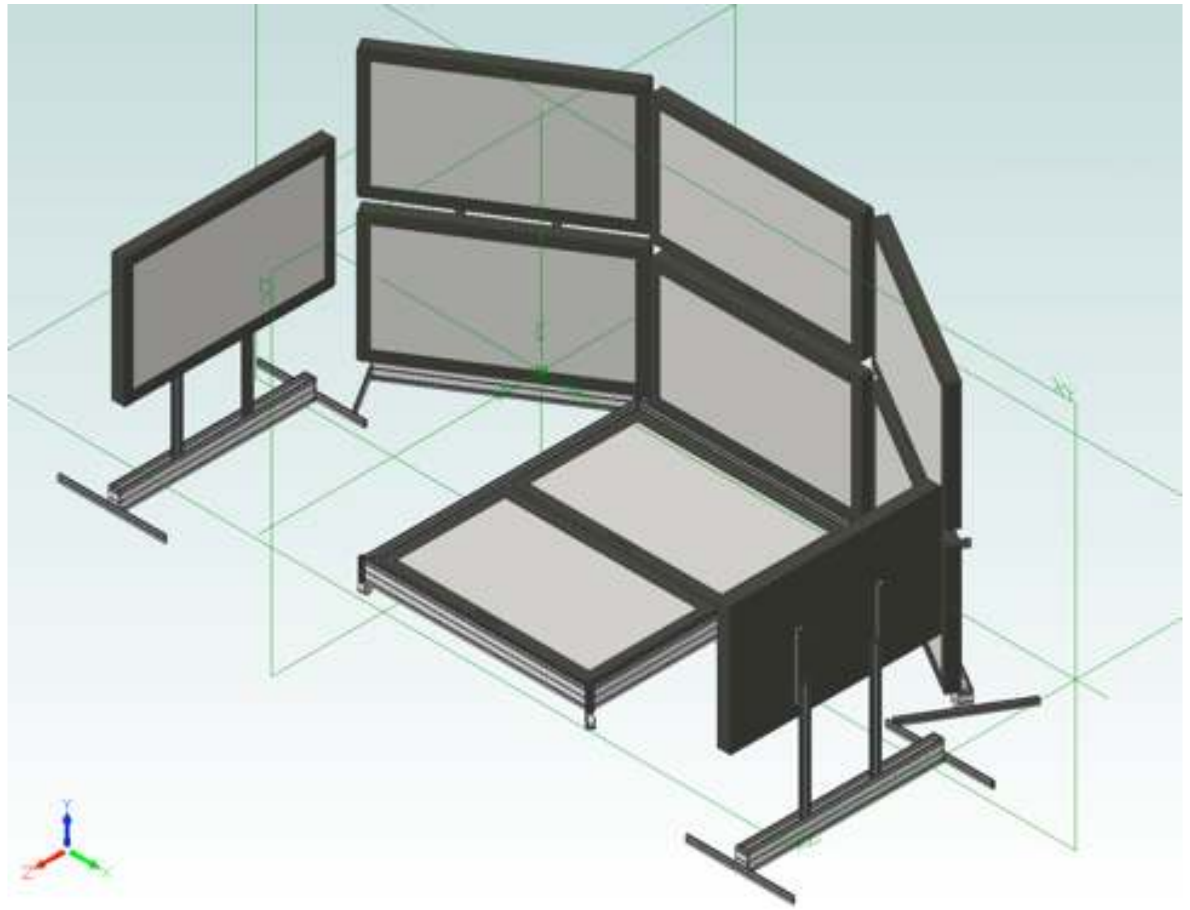
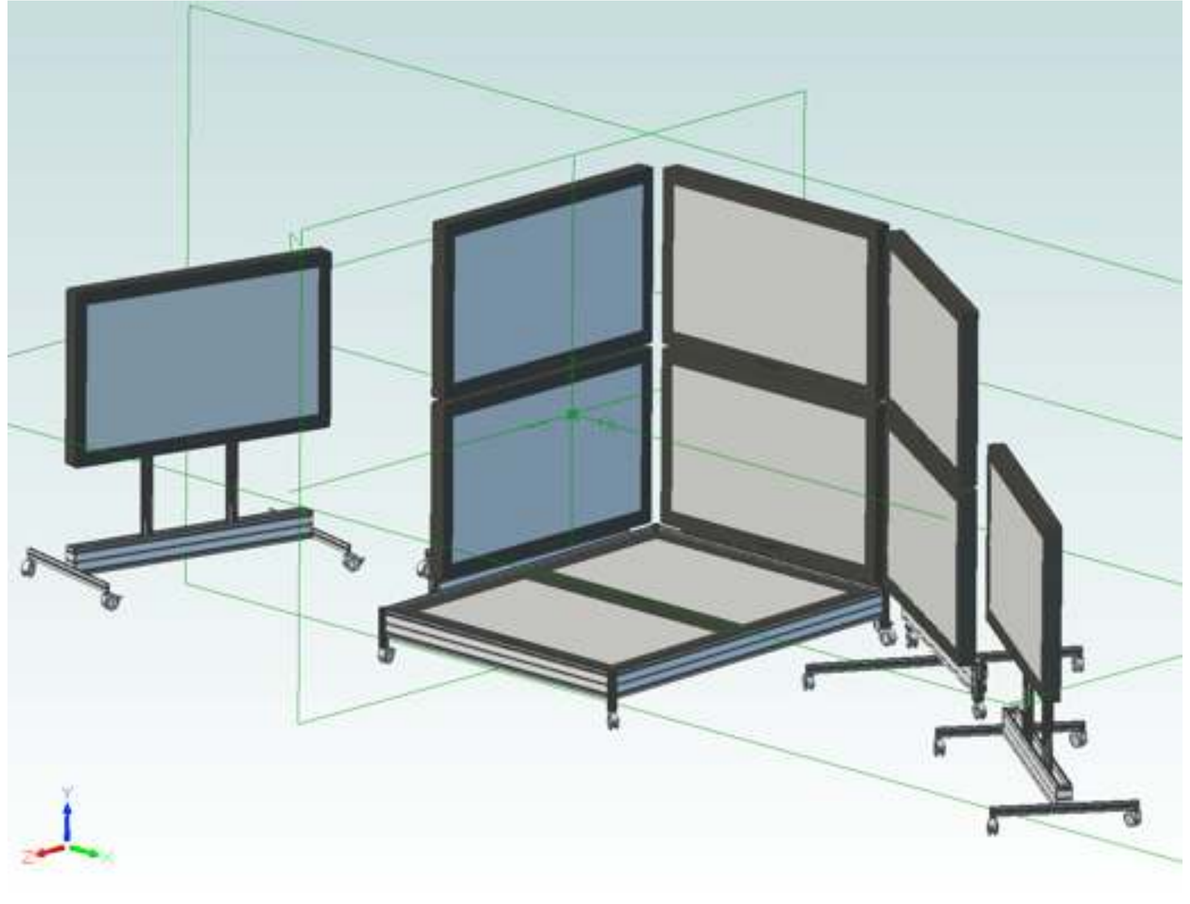

Personal Domes

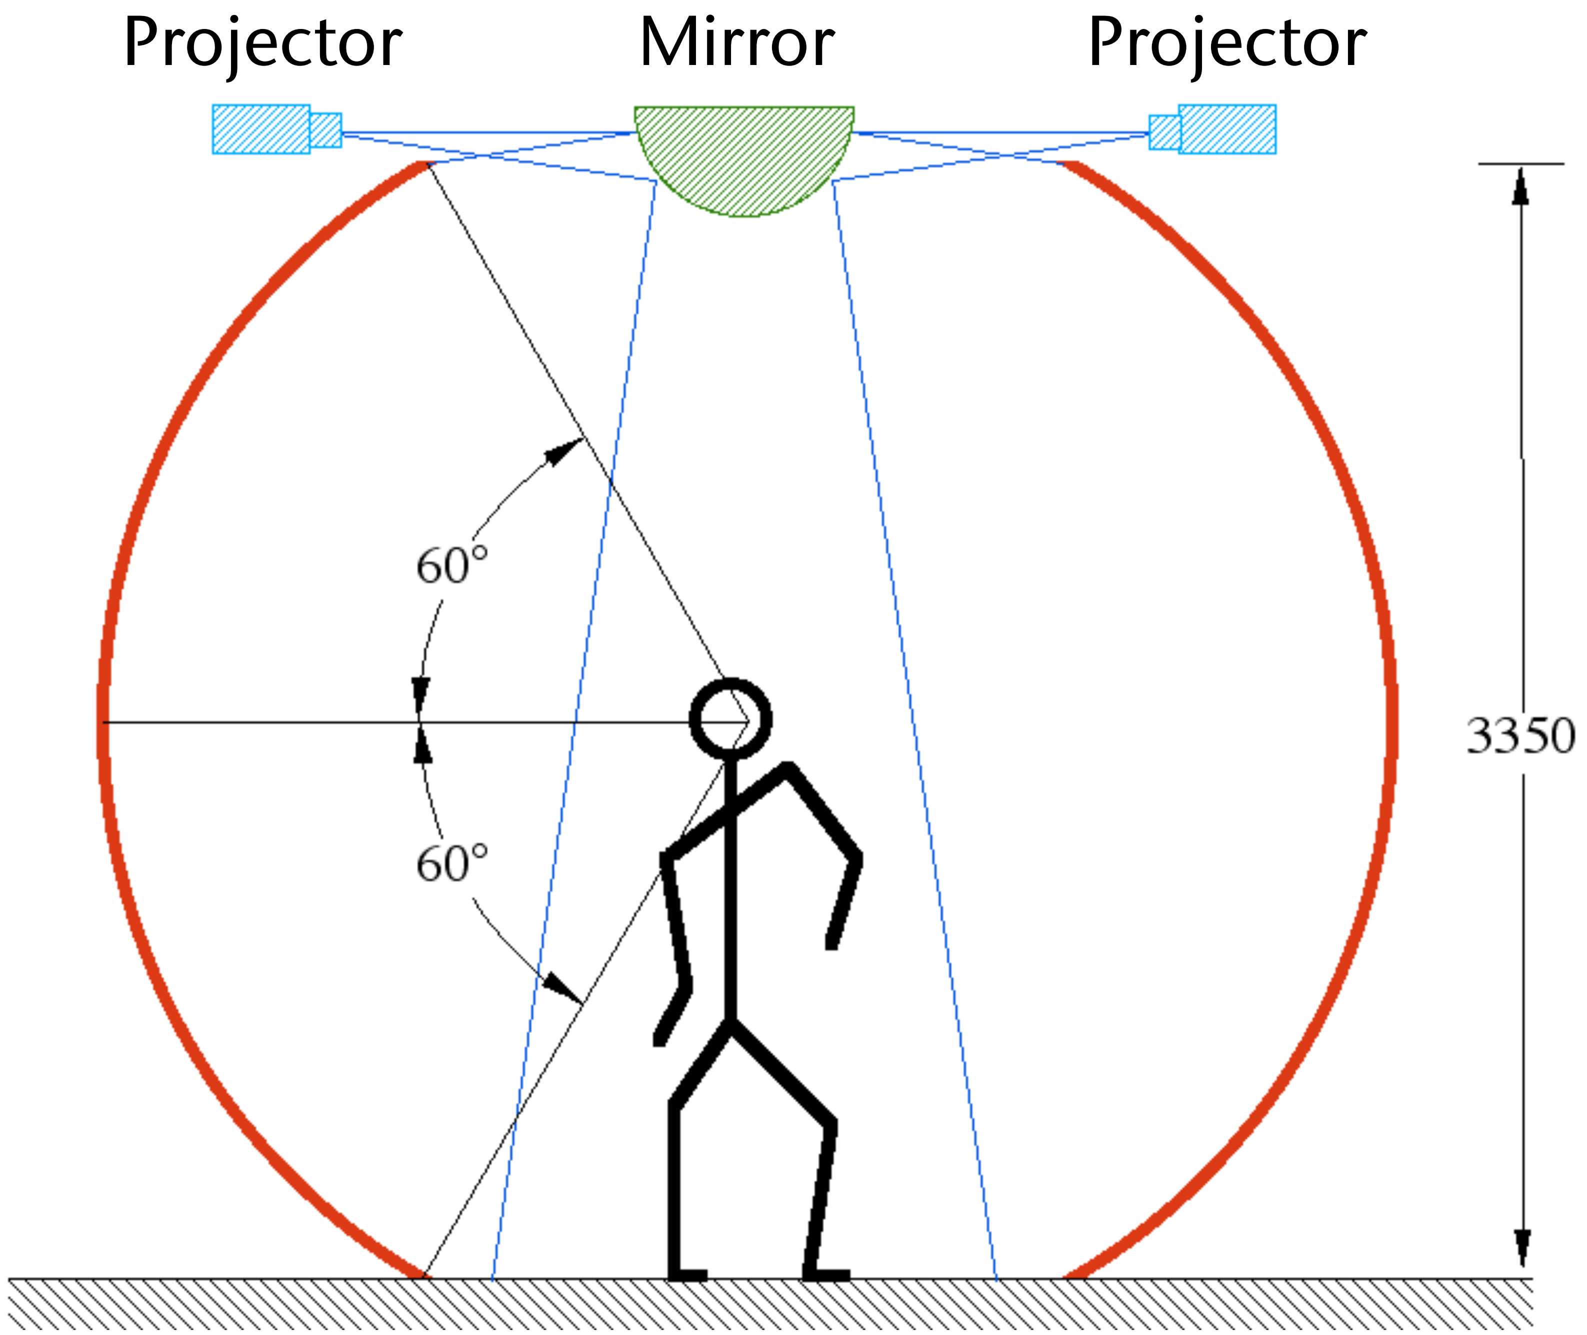

Source: Paul Bourke, University of Western Australia, http://local.wasp.uwa.edu.au/~pbourke/


GIZMODO

A Modern "Sensorama"

Immersa-Dome from Aardvark Applications

Advantages and Disadvantages of IPTs

- Advantages:
  - Large resolution
  - Large *field-of-view*
  - "Non-invasive"
  - No isolation of the real world
  - Can accomodate several users who can see each other
  - *Cave*: turning the head results in small changes of the images
    - *Problem of latency is reduced / not so prominent*
- Disadvantages:
  - Size
  - Price (lots of projectors, lots of graphics cards)
  - Precision, calibration
  - Potentially *stereoscopic violation*
  - Correct view only for one *viewer* (unless a massive amount of hardware is used)

Retinal Displays

- Idea:
  - Use the human retina as the display surface directly (all images from the outer world end up there anyway)
  - Use a laser to write the image by scanlines into the eye
- Advantages:
  - High contrasts, high brightness
  - Good for *see-through* displays, bad for VR

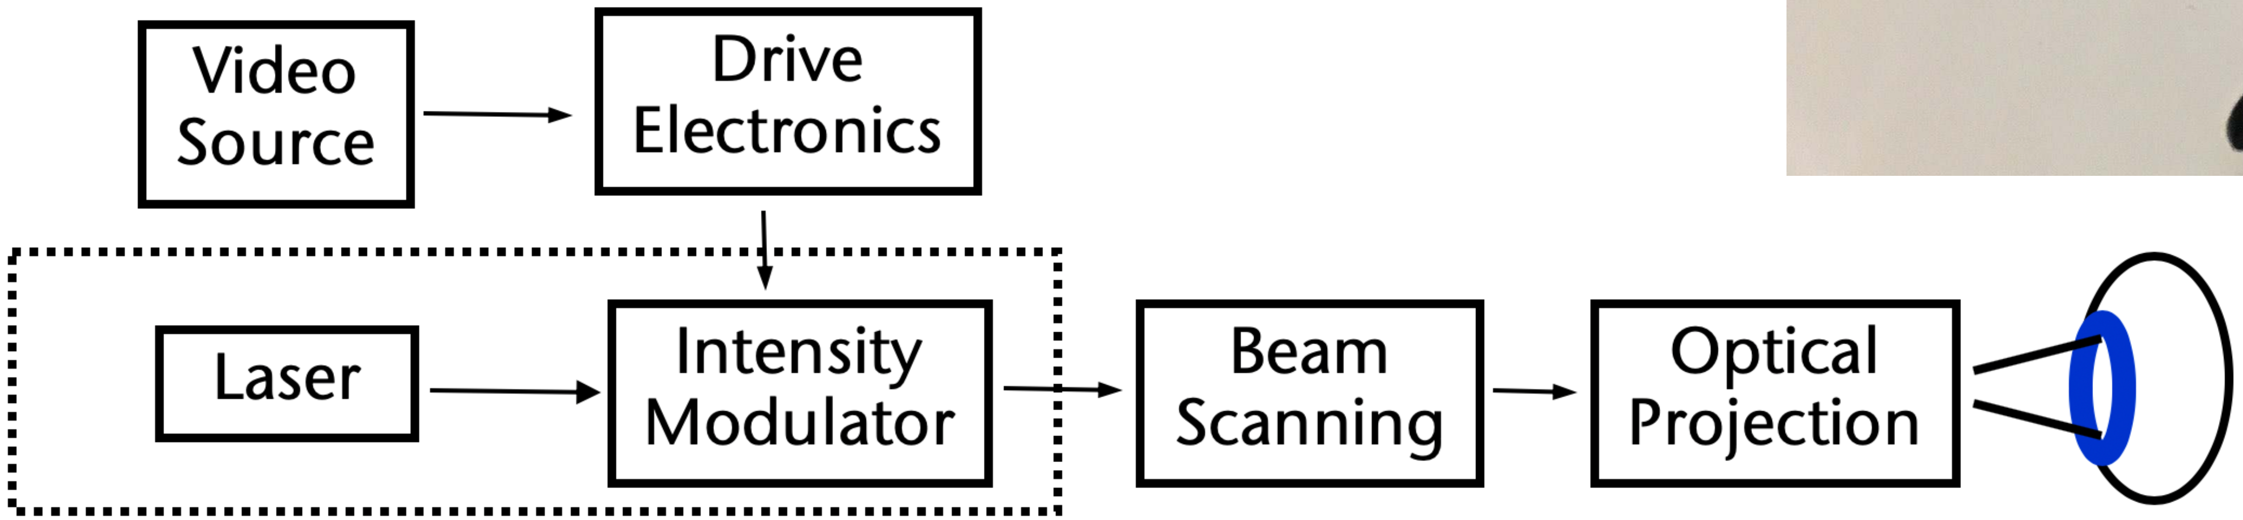

Holographic / Volumetric / POV Displays

- Volumetric = real 3-dimensional image, i.e., image that occupies a 3D space
- Usually based on a rotating disk (utilizing Persistence of Vision)
- Example volumetric display:
  - 198 x 768 x 768 ≈ 100 million voxels
  - Frame rate: 20 Hz
- Hologram / Holographic display: a 2D display that can recreate a light field as if it was coming from the real 3D object
-

Example of Volumetric Display

Voxon

- Advantages:
  - Provide correct perspective/view from *every* angle!
  - Coherence between accomodation and convergence
  - Depth of field (Tiefen(un-)schärfe)
- Holographic displays: algorithmic computation of holograms
- Problems:
  - Staggering amount of computational work
  - Colors
- Volumetric displays: voxels are projected onto a a rapidly rotating surface covering a volume
- A.k.a.: Persistence of Vision Displays
- Problems:
  - Size of data (e.g. 100 mega-voxels = 1000x1000x100 display resolution)
  - Occlusions?

Mid-Air Volumetric Displays

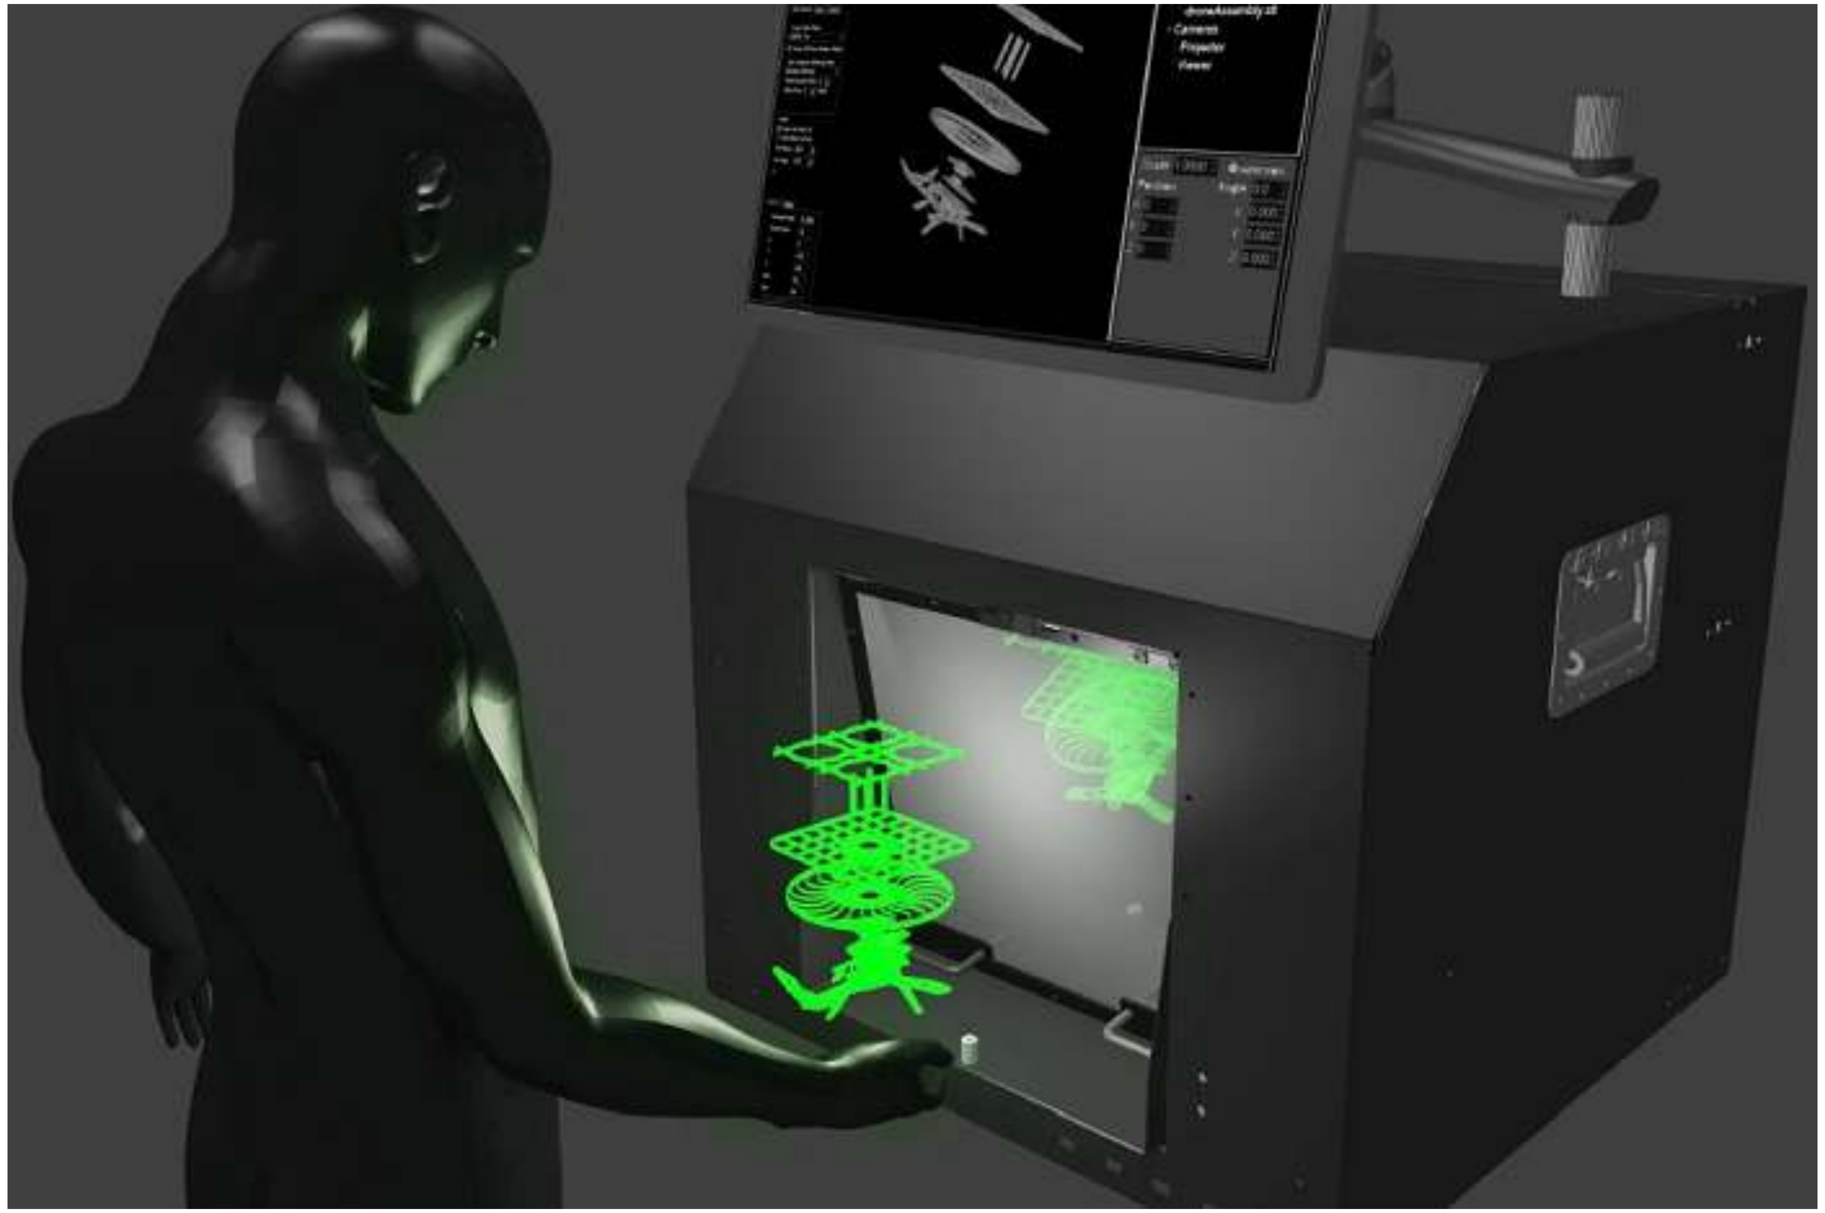

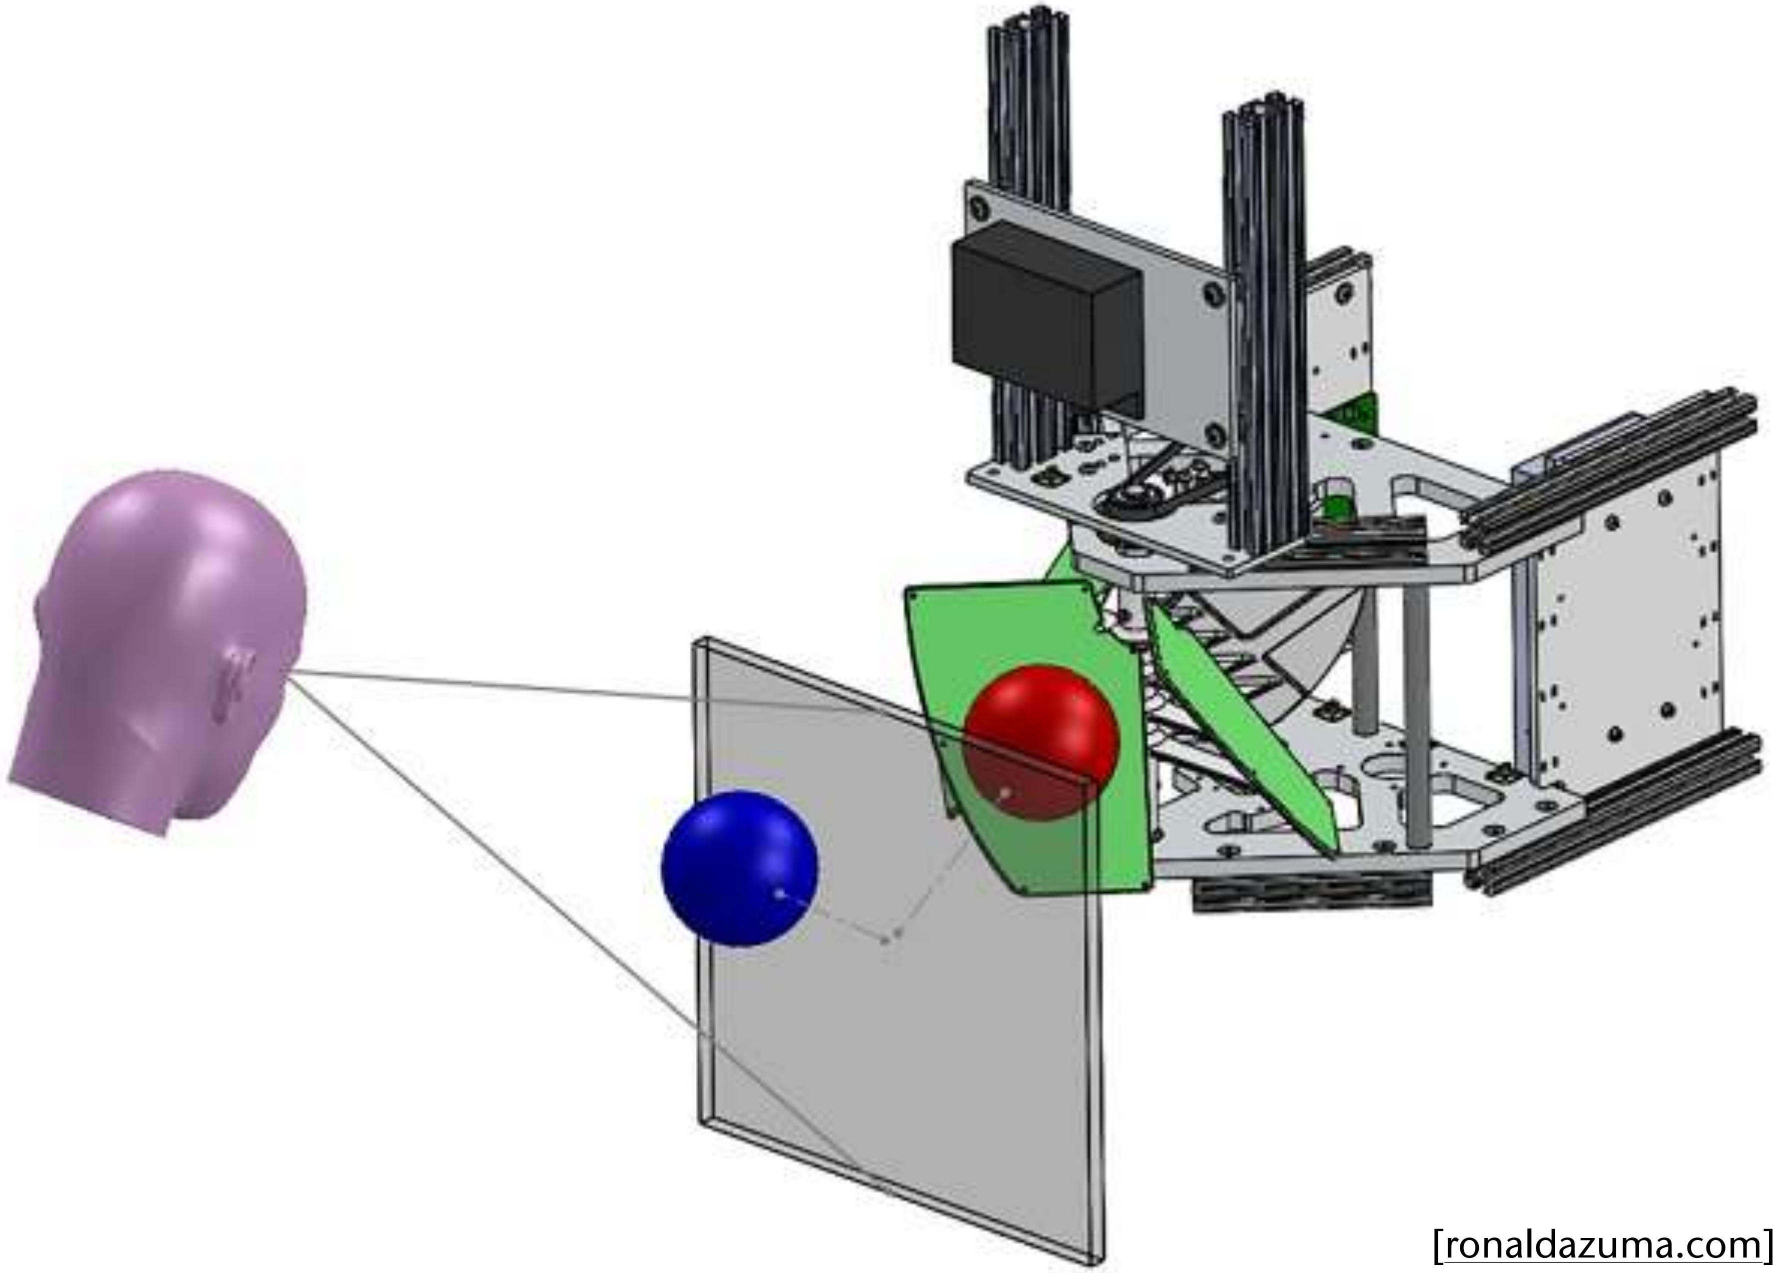

[ronaldazuma.com]

The display is not a holographic display. Instead, it is an optically re-imaged volumetric display.
A fast projector is synchronized with the revolving planes. At each instant, the projector illuminates the parts of the plane that intersect the virtual 3D object. This generates a true 3D hologram.
How does this appear in front of the display? We achieve that by combining the volumetric display with an *optical reimaging glass* made of large numbers of tiny mirror elements. These create a real optical image in front of the glass.

Bremen

Not a Hologram!

Lightfield Displays

Automultiscopic Display

- Like a lightfield/holographic display, but views (i.e., perceived images) differ only along horizontal viewpoint changes

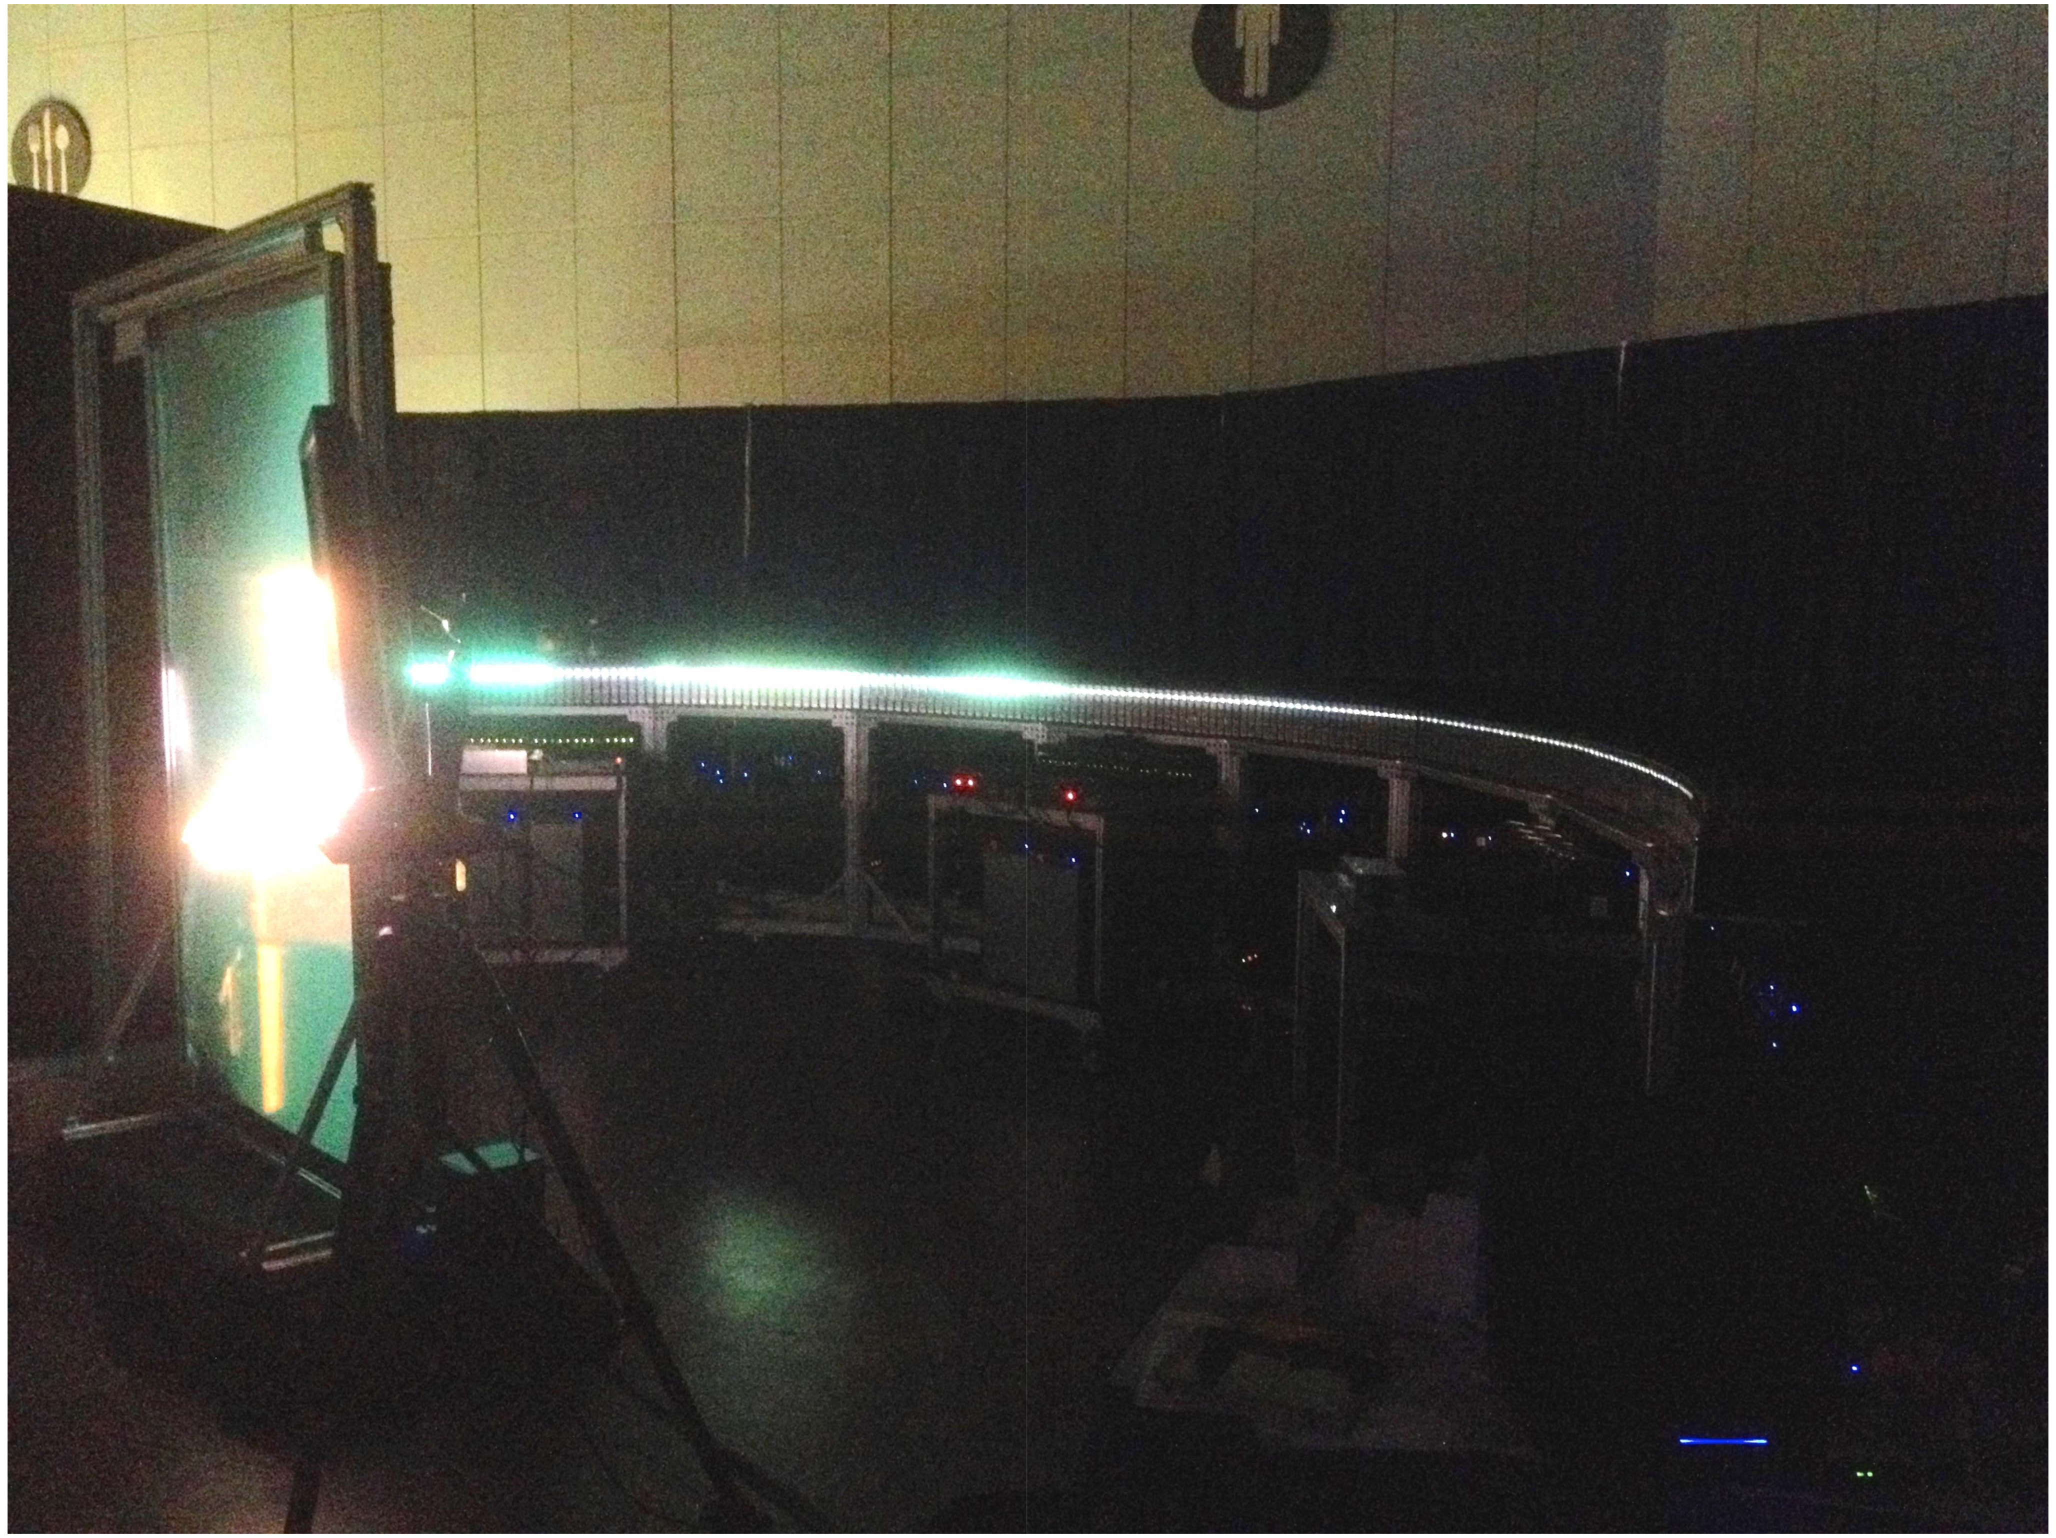

- A special screen transmits images from from each projector *only* in one direction with a very small scattering angle (1°)

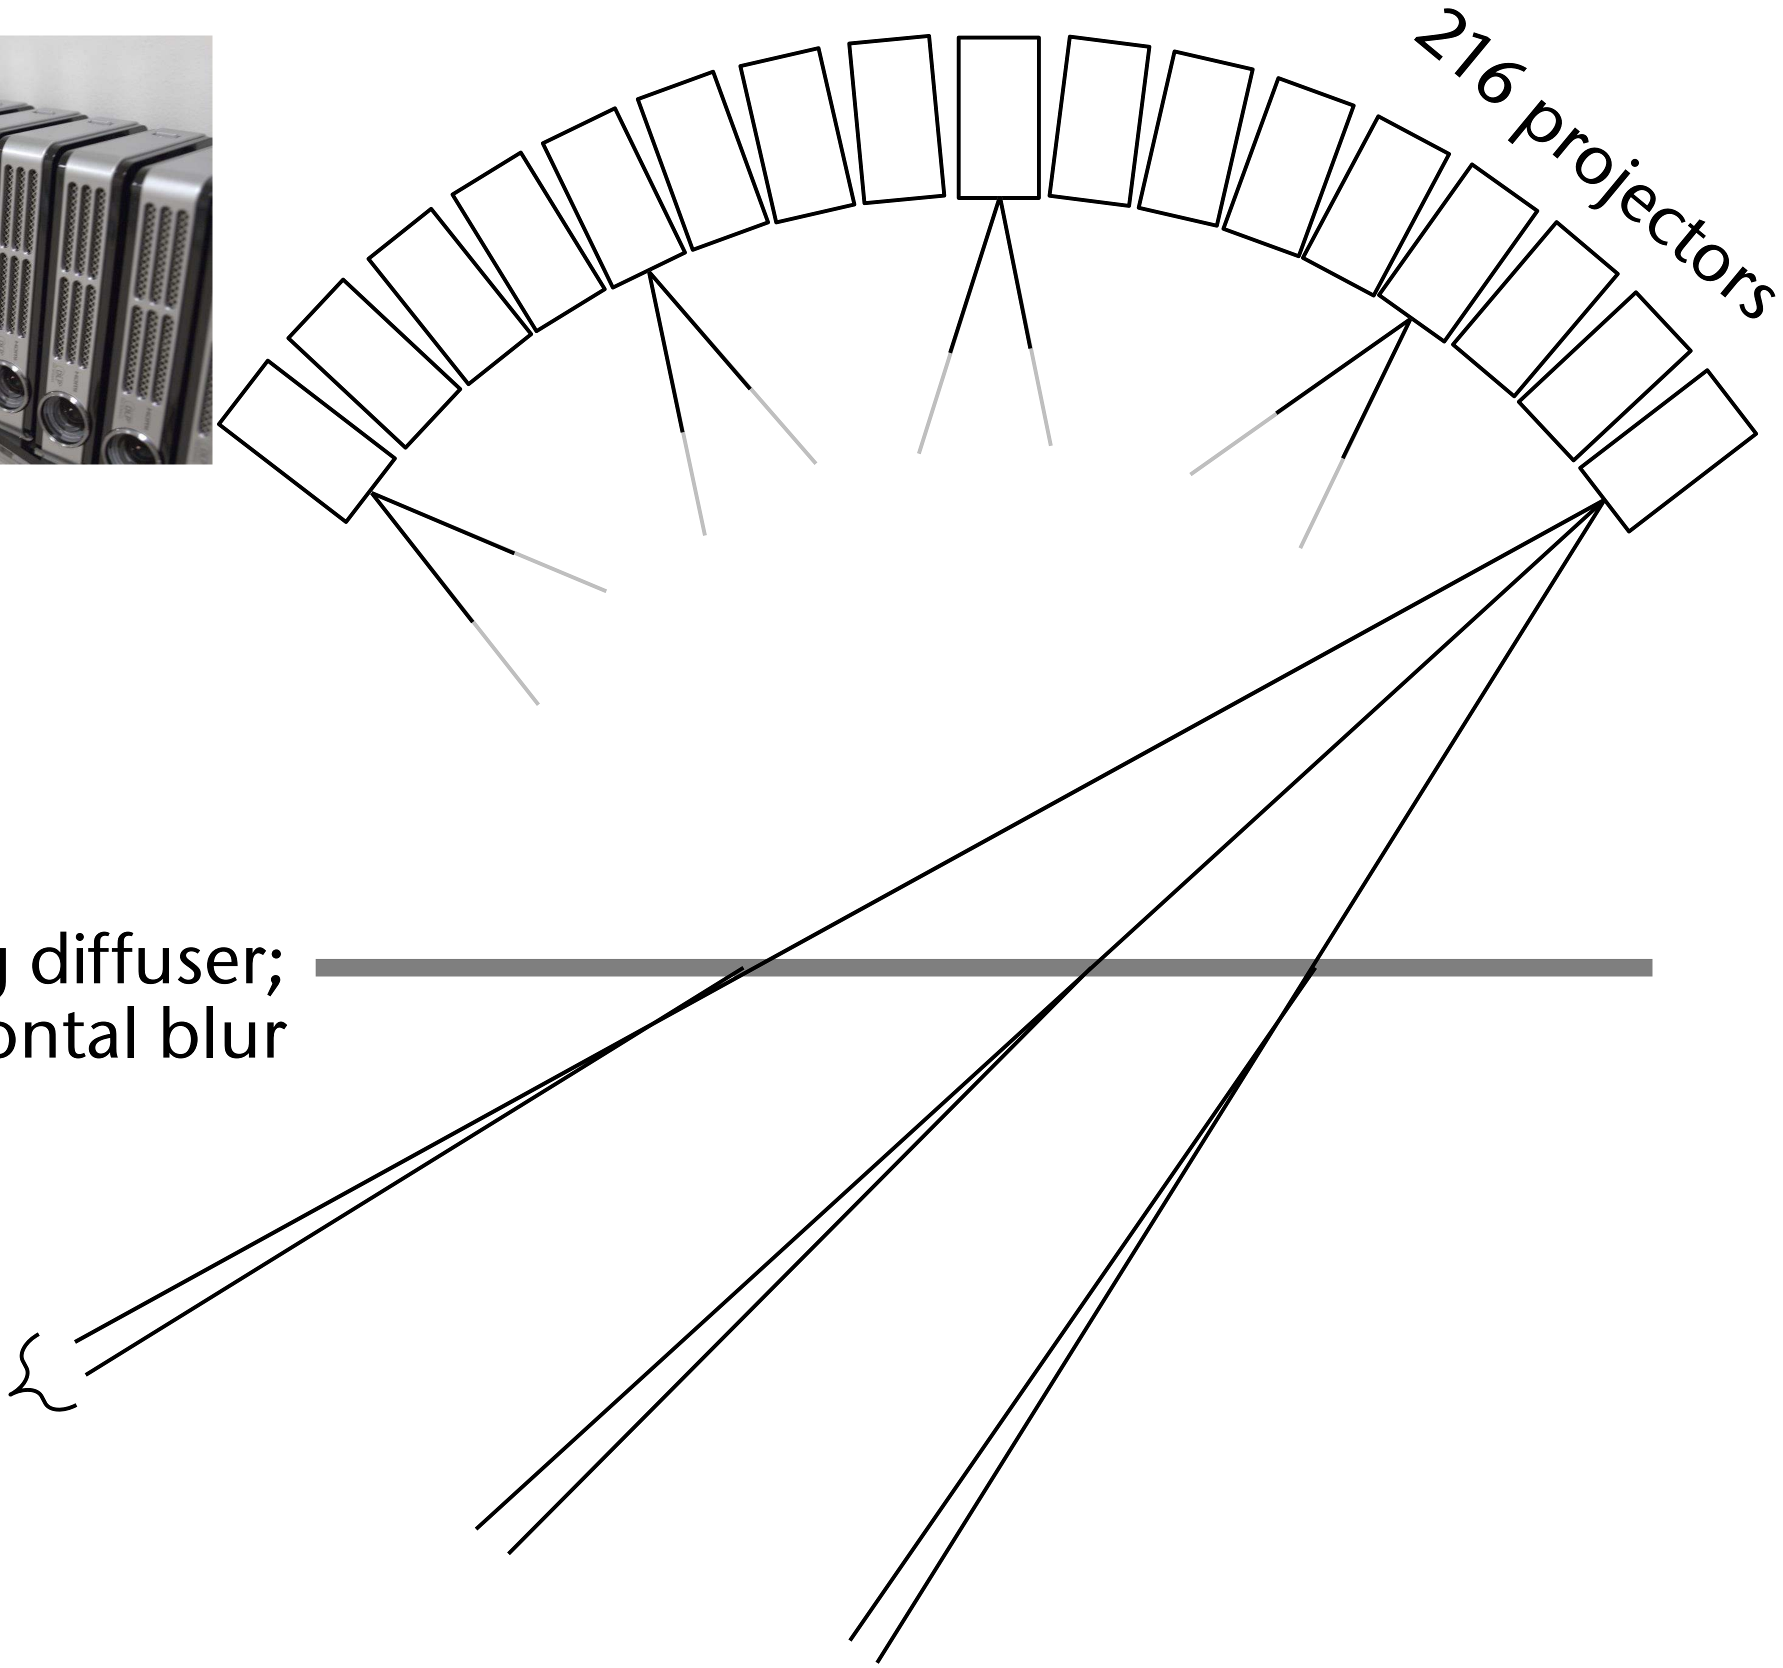

Screen = anisotropic light shaping diffuser;
scatters light vertically, maintains narrow horizontal blur

1° degree horizontal blur

Challenges

- Lenticular screen with very small horizontal diffusion angle:
  - From a specific viewing direction, the light from a single projector appears as a single stripe of light

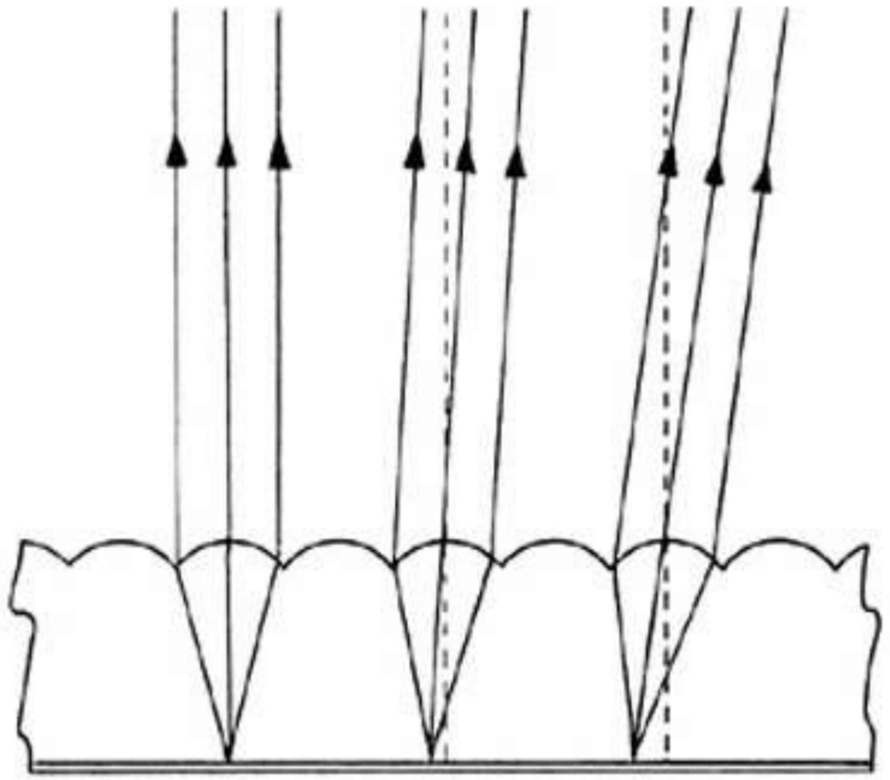

- Example bandwidth: 1920 x 1080 x 24 bits x 60 FPS x 216 cams = 80 GB/sec
- Synchronization between all GPUs (swapbuffers) and all projectors (VSYNC)
- If number of cameras < number of projectors $\longrightarrow$ video streams for "in-between" projectors must be interpolated from neighboring streams

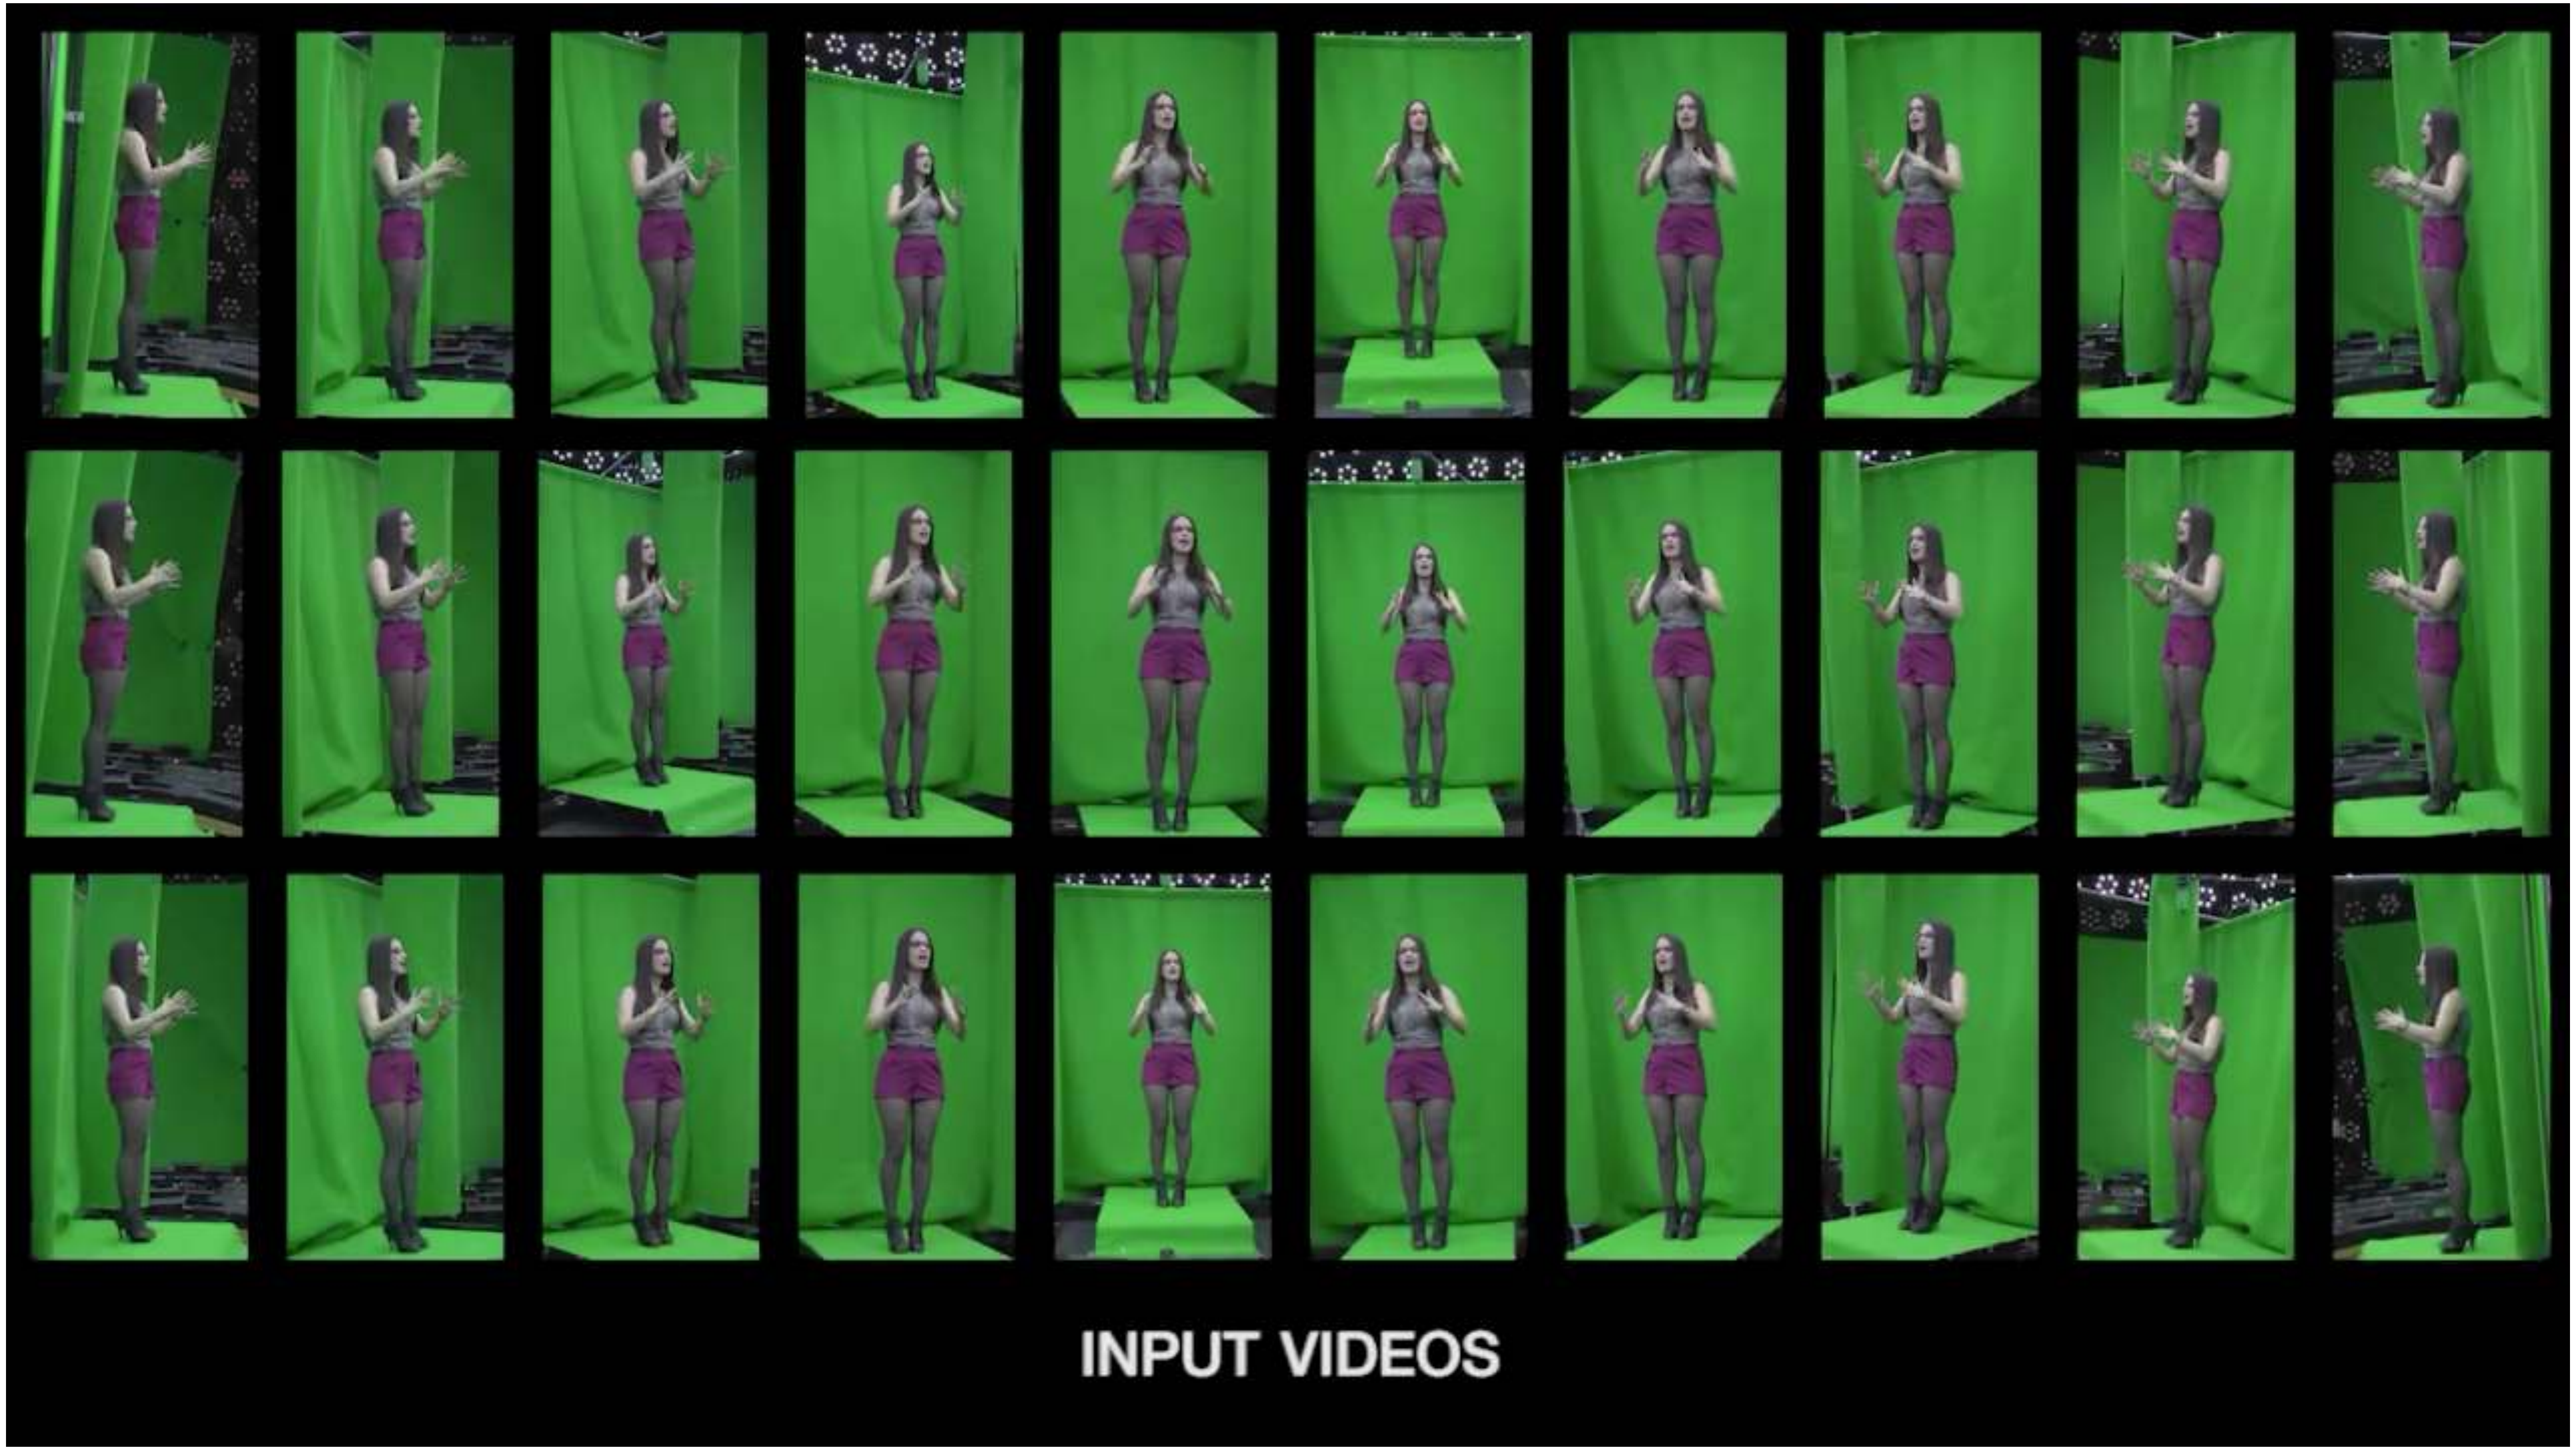
INPUT VIDEOS

- Advantage: unlimited number of viewers
- Disadvantages:
  - Expensive (lots of projectors), and needs lots of space
  - Does not work with tilted heads (eyes must be aligned with the lenticular lenses)

Unconventional Displays and Display Surfaces

Olwal et al.: Consigalo

DisplAir

- The "Janus" display of KAIST, Korea: each person on either side gets their own, possibly different image
  - Utilizes persistence of vision
  - See-through display with touch interaction for collaboration

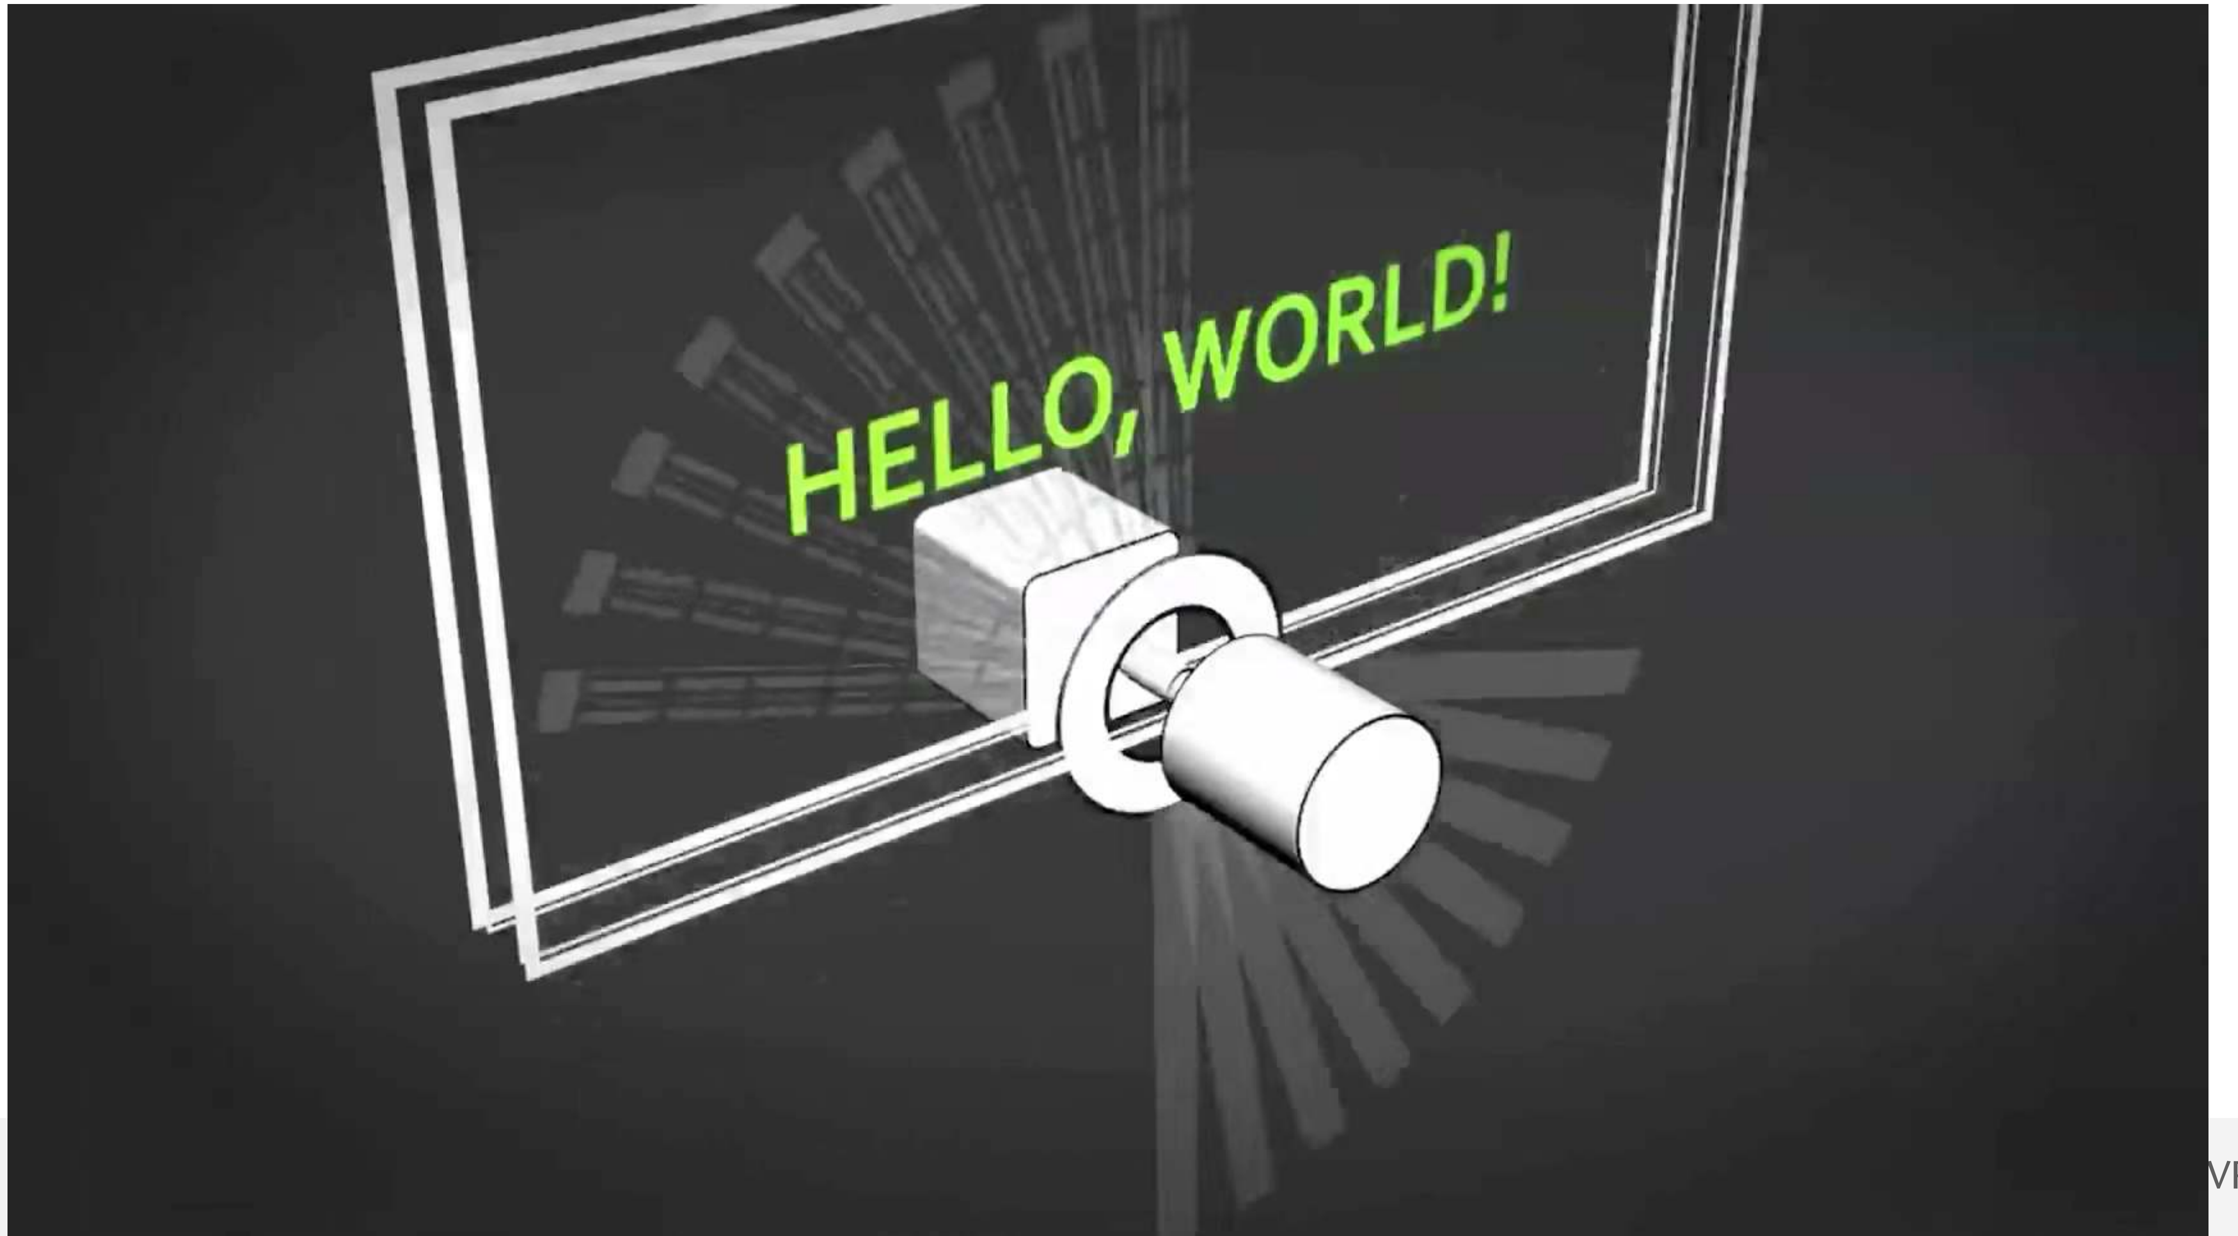

- "Everywhere displays":

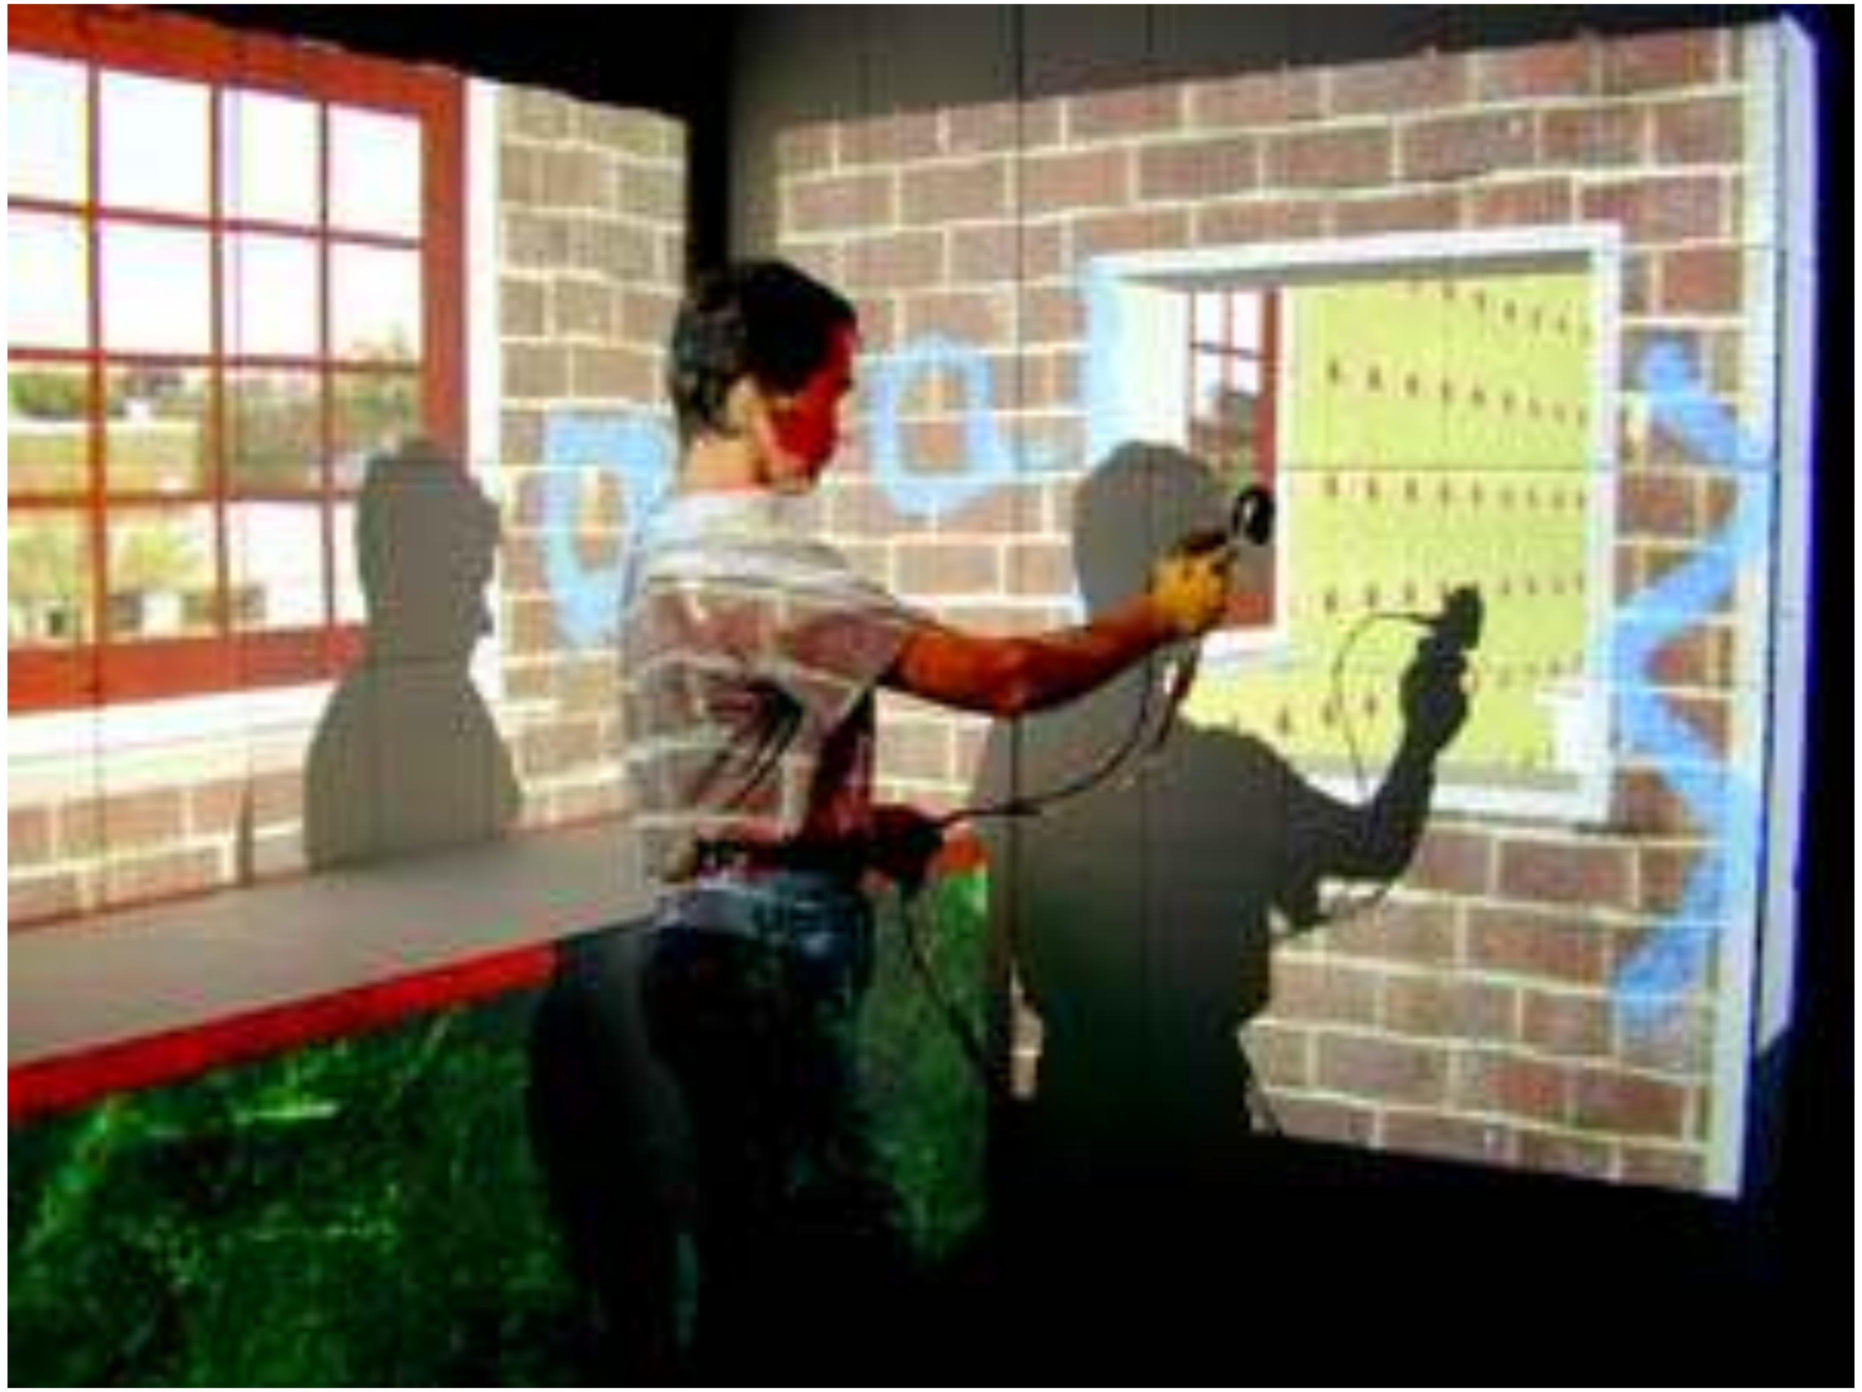

What Would be Your Favorite/Best VR Display?

https://www.menti.com/hht2ei5jxt

How to Project Stereo With Only One Display Surface?

- One channel, two senders & receivers -> need some kind of multiplexing

1. Temporal Multiplexing ("active stereo"):
   - Typically 1 projector (e.g. monitor)
   - Project/render alternating left/right image
   - Synchronously, switch left/right glass of *shutter glasses* on pass-through
   - Shutter glasses run with 120 Hz -> 60 Hz framerate

2. Multiplexing by polarization ("passive stereo"):
   - Usually 2 projectors displaying on same surface
   - Project left/right simultaneously but with different polarization of the light
   - Polarization glasses let only left/right images pass, resp.

Different Kinds of Polarization

1. Linear polarization:

- The transversal wave propagates only in one, fixed plane

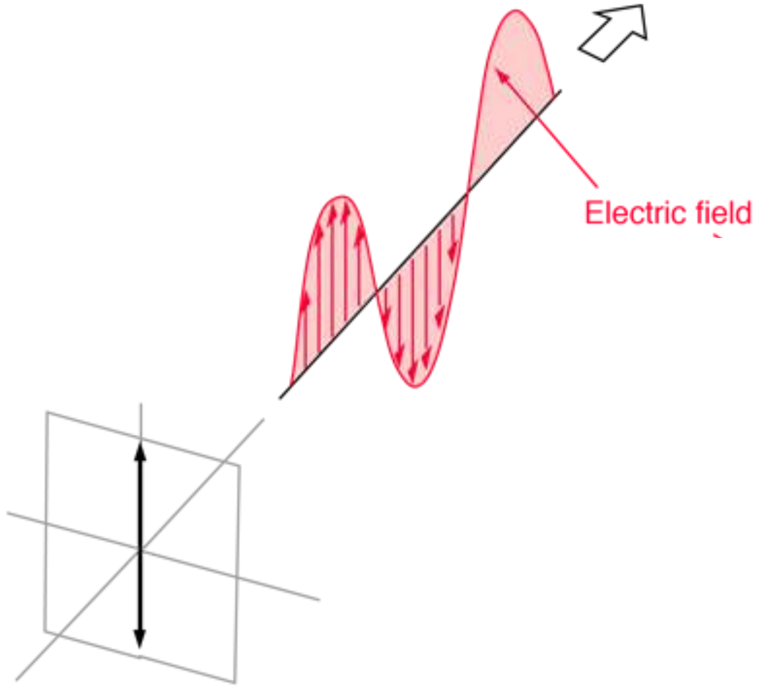

2. Circular polarization:

- Generated by two orthogonally linearly polarised waves that have a phase shift of $\lambda/4$; sum is a wave where electric vector rotates
- Left-handed / right-handed polarization

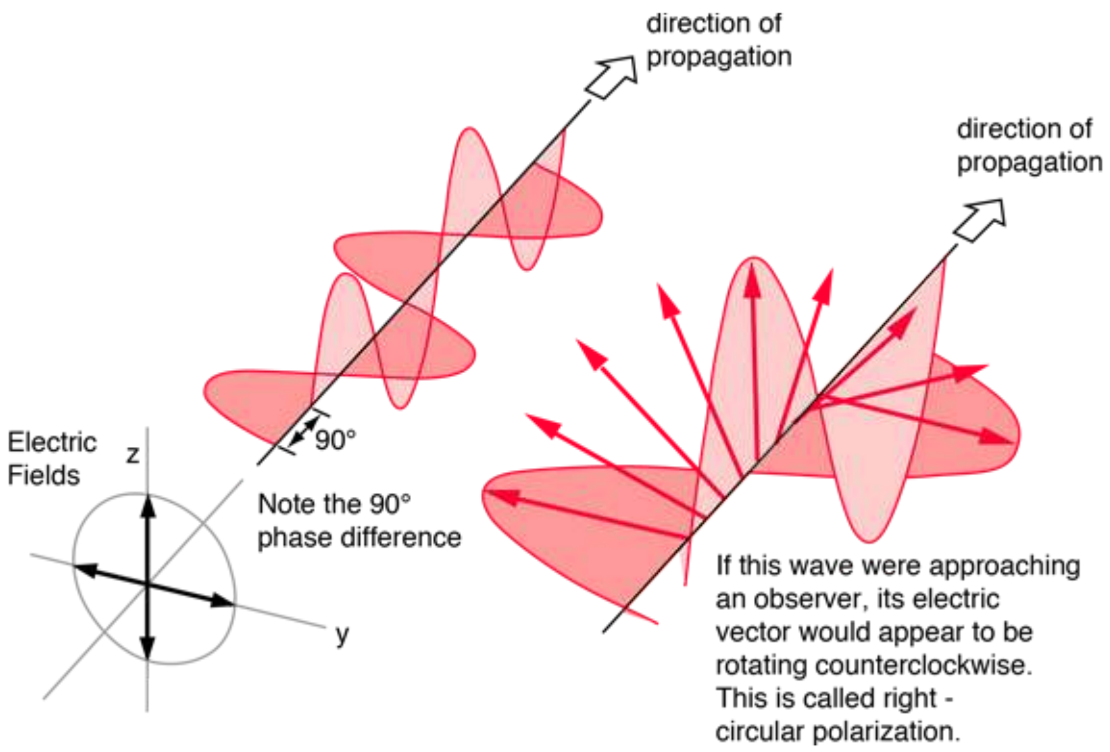

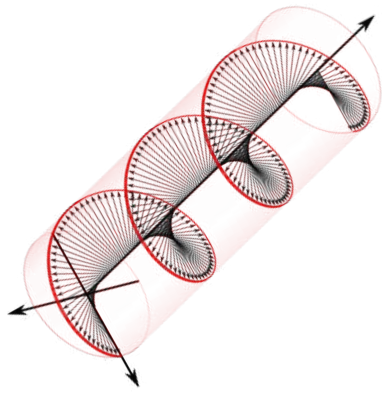

Demo

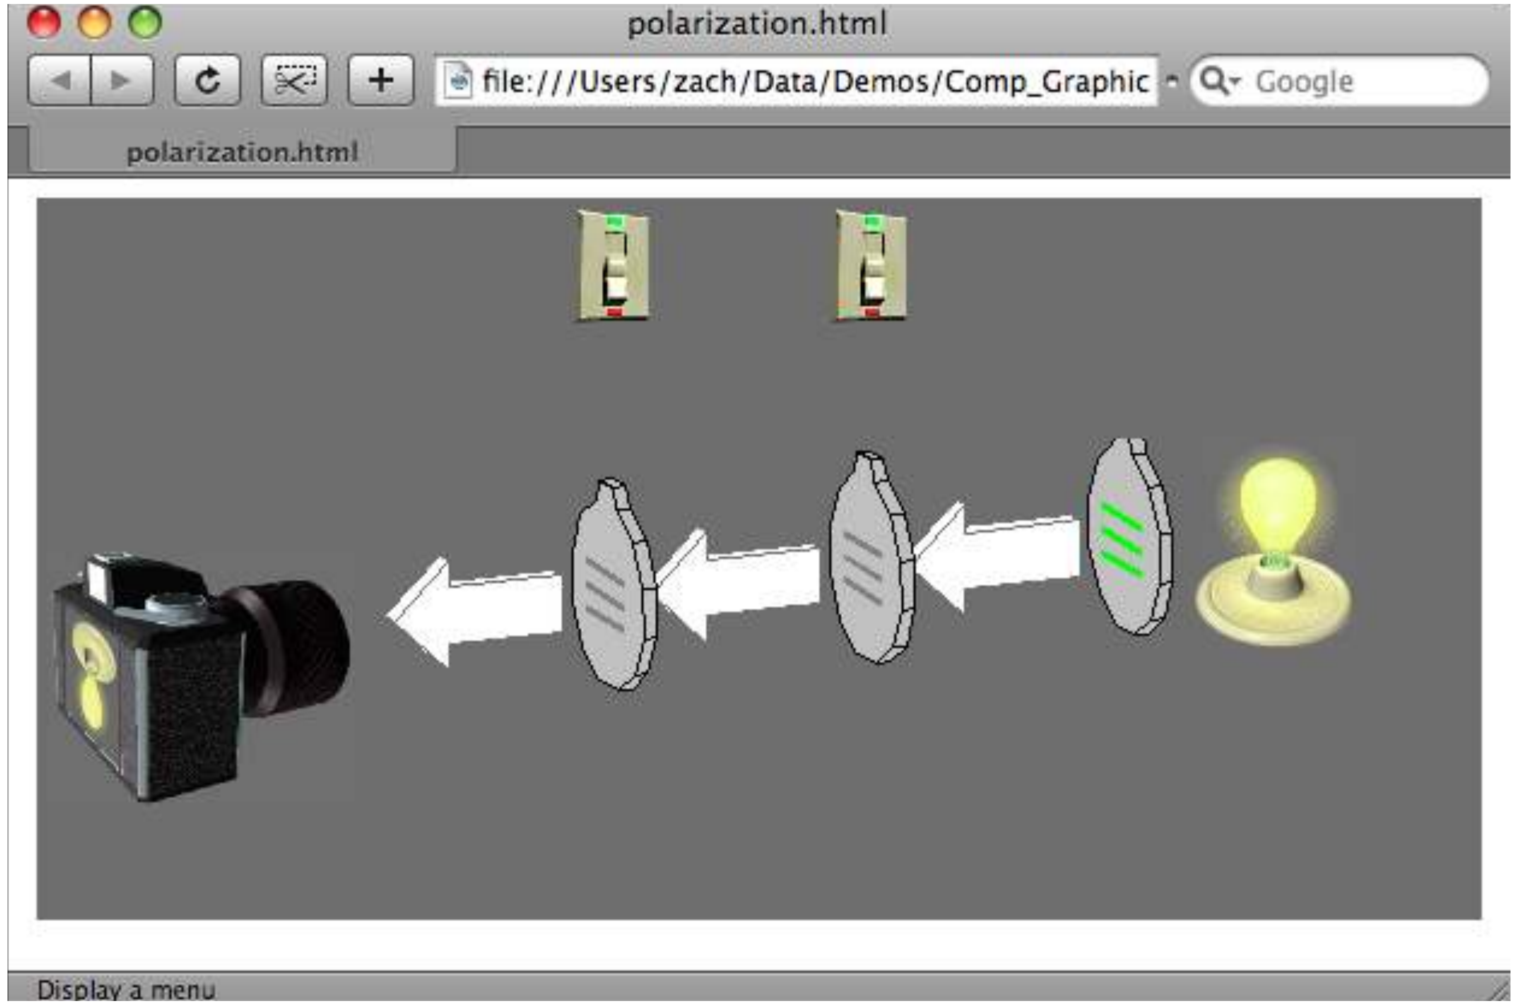

http://www.colorado.edu/physics/2000/applets/polarization.html

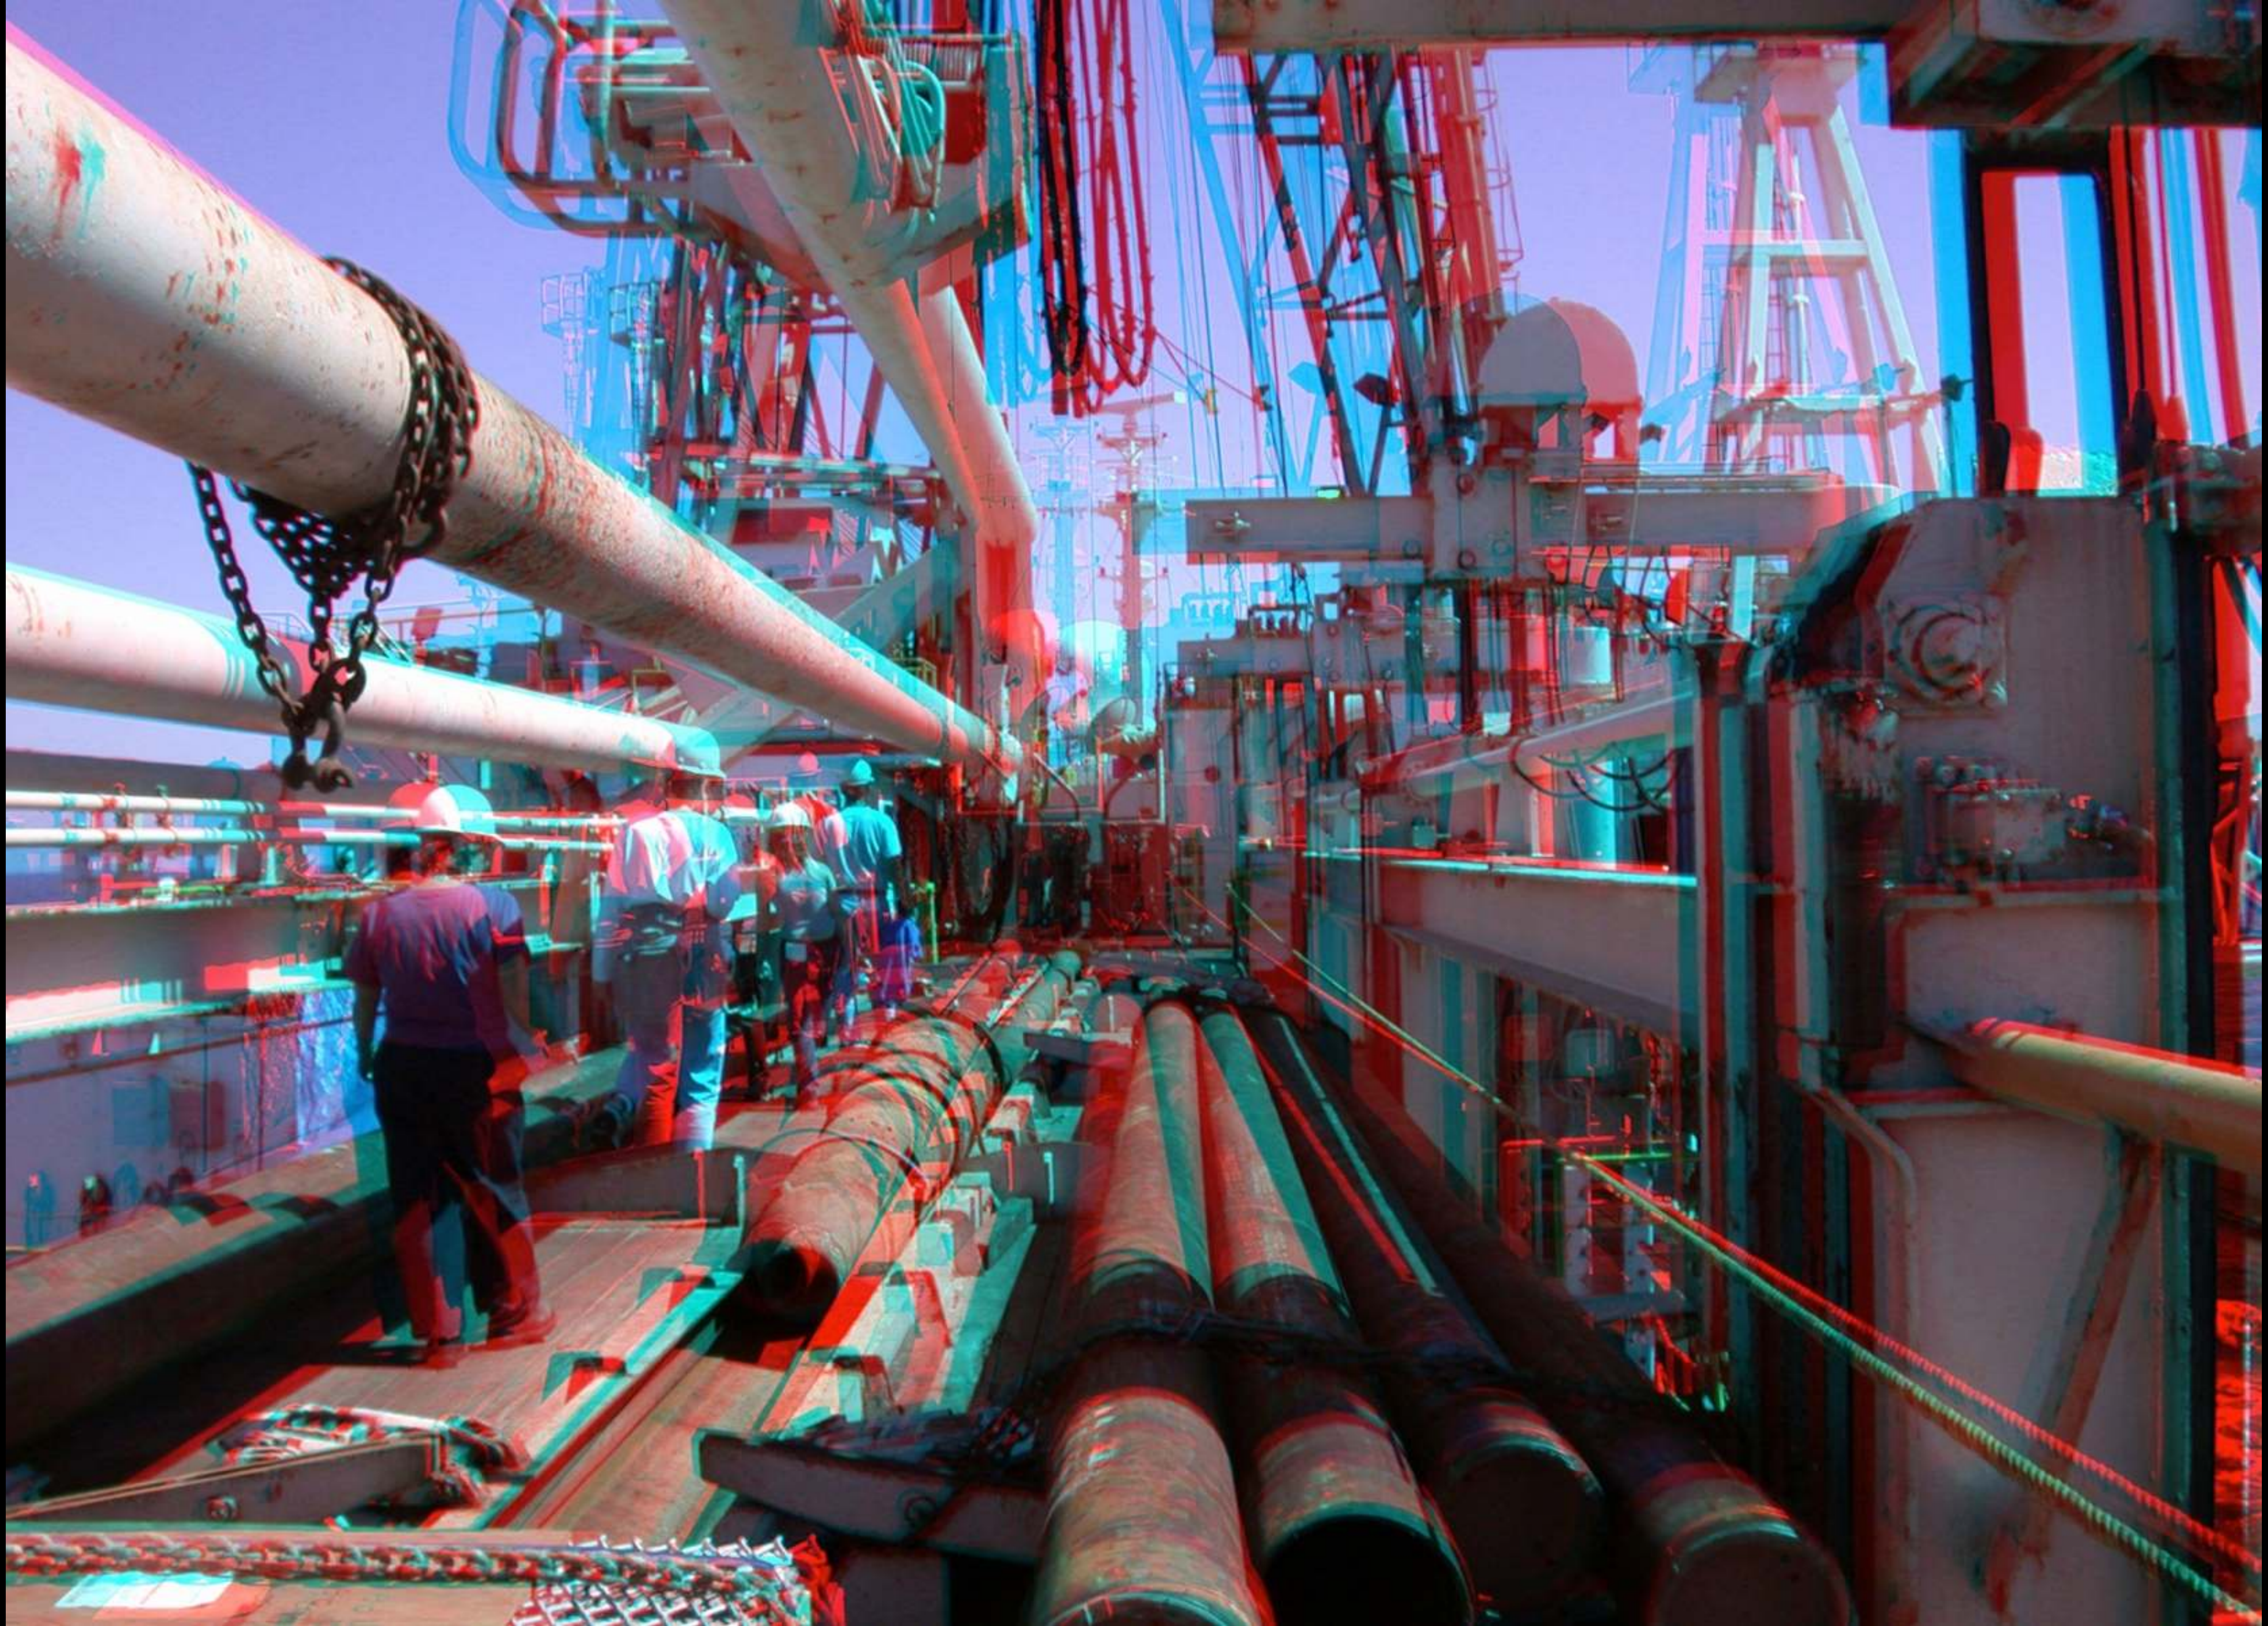


RESOLUTION

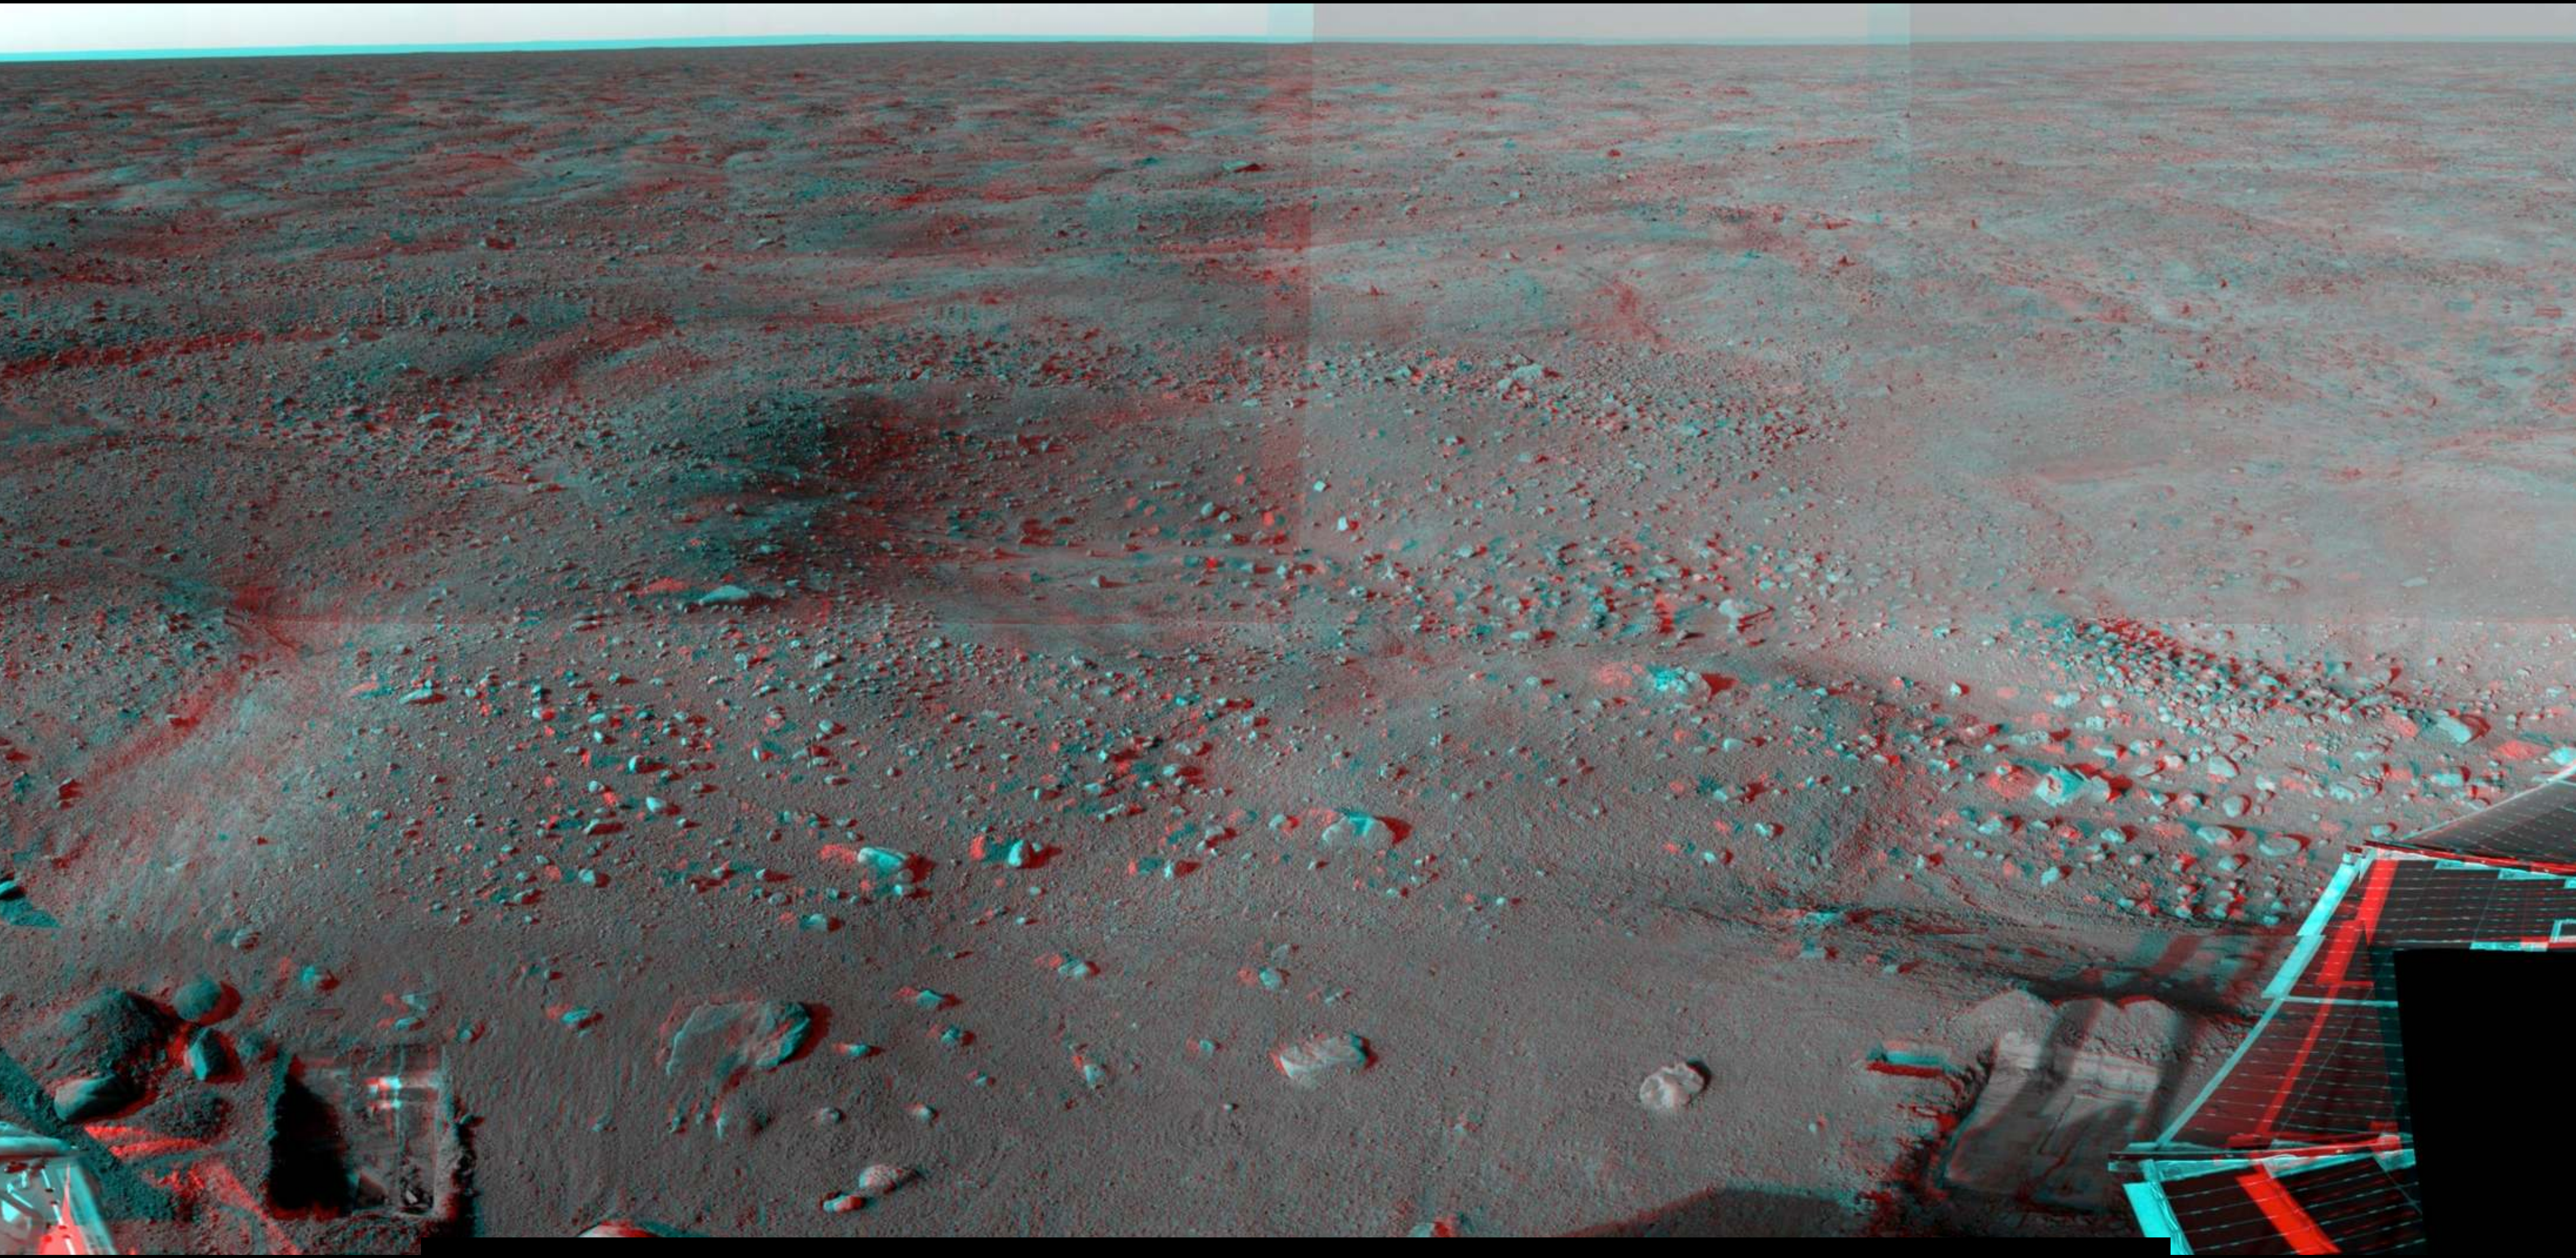

Scrat's CONTINENTAL CRACK-UP

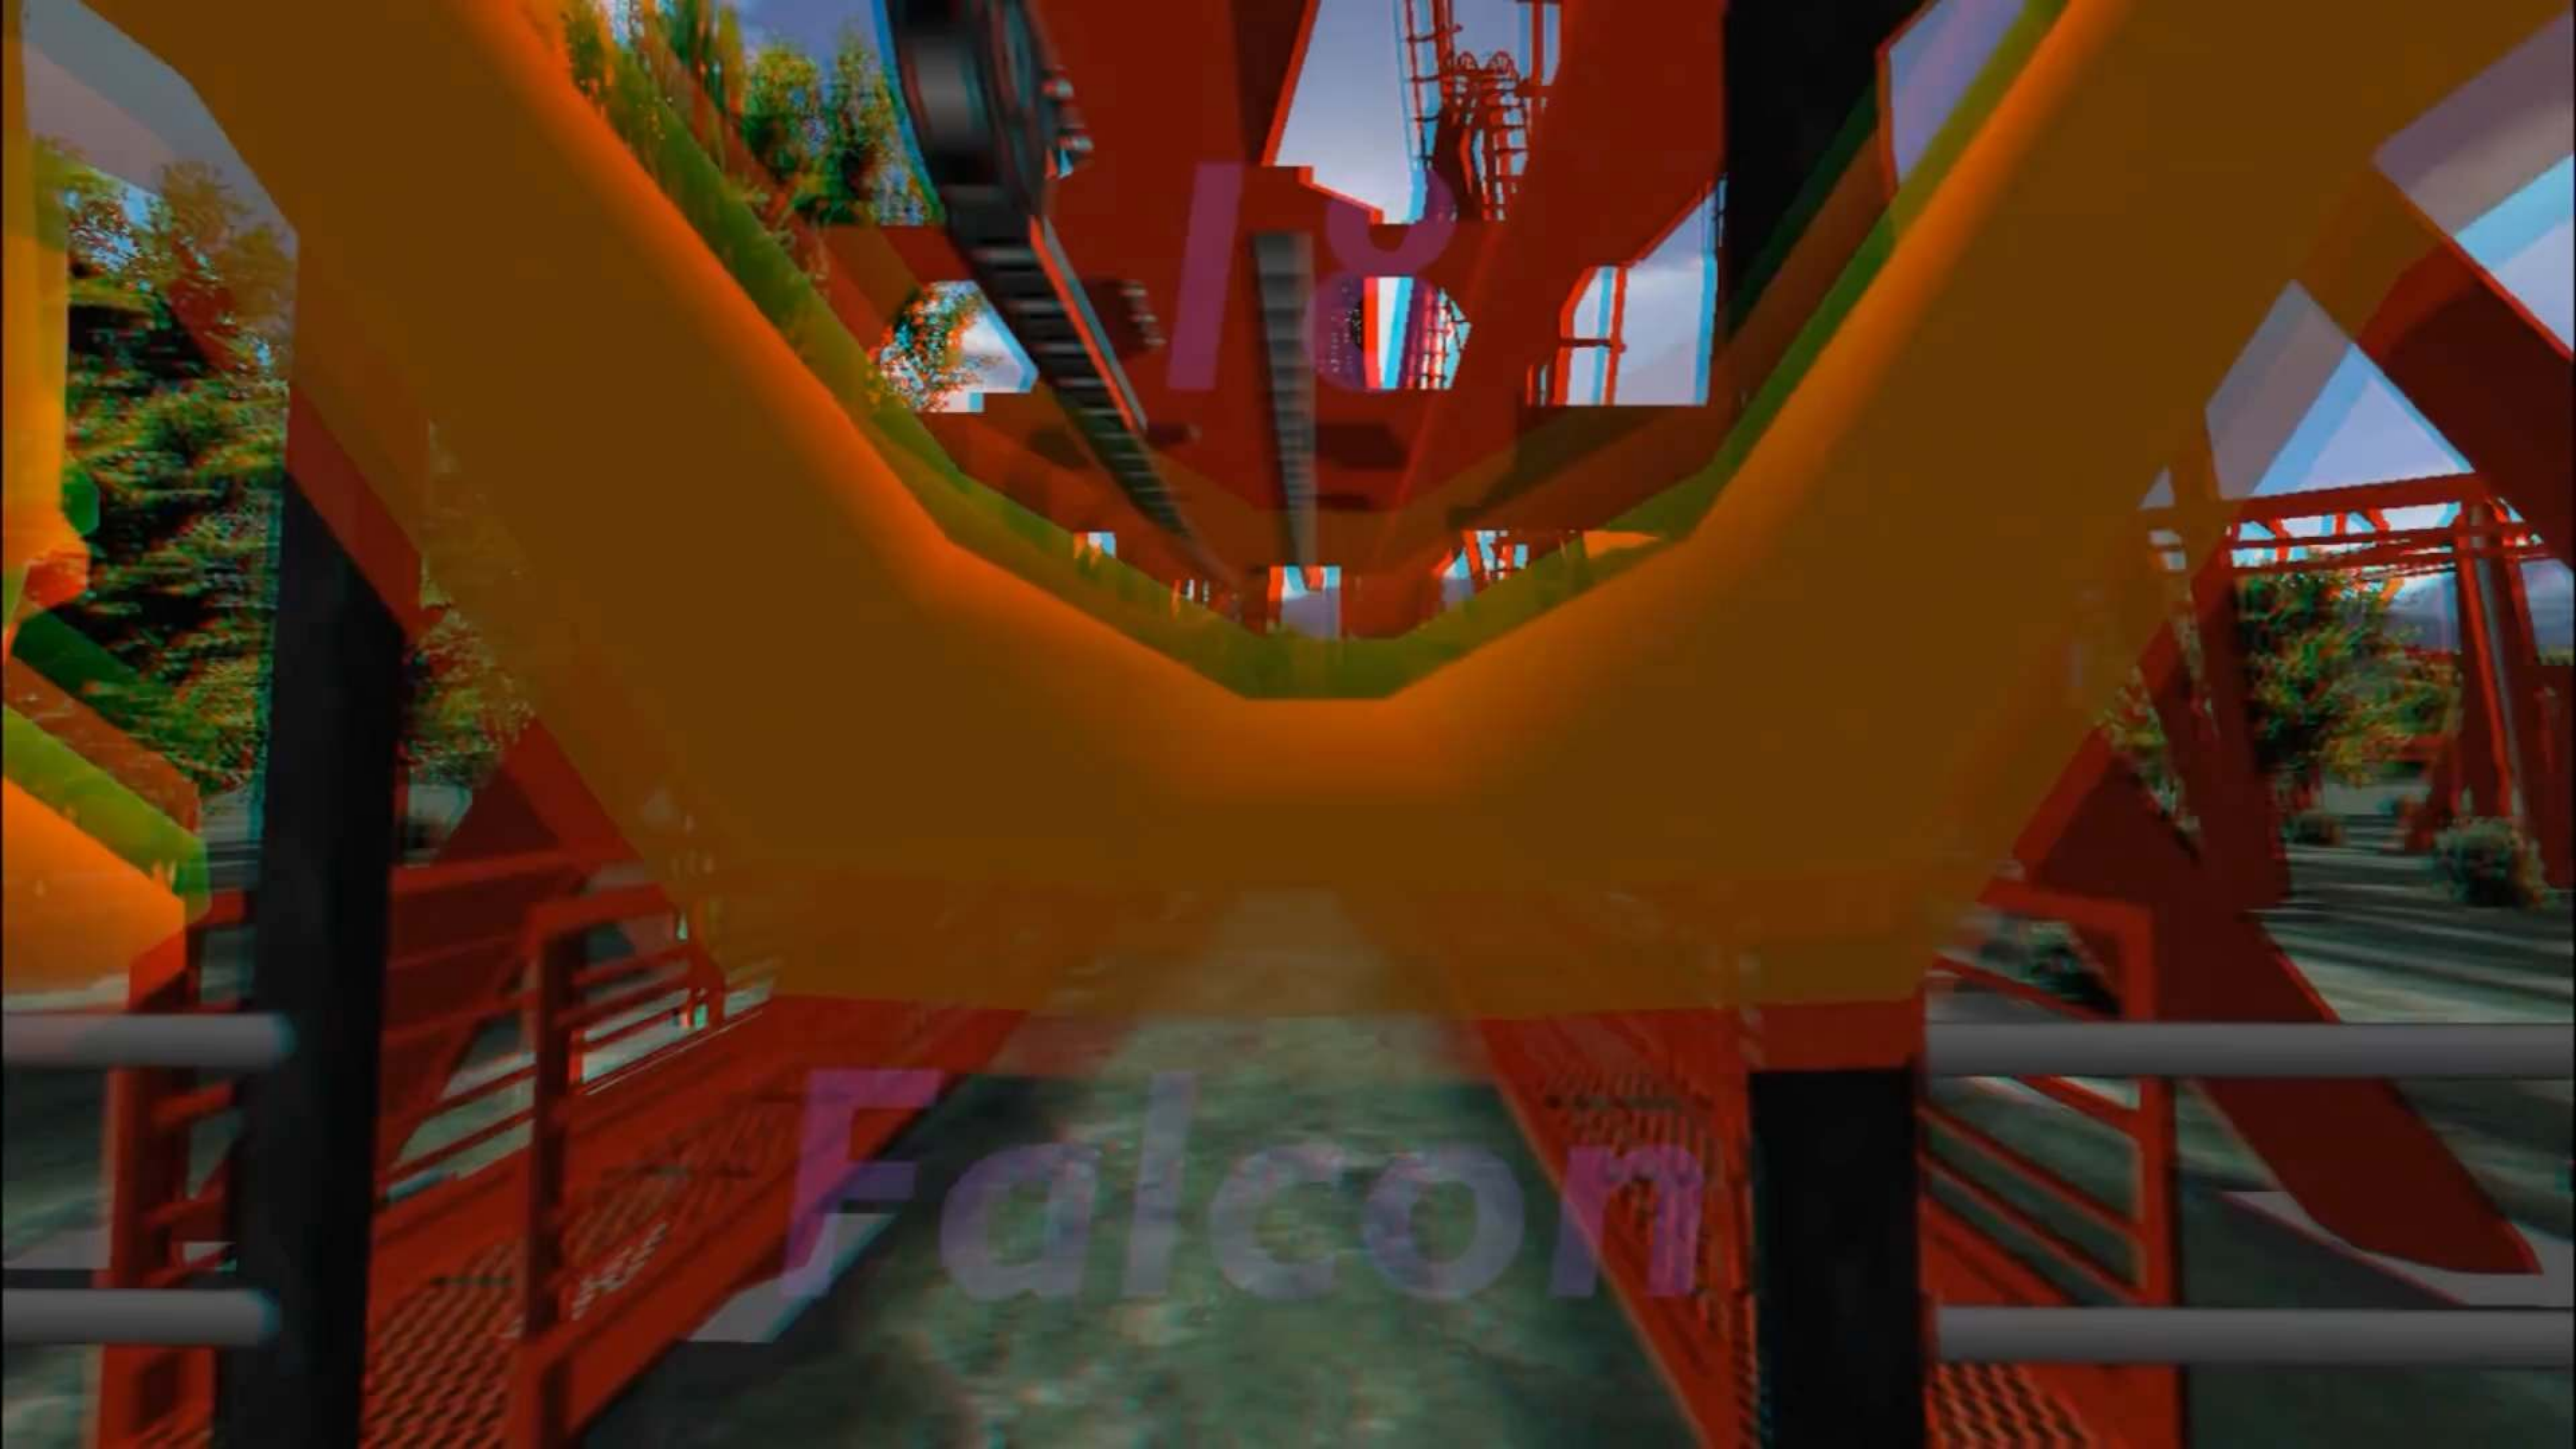
Falcon

Creating Anaglyph Images

- Separation by color filters
  - Convention: red = left eye, cyan = right eye
- For monochrome images:
  - Render left & right images and convert to grayscale $\longrightarrow$ L, R images
  - Merge into anaglyph image $I(r,g,b)$ by assigning

$$I(r) = L \; , \quad I(g,b) = R$$

- For full color anaglyph images:
  - Render left & right images, but do not convert to grayscale $\longrightarrow$ L, R
  - Merge into anaglyph image:

$$I(r) = L(r) \; , \quad I(g,b) = R(g,b)$$

Multiplexing by Wavelength (Infitec)

- Generalization of anaglyph stereo:
  - Partition whole spectrum into 6 (narrow) bands
  - Left & right eye get filters with *interleaving* band passes
  - Other names: Dolby3D, spectral comb filter
- Tricky part: color fidelity

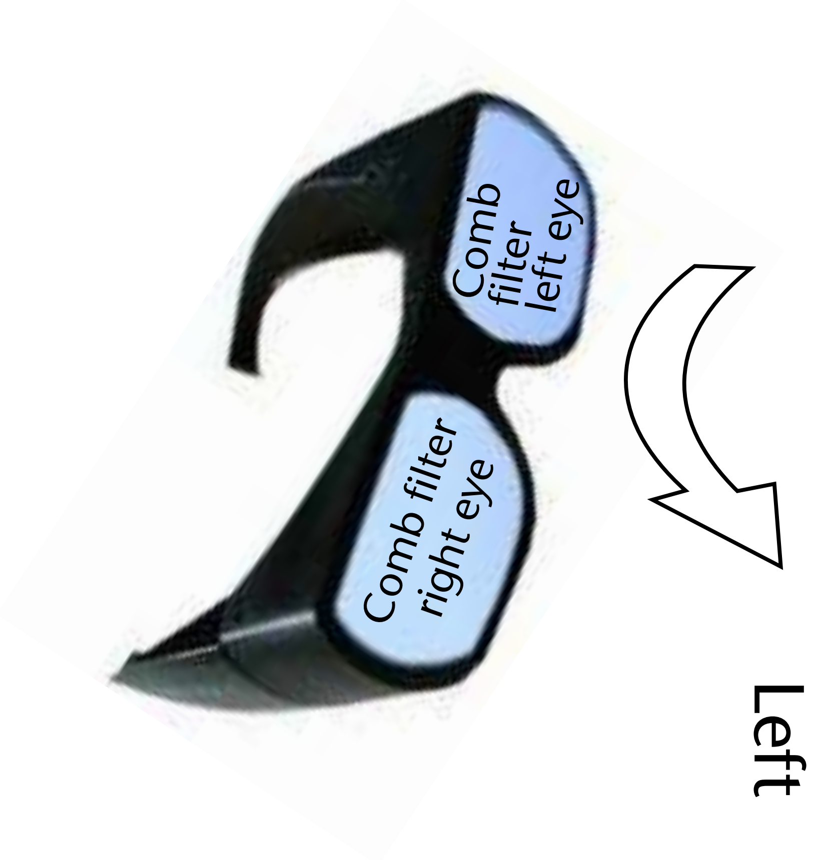

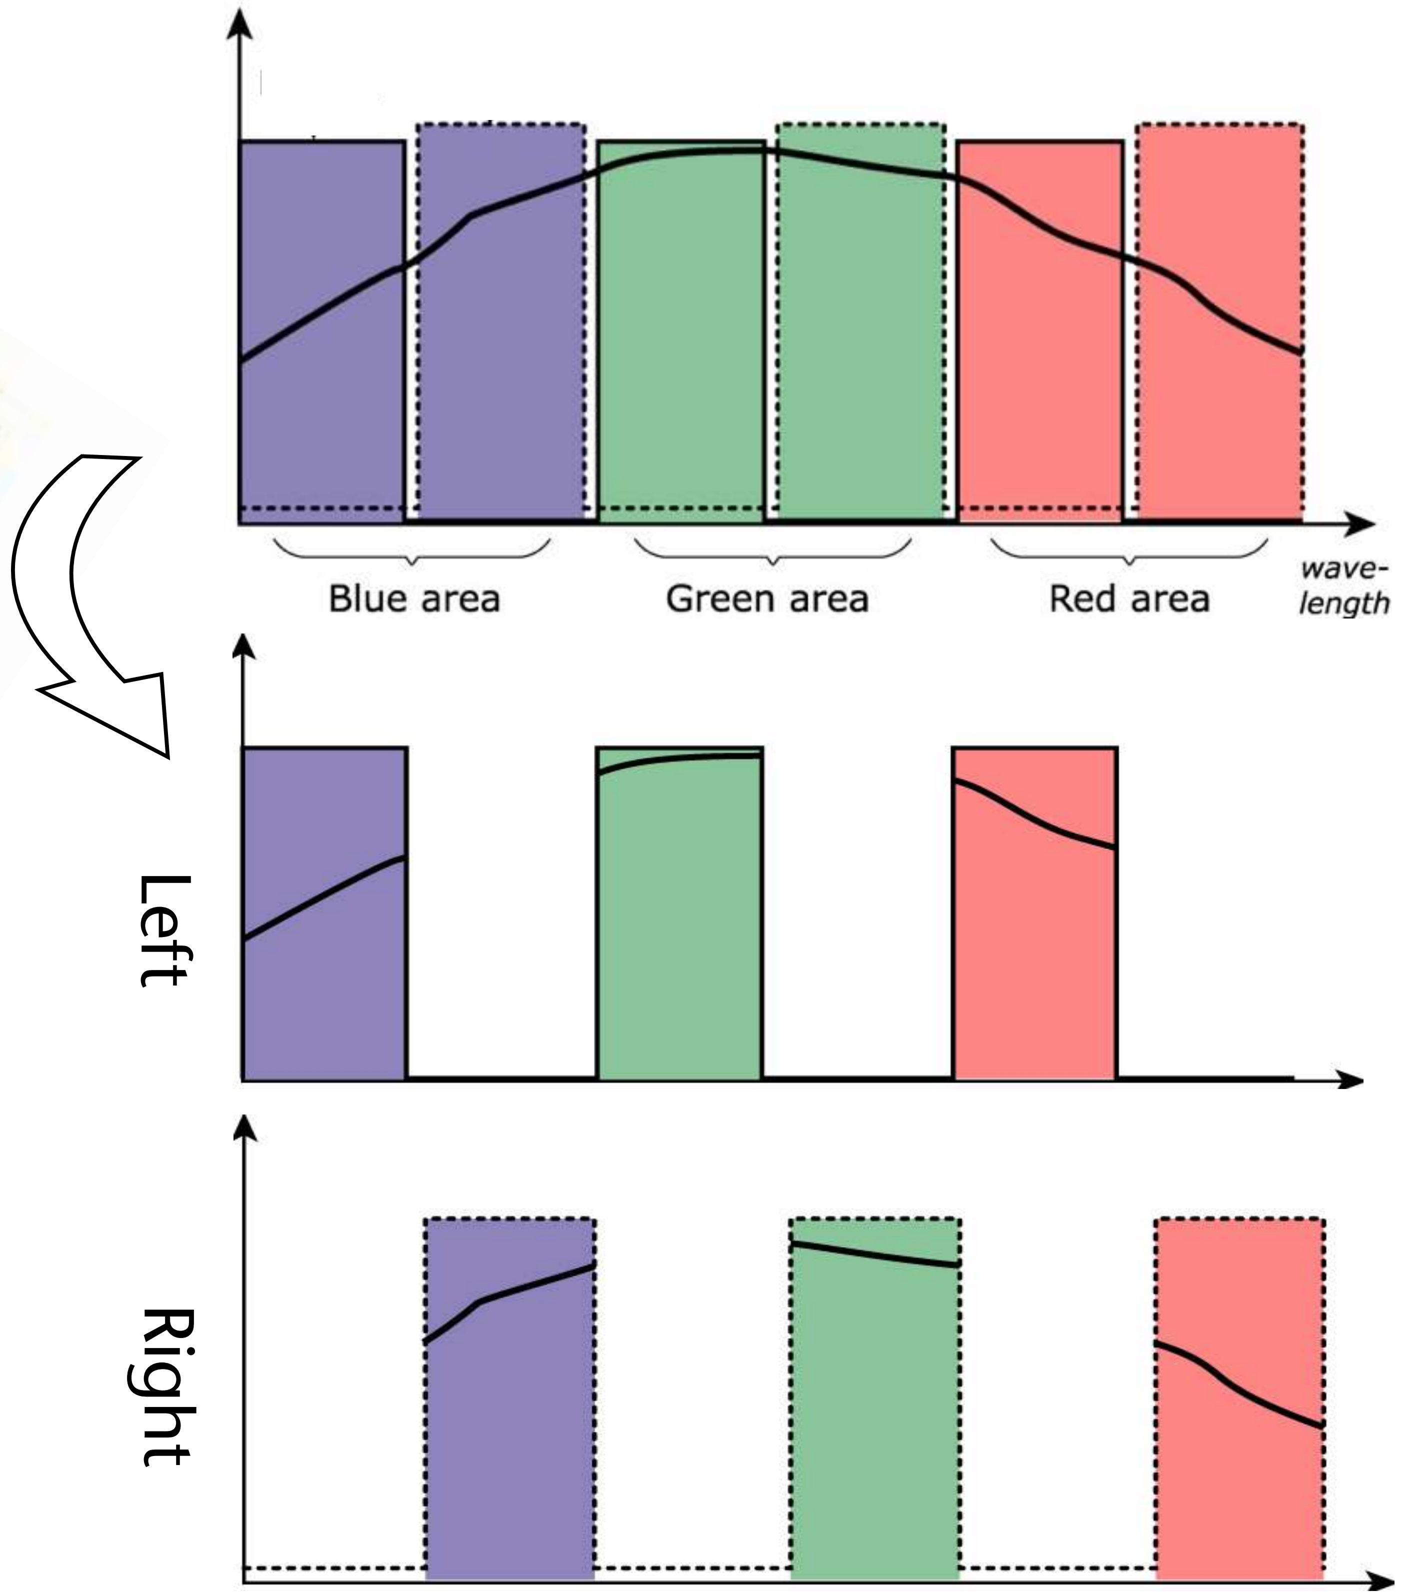

Improvement: Utilize Color Metamerism

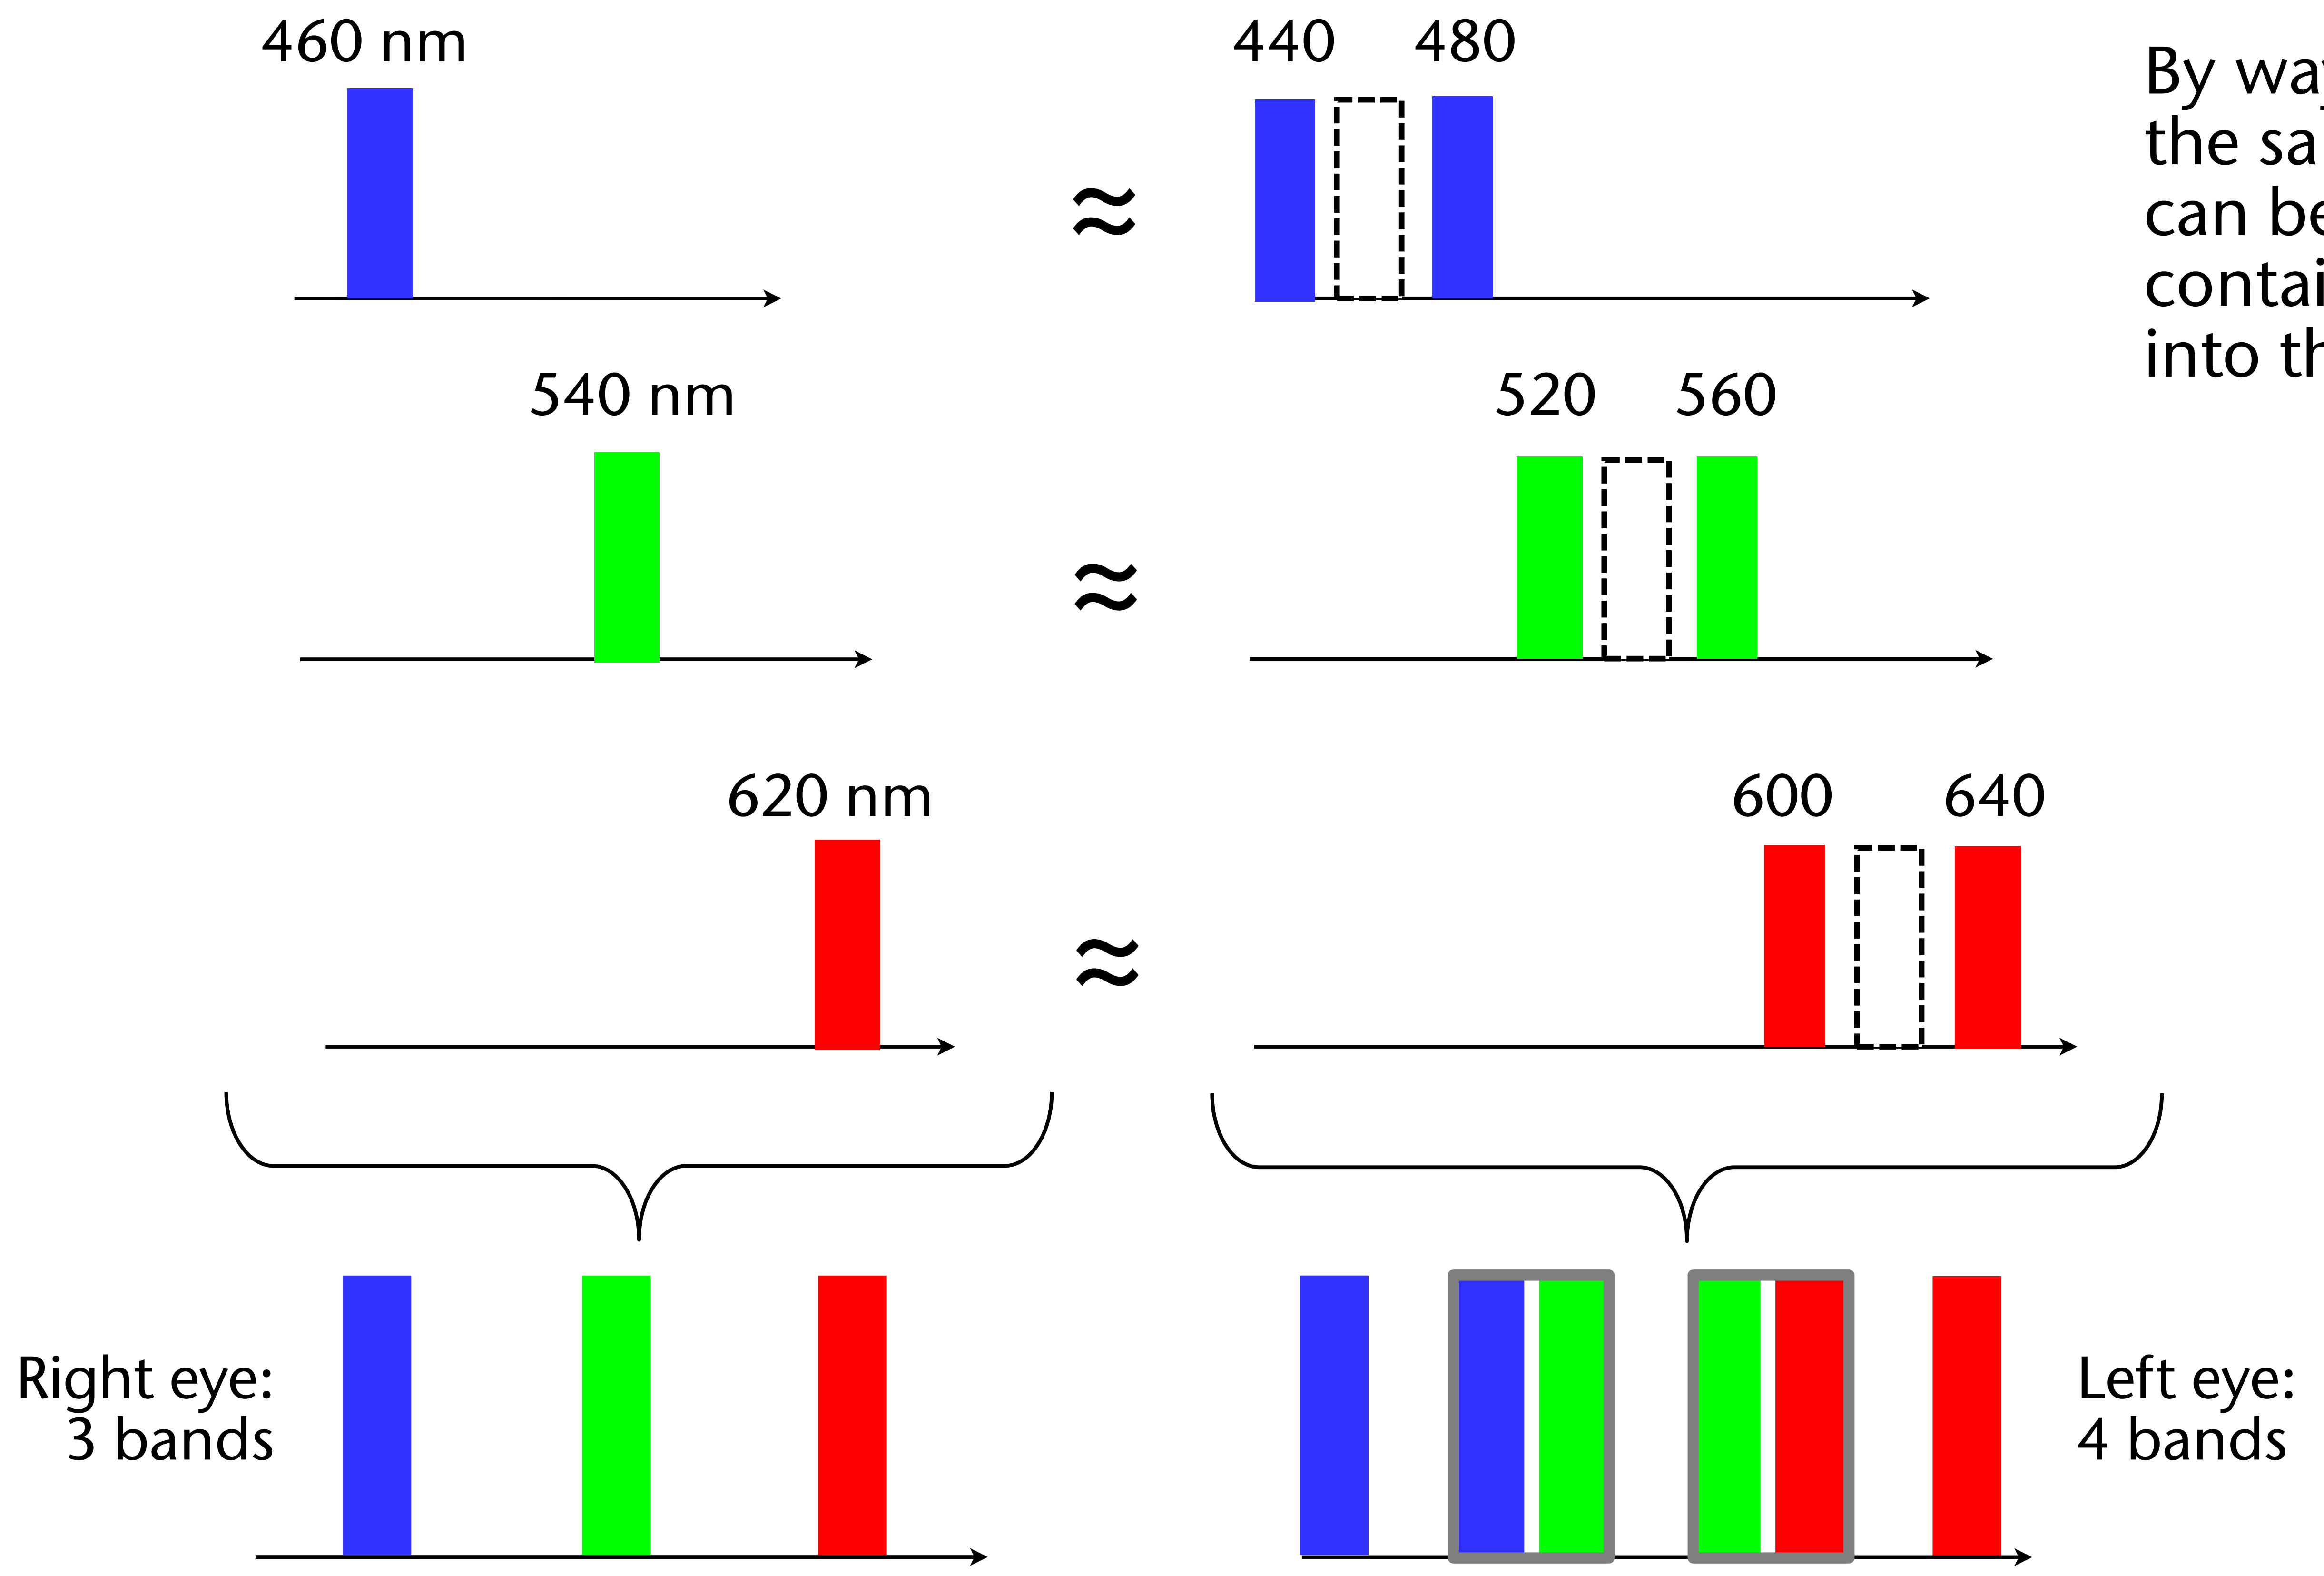

By way of metamerism,
the same "color" of 460 nm light
can be created by shining light
containing 440 and 480 nm
into the eye

Courtesy Infitec (infitec.net)

Spatial Multiplexing

- Projection surface is partitioned among users
- Consequence: interdependence between
  - Size of the *view frustum*
  - Working volume of users
  - D & radius of hole

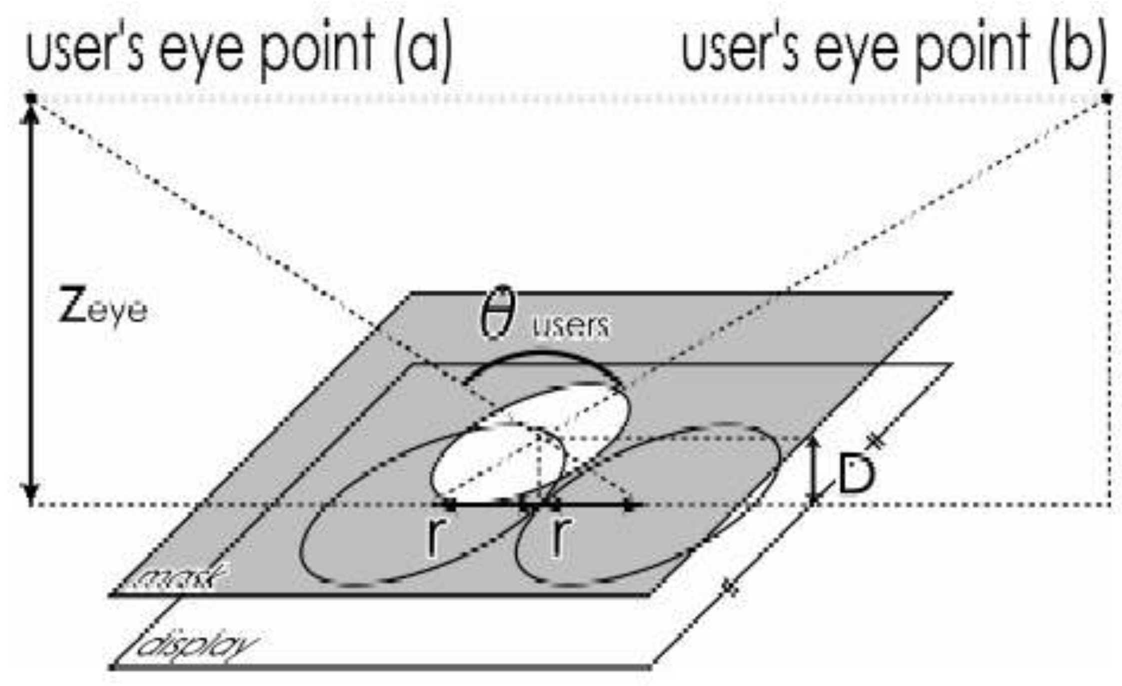

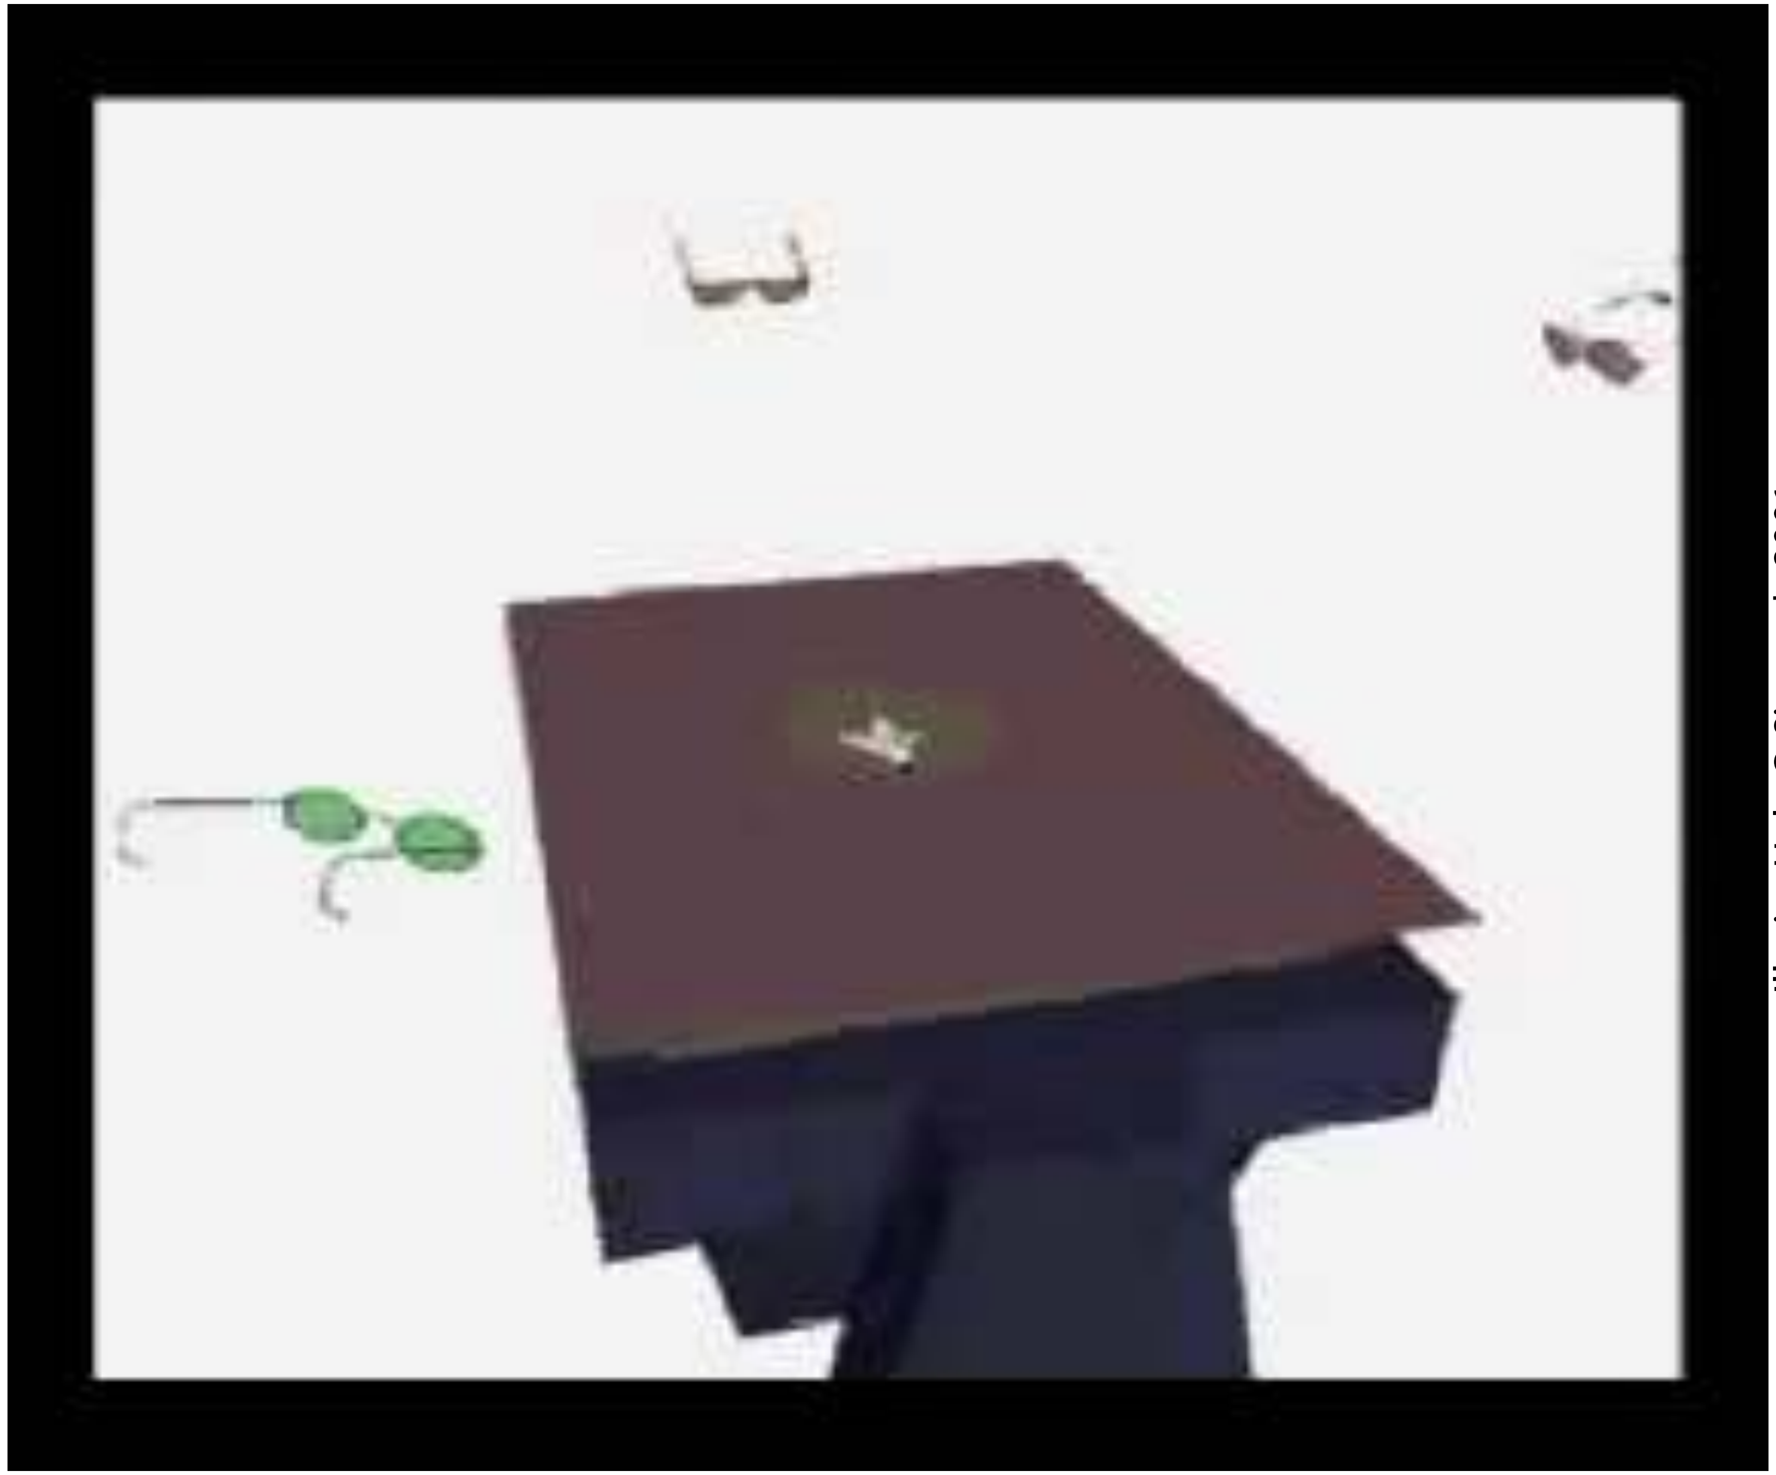

Curiosity: Autostereogram (Single Image Stereogram)

- "Magic Eye" images are patterns constructed such that seemingly corresponding points (within *same* image) convey the desired depth

This is what your eyes focus, but with "parallel" axes (accommodation = near, convergence = far)

Underlying "depth image"

Curiosity: Stereoscopic Effect Based on the Pulfrich-Effect

- The Pulfrich effect:
  - Dark stimulus in the eye arrives later in the brain than a bright stimulus
  - Discovered by Carl Pulfrich, German physicist, 1922
- Viewing instructions: put sunglasses or similar darkening filter over *one* eye, the other eye remains naked

http://www.youtube.com/watch?v=1mnWI_u_zBg

Demo Video: Looking Through a Window ..

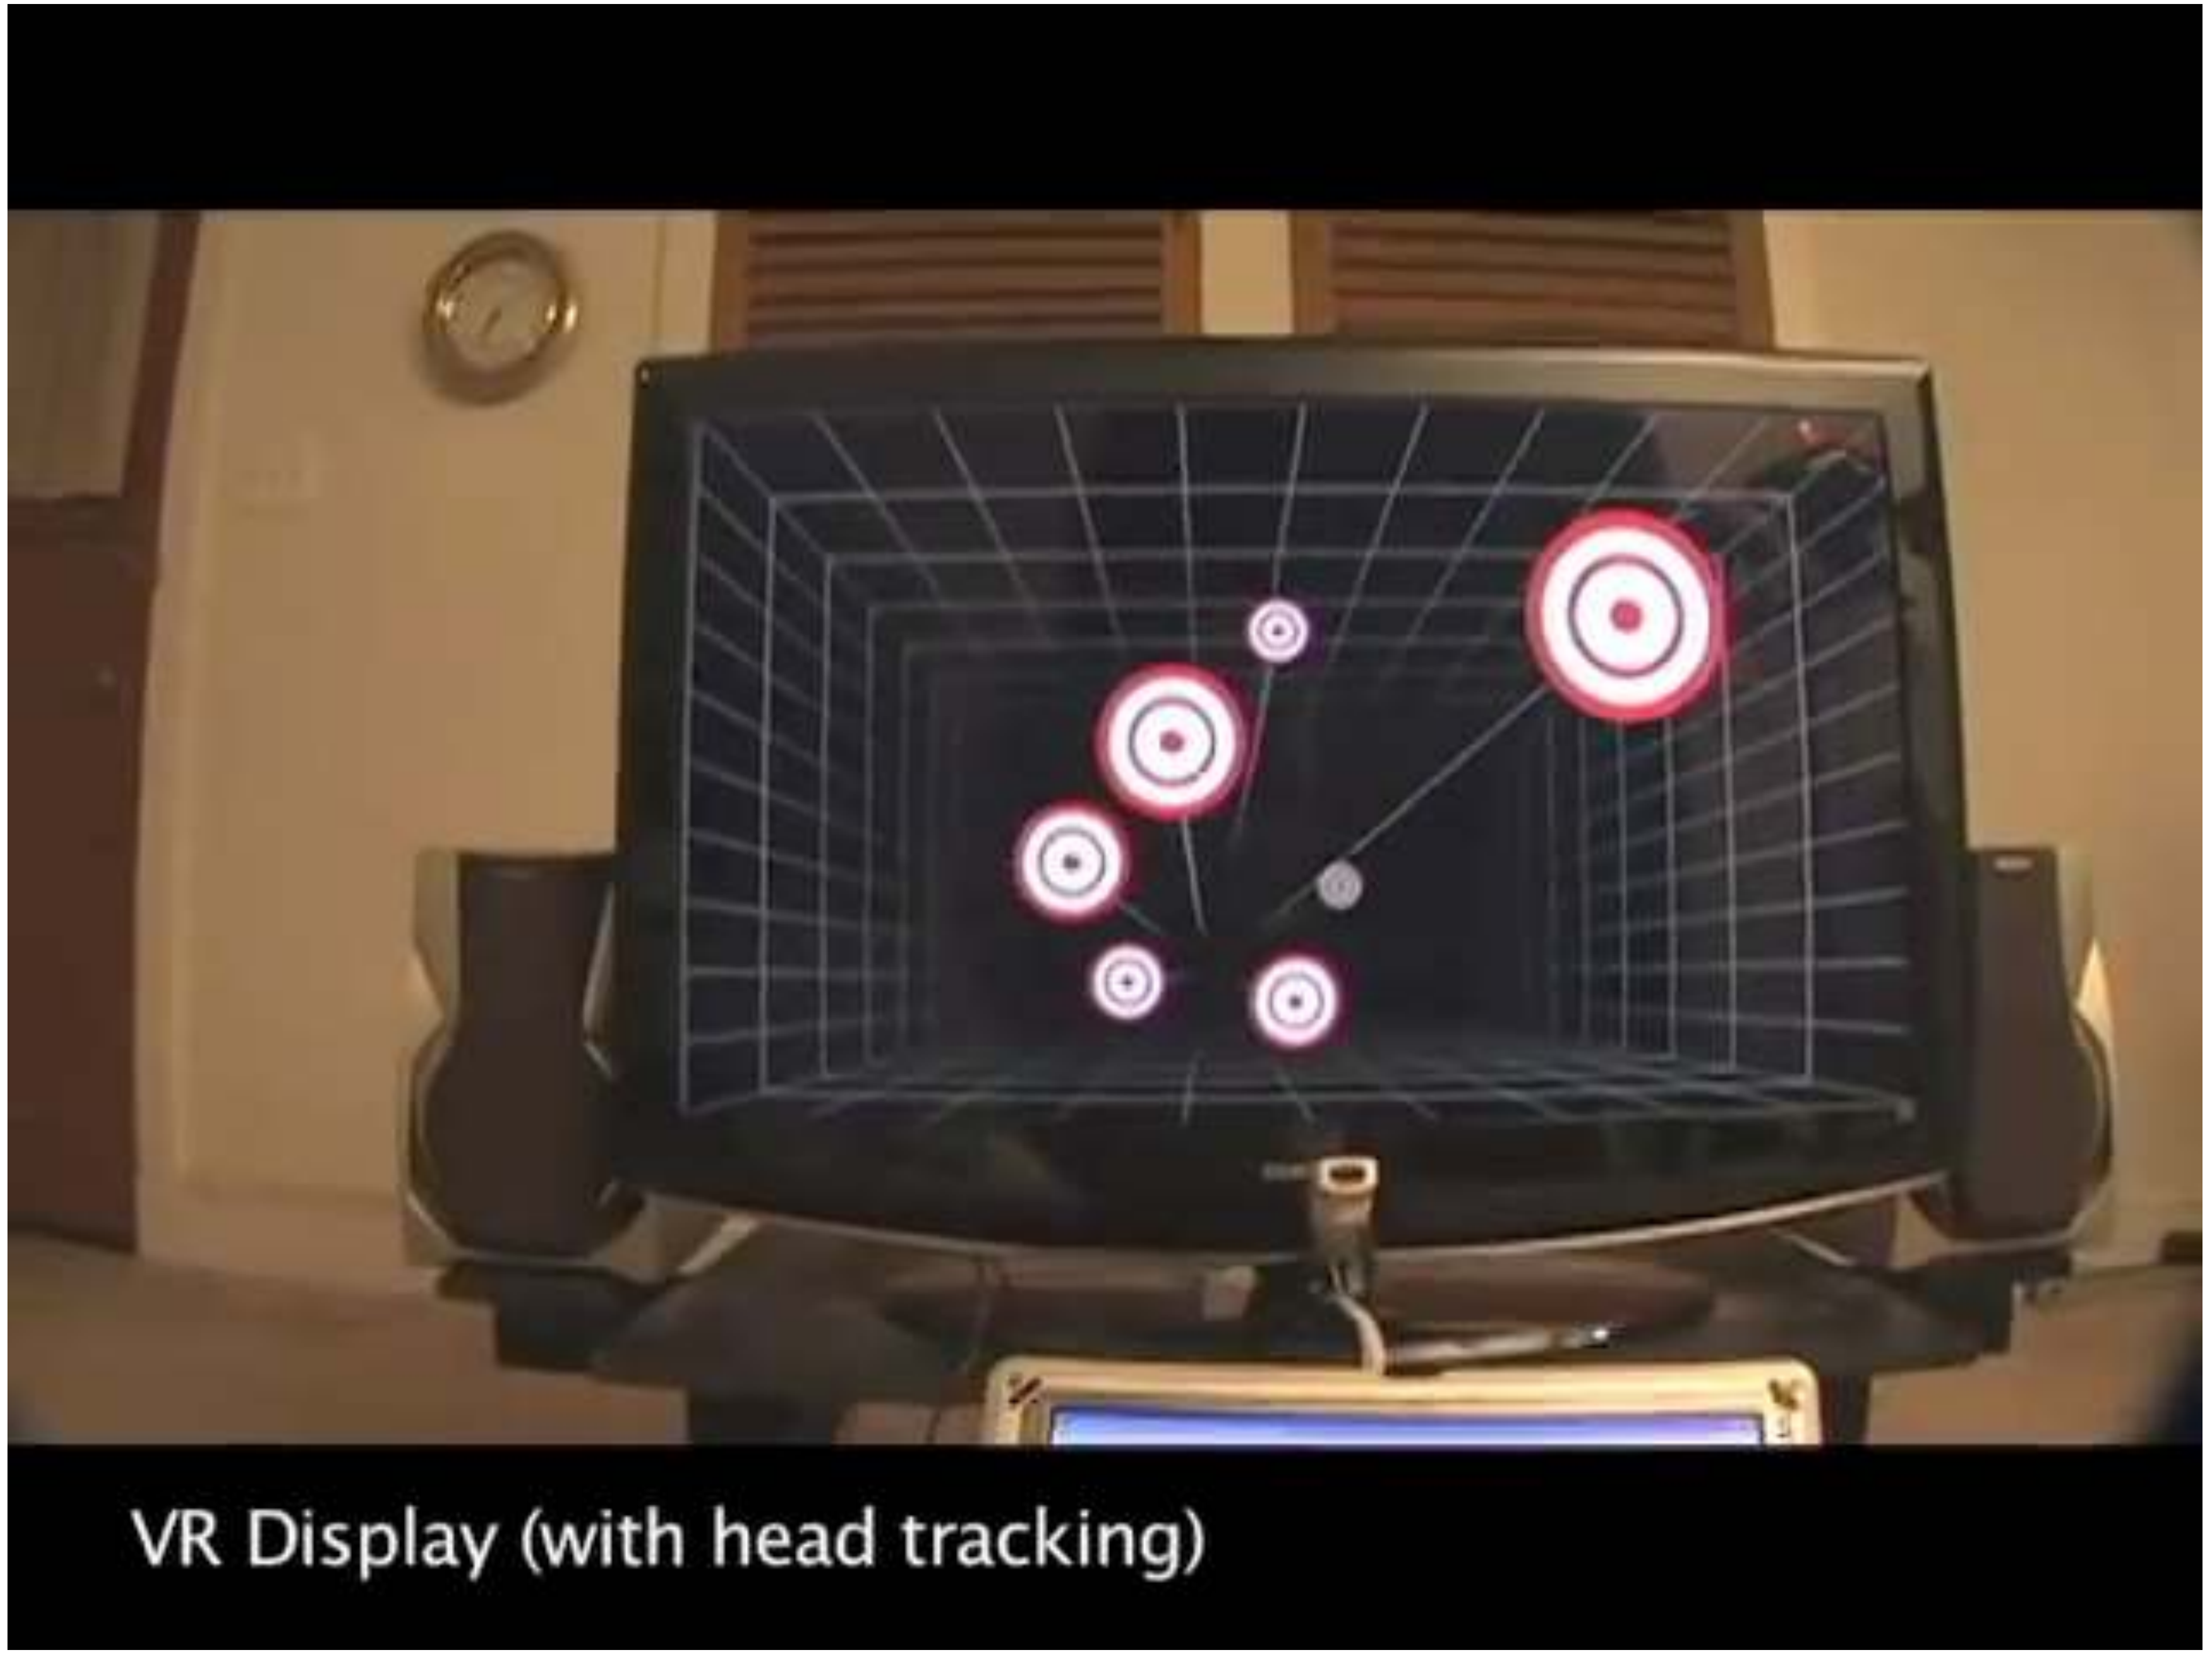

Recap: Perspective Projection in Graphics API's

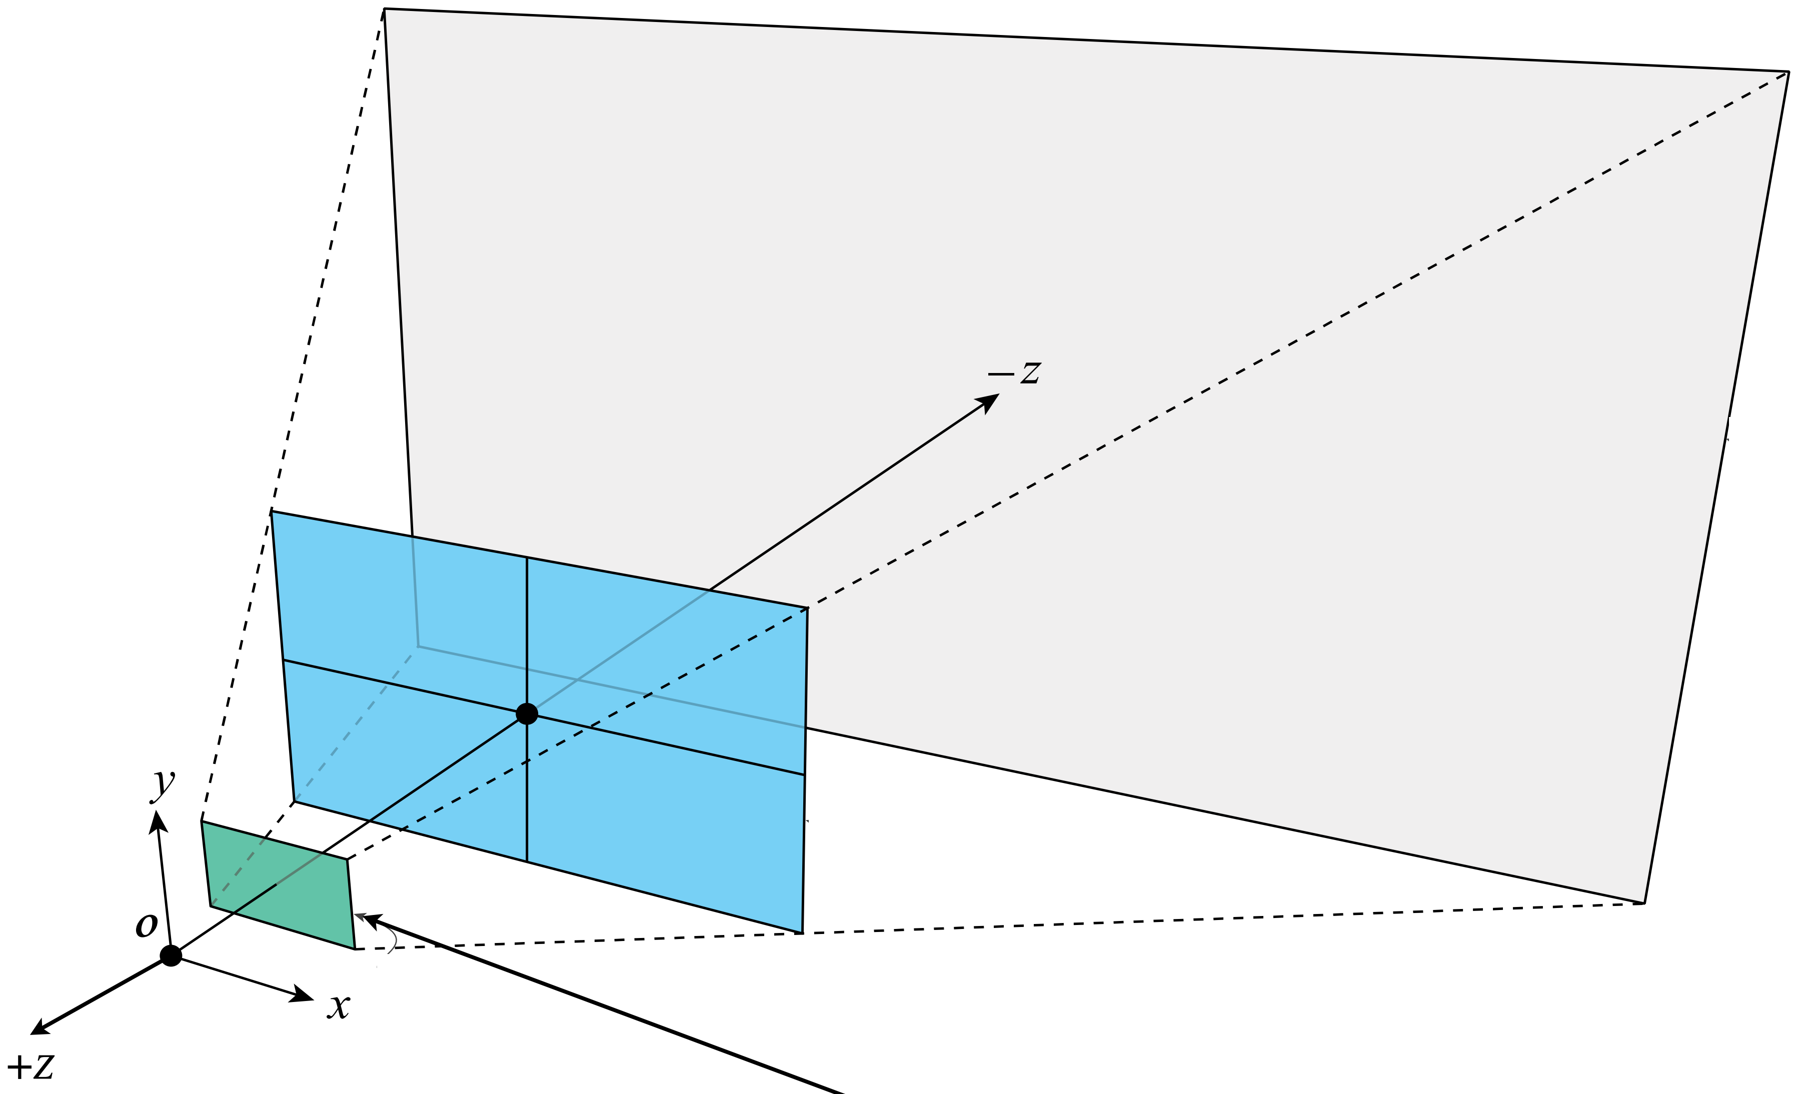

Near plane (near clipping plane) at $z = -n$
Projection plane, also

Stereoscopic Projection (aka. Stereo Rendering)

- Stereo parallax on the screen ⟶ disparity in the eyes

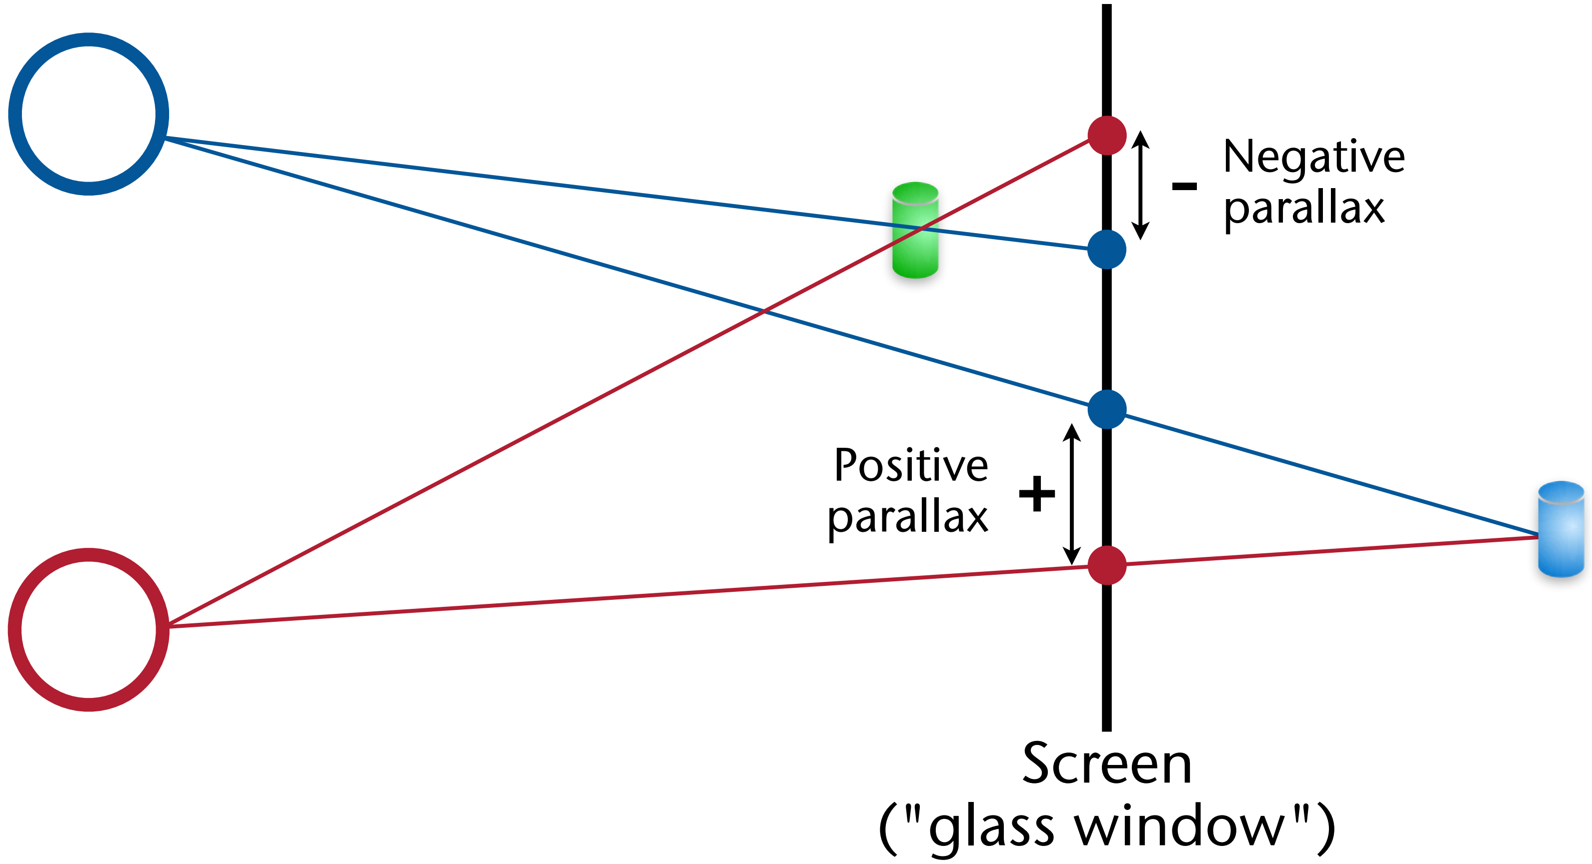

- Wrong way: converging view vectors
  - Problem: vertical parallax!

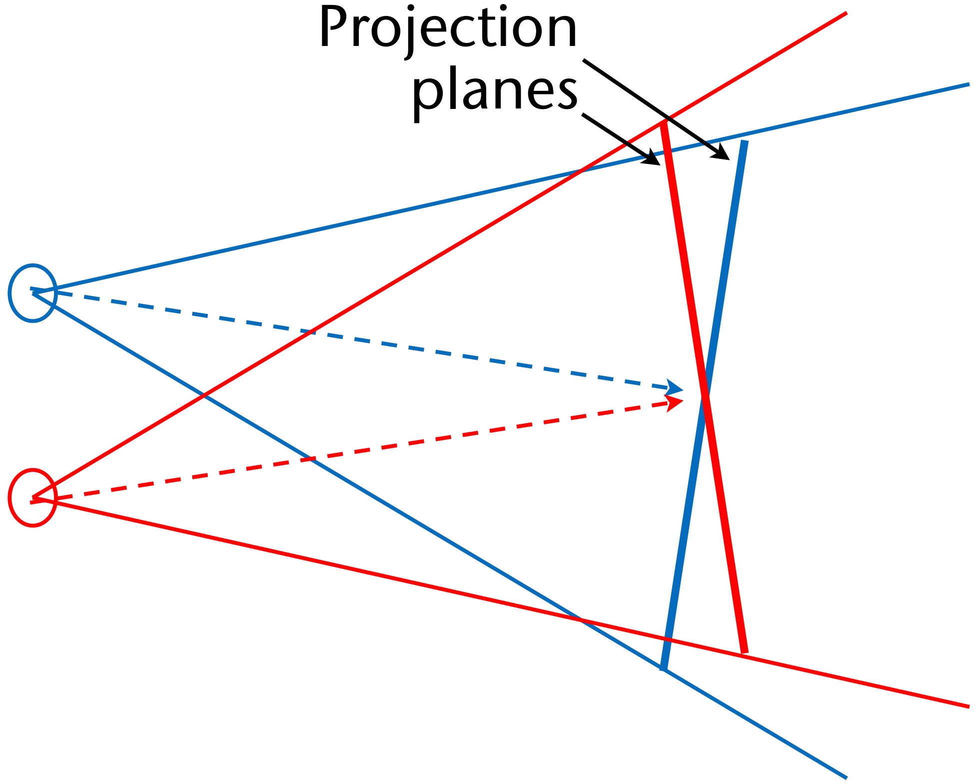

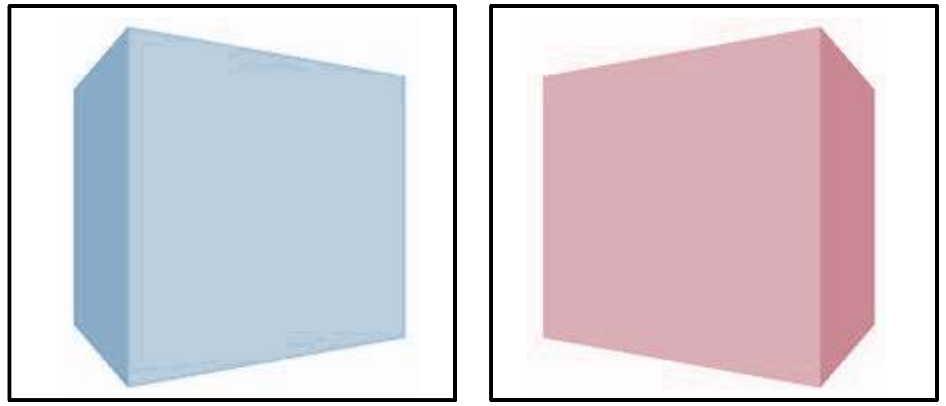

Wrong Stereo → Vertical Parallax Causes Severe Eye Strain

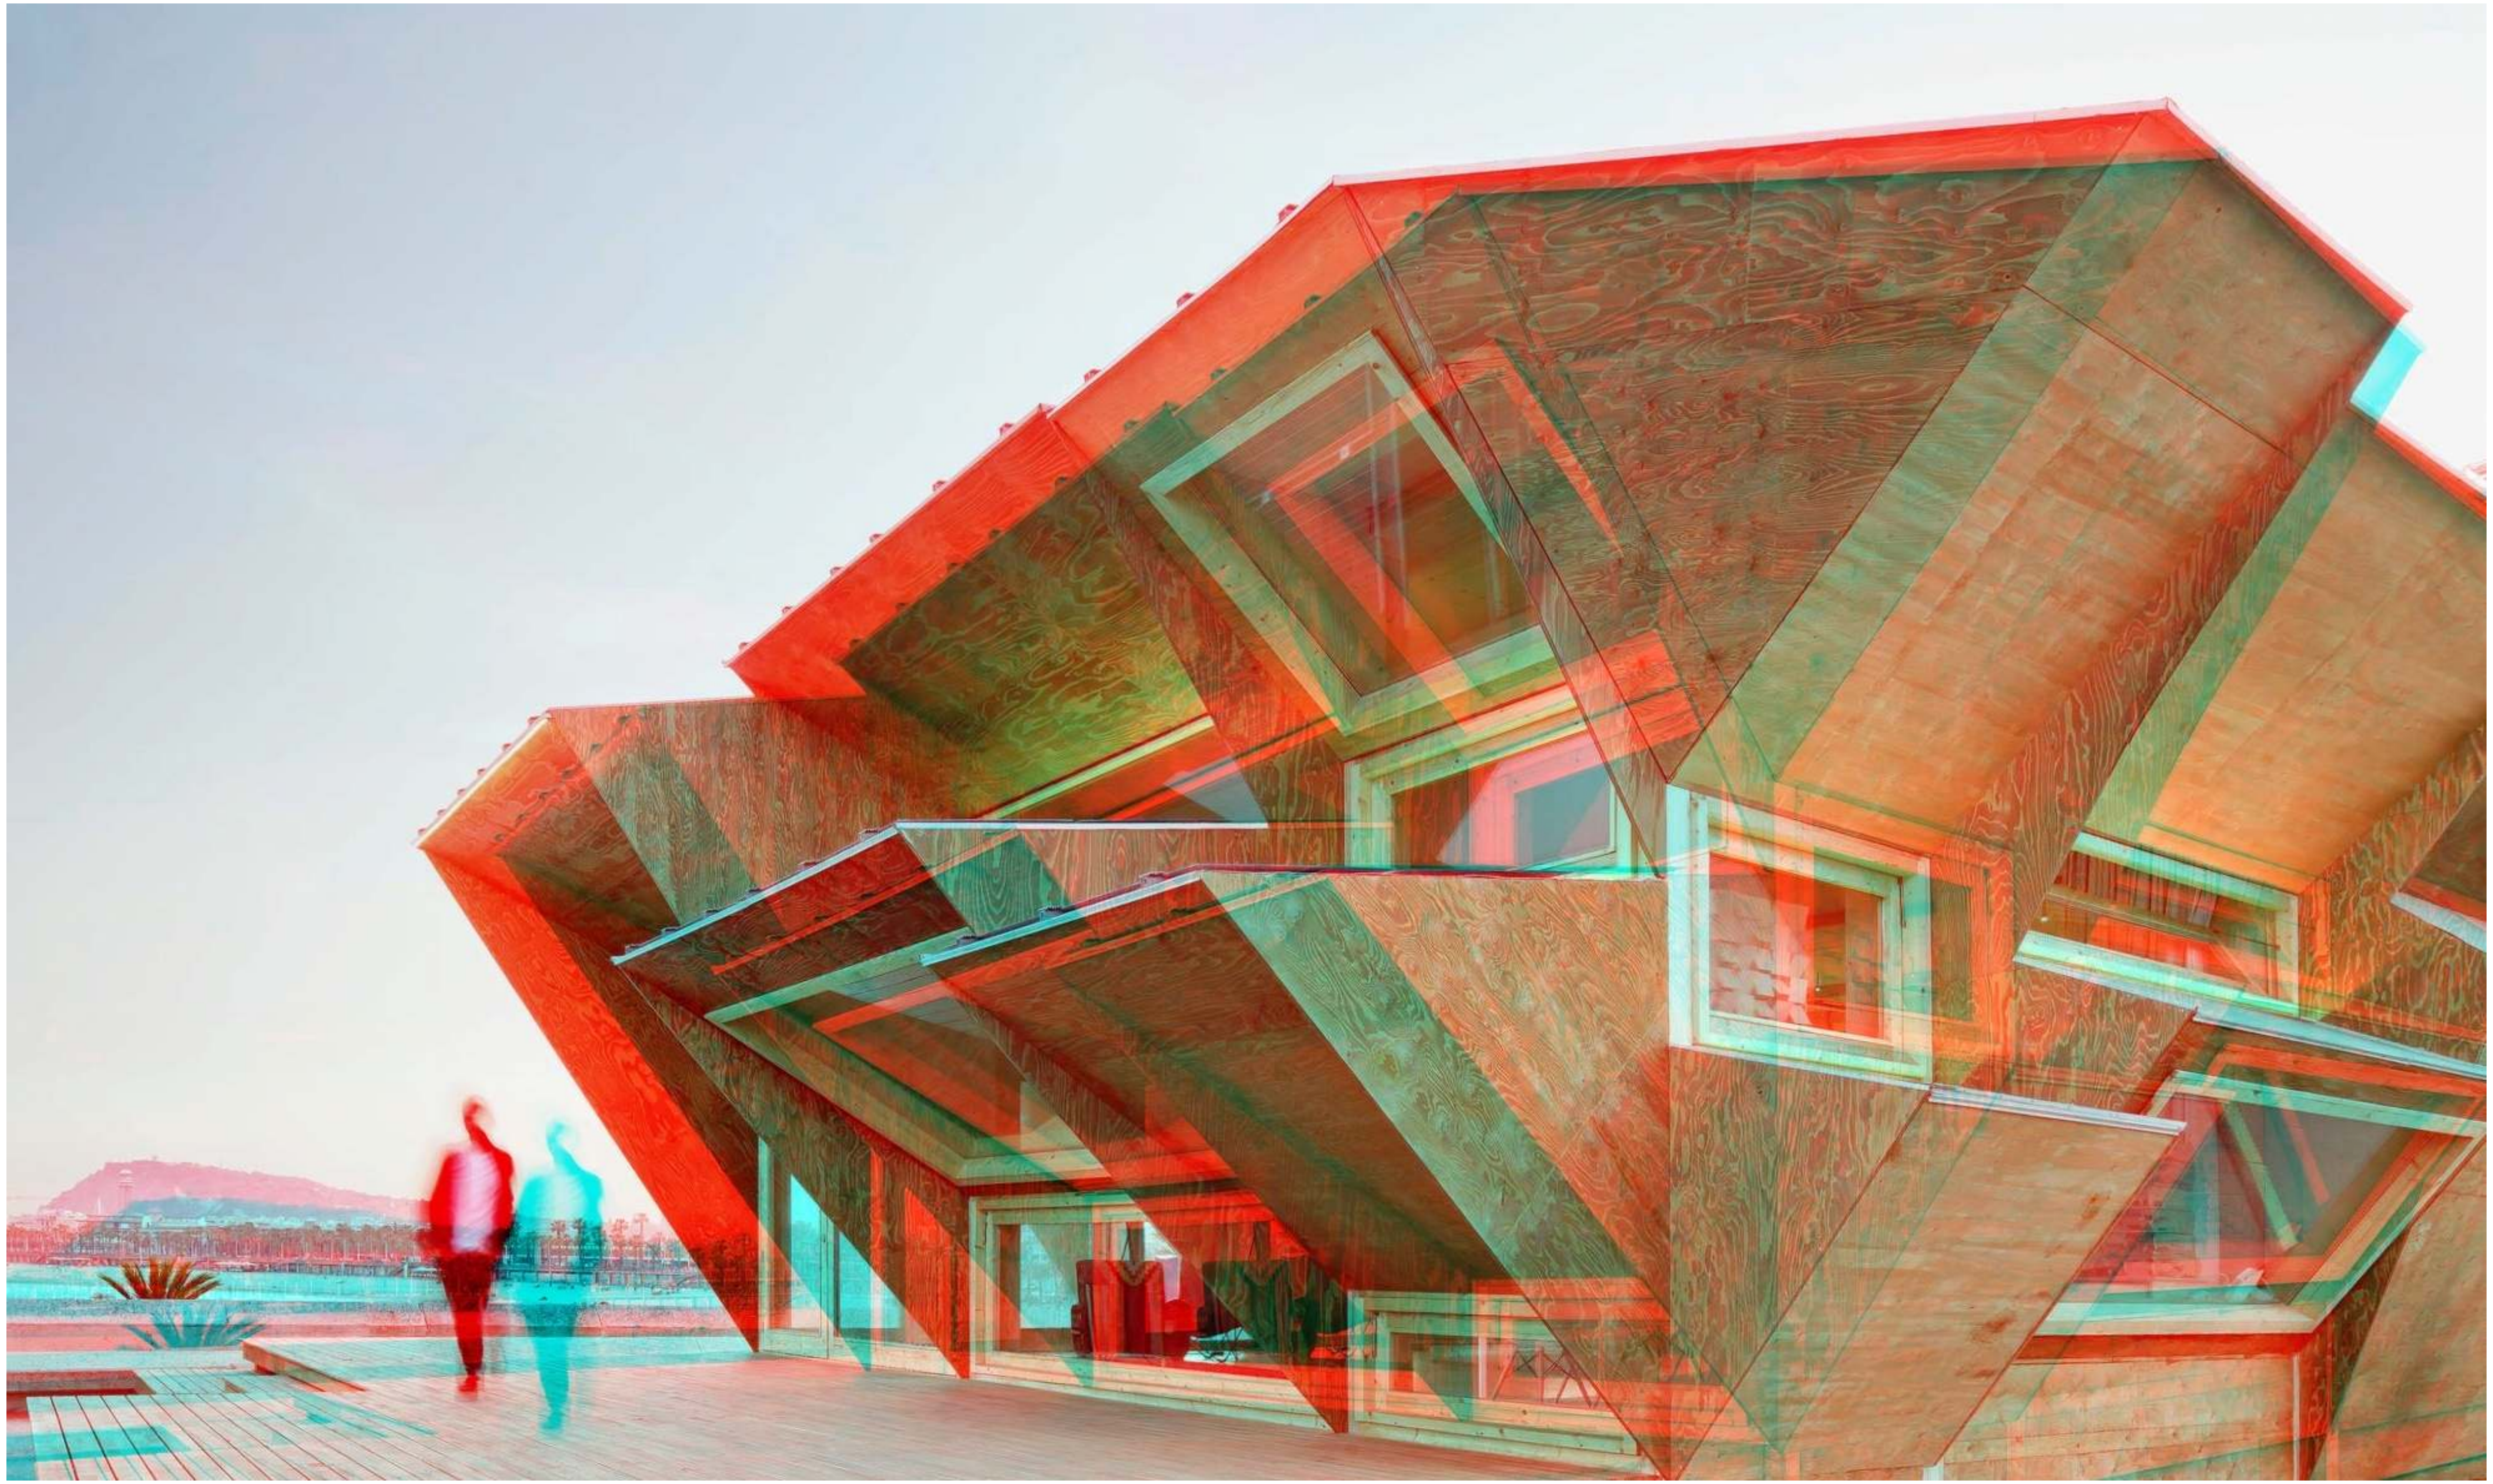

Correct Stereoscopic Projection for IPT's

- Parallel viewing vectors
- Important stereo parameters:
  - Cyclop's eye (center between left and right eye)
  - Eye separation, aka. interpupilary distance (IPD)
  - Zero parallax plane (ZPP), aka. "fusion distance" or "horopter"
- Off-center perspective projection (a.k.a. "off-axis projection")

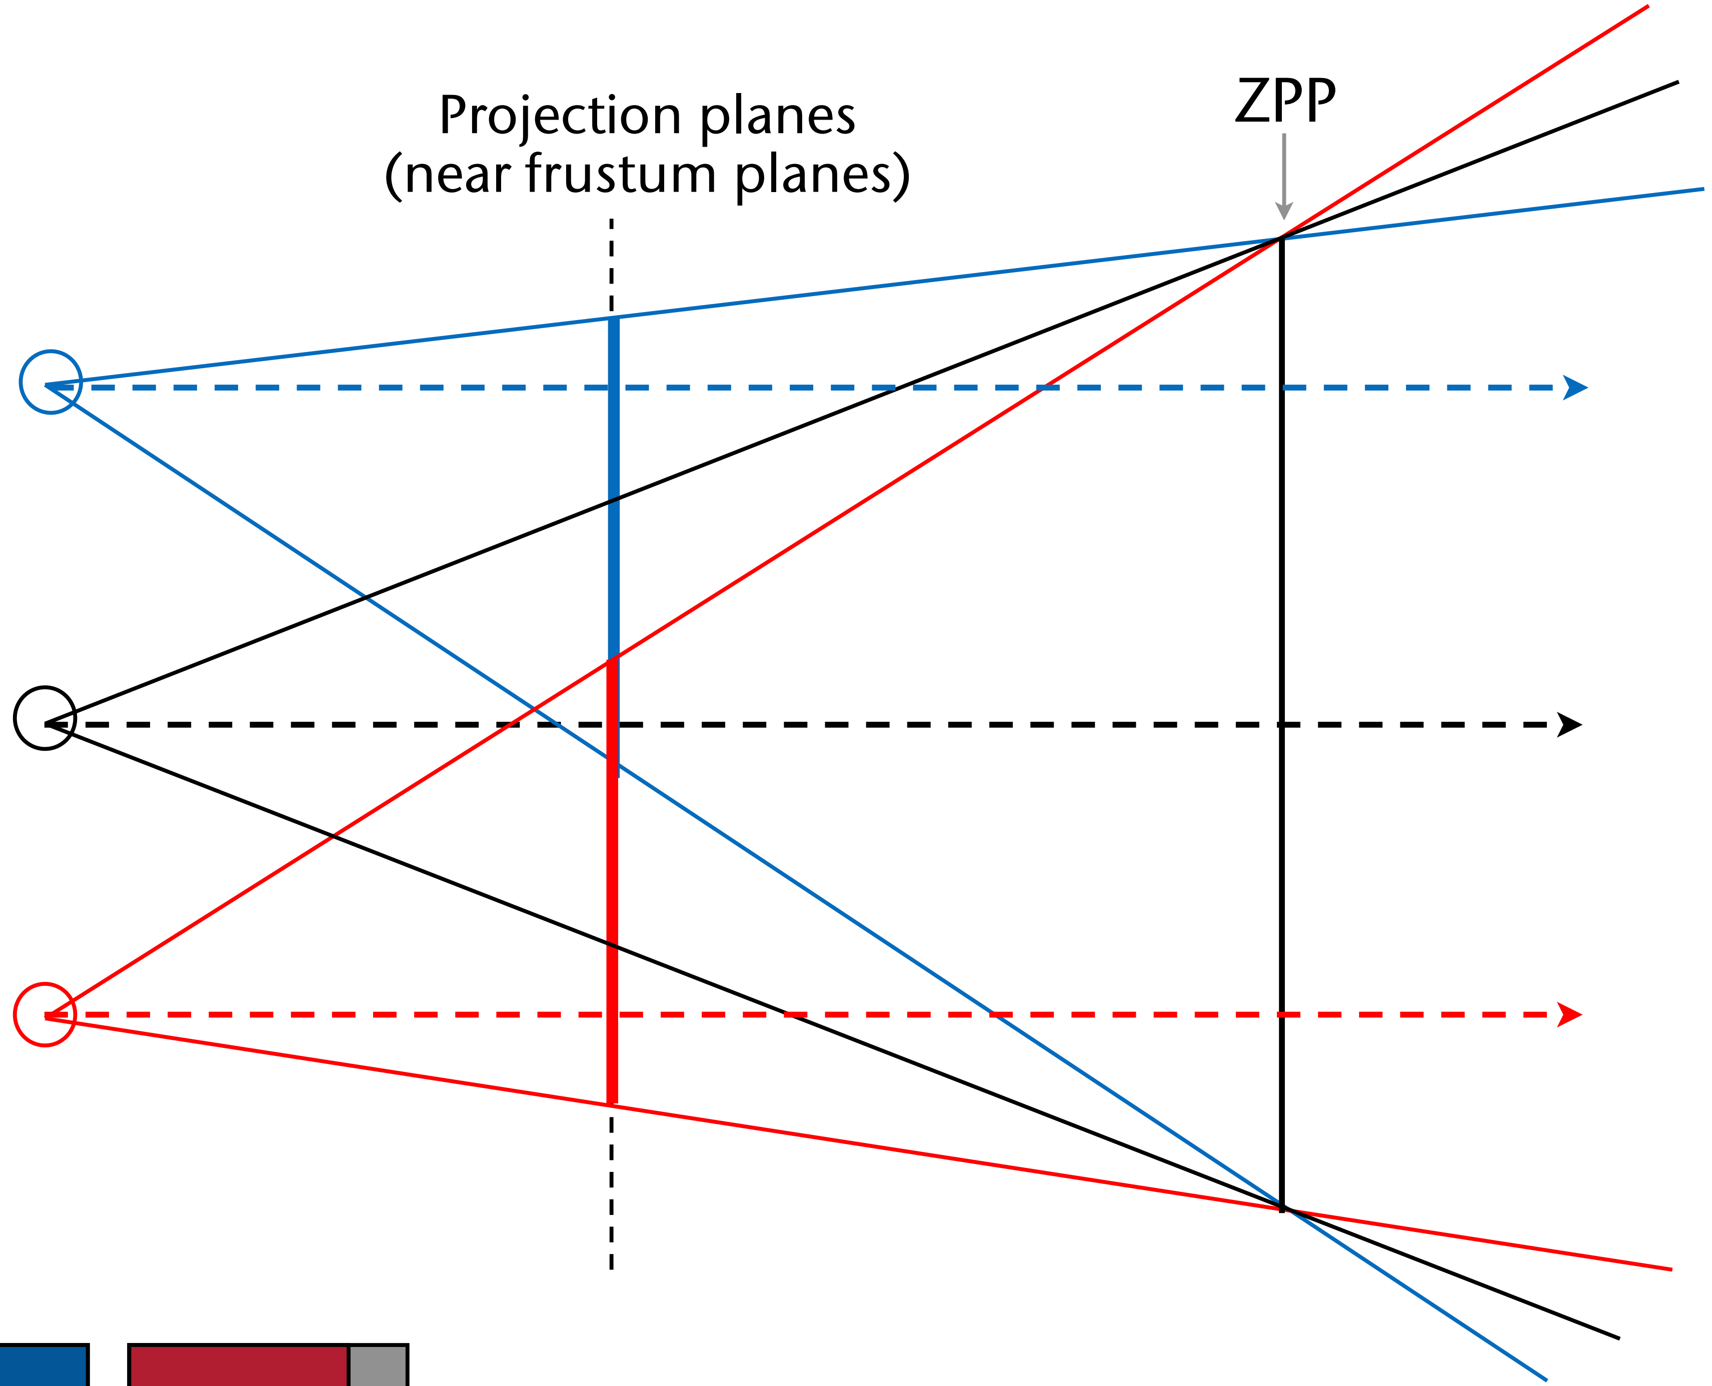

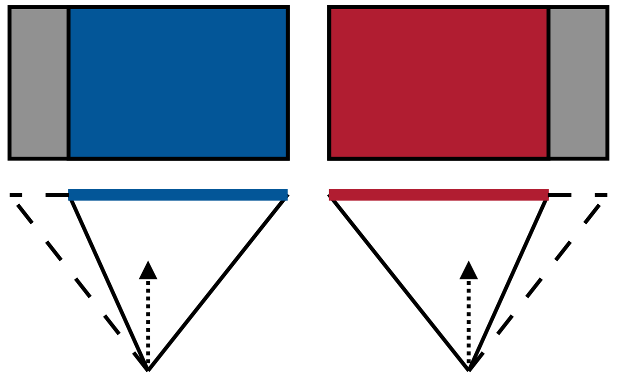

What is the Difference Between Stereo Parallax and Disparity?

https://www.menti.com/hht2ei5jxt

Where is the ZPP in the Anaglyph Images?

Thought Experiment

- Imagine a single line emanating from 1m in front of you, away from you to infinity
- What stereo image do you get?
- What happens, if the IPD increases?
- What happens, if you move the ZPP closer or further away?

Specification of the Projection Screen in Virtual Space

- Screen = window into virtual space → screen must be fixed in virtual space → virtual space is decoupled from user's movement in physical space!
- Specify the screen as a polygon in virtual space, e.g., by its corners
- Specification of the view frustum: distance of left/right/top/bottom edges of the screen from the "midpoint" (= point on screen closest to viewpoint) - **measured on the near plane**!
  - Compare `glFrustum()` in OpenGL

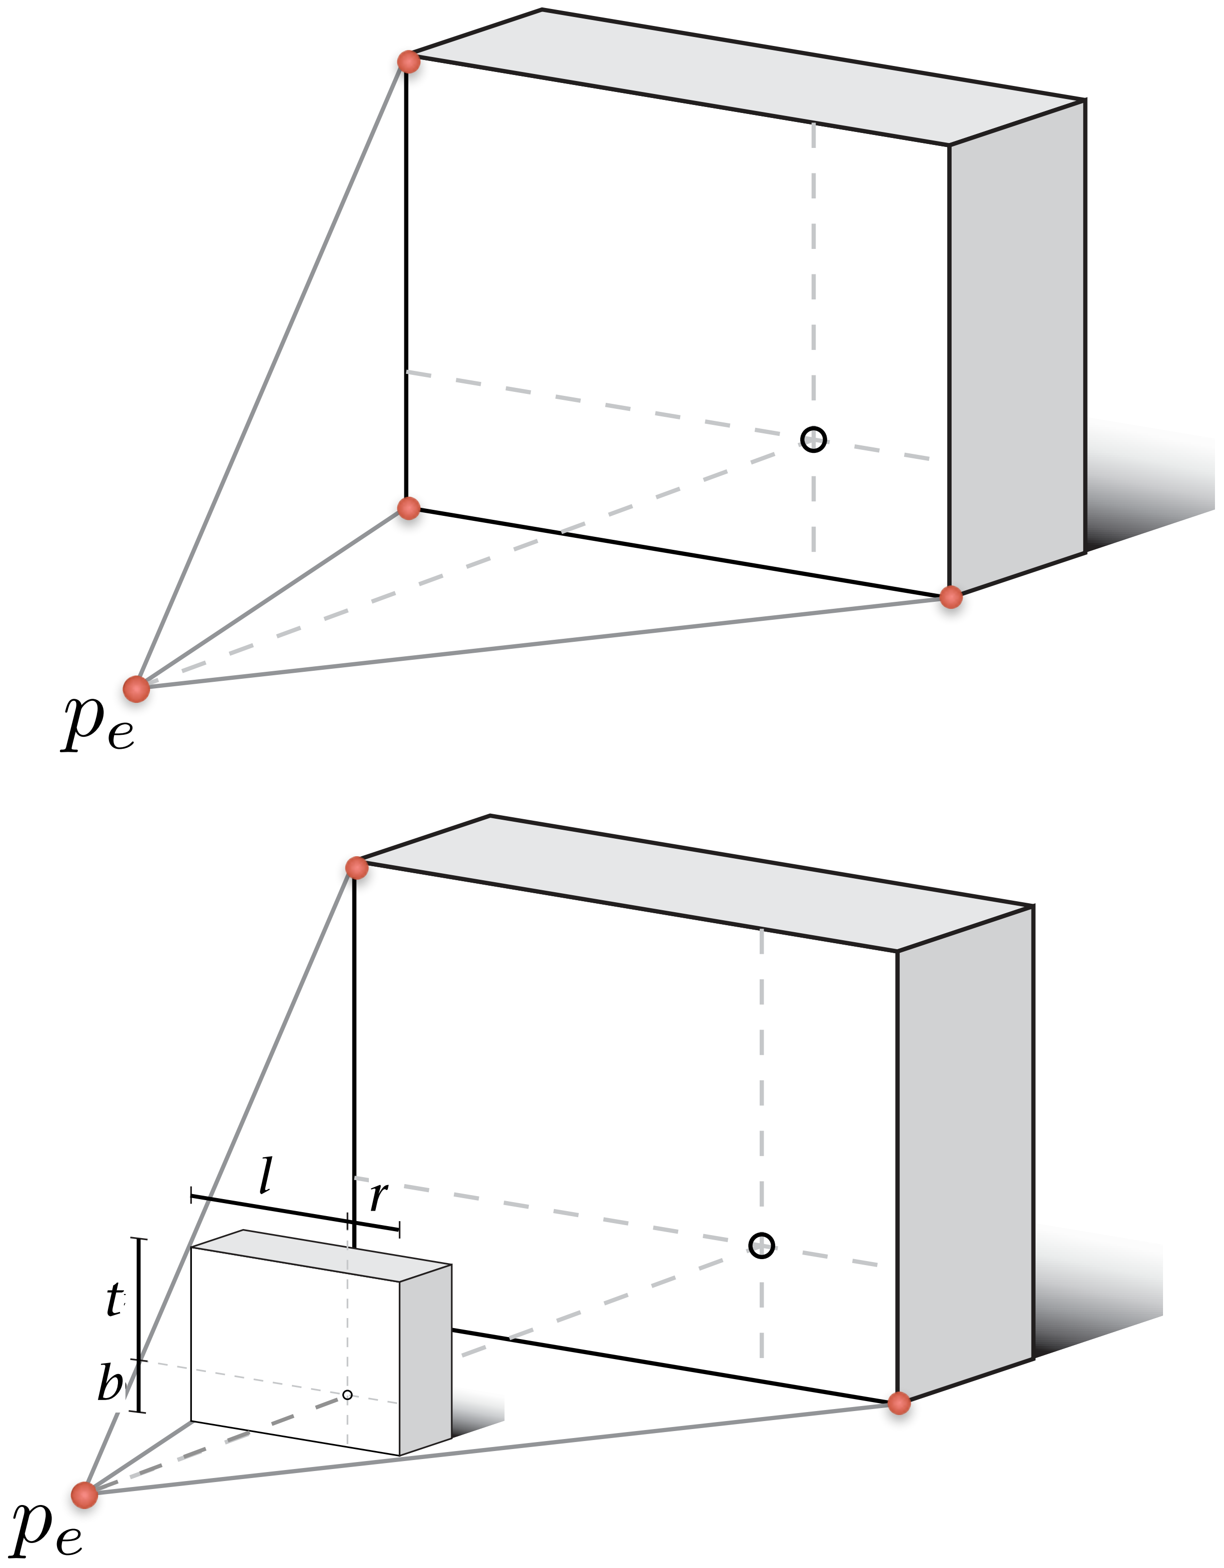

Computation of the View Frustum for IPT's

- Given: vertices of the IPT screen (e.g., powerwall) in virtual space $\longrightarrow$ vectors $\mathbf{v}_a$, $\mathbf{v}_b$, $\mathbf{v}_c$
- Assumption: ZPP passes through "virtual IPT screen" $\longrightarrow$ z_0 in virtual space $= -\mathbf{v}_n \cdot \mathbf{v}_a$
- Using similar triangles:

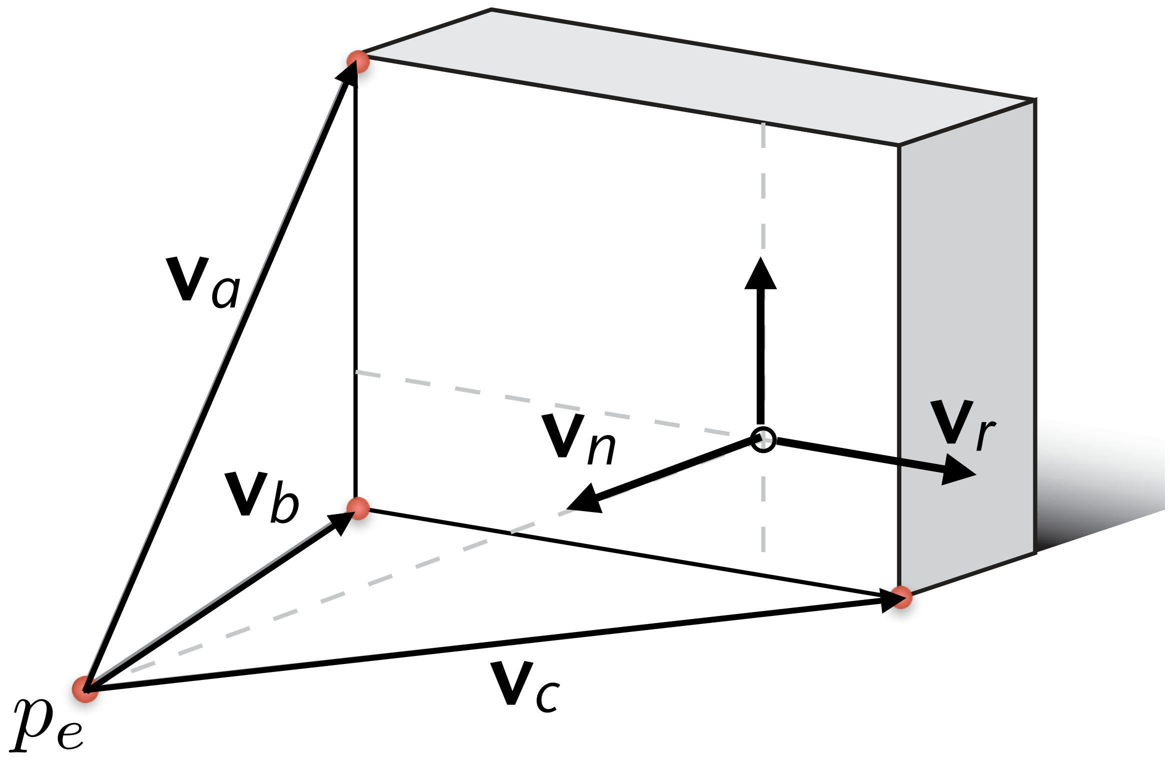

$$\frac{z_0}{n} = \frac{-\mathbf{v}_r \cdot \mathbf{v}_a}{l}$$

$$l = -\mathbf{v}_r \cdot \mathbf{v}_a \frac{n}{z_0}$$

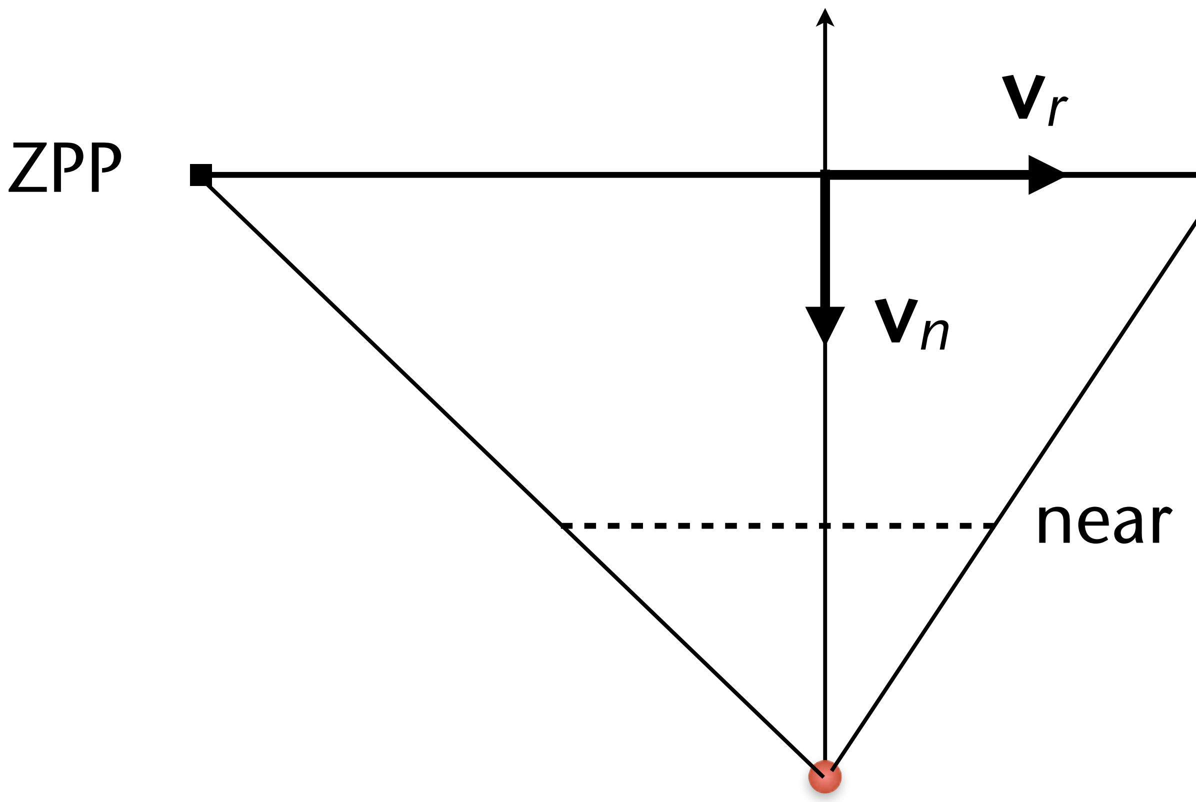

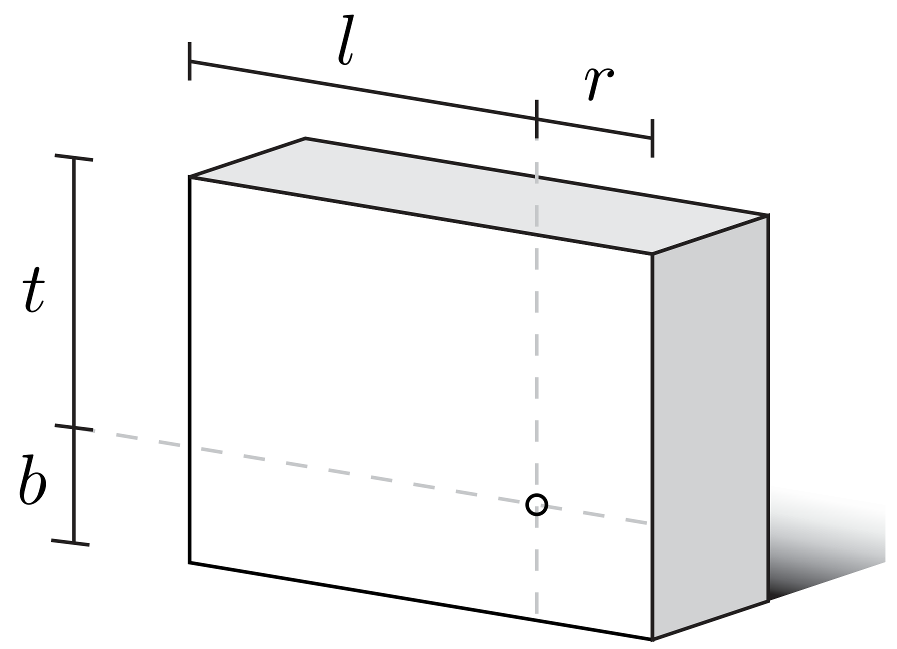

Alt. Computation of the View Frustum

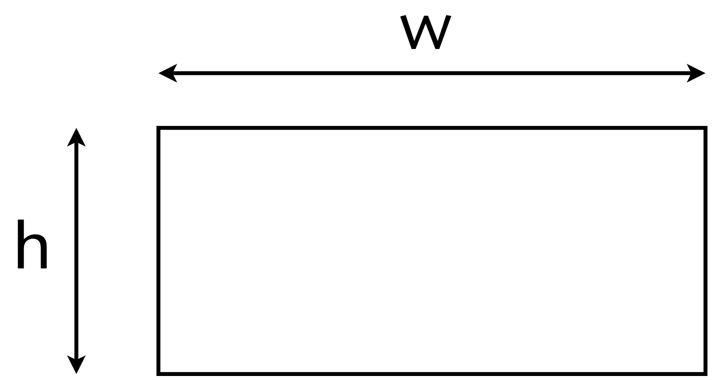

- Given: i = interpupilary distance ÷ 2, w/h = aspect ratio, α = horizontal FoV , n = near plane , z_0 = zero-parallax
- Task: determine **`left`**/**`right`**/**`top`**/**`bottom`**
- Assumption (for now): no head tracking
  → cyclop's eye is in front of the center of the viewport
- Example: compute **`left`** for left eye

$$l_c = n \tan \frac{\alpha}{2} \qquad l' = i \frac{z_0 - n}{z_0}$$

$$l = l_c + l' - i = l_c - i \frac{n}{z_0}$$

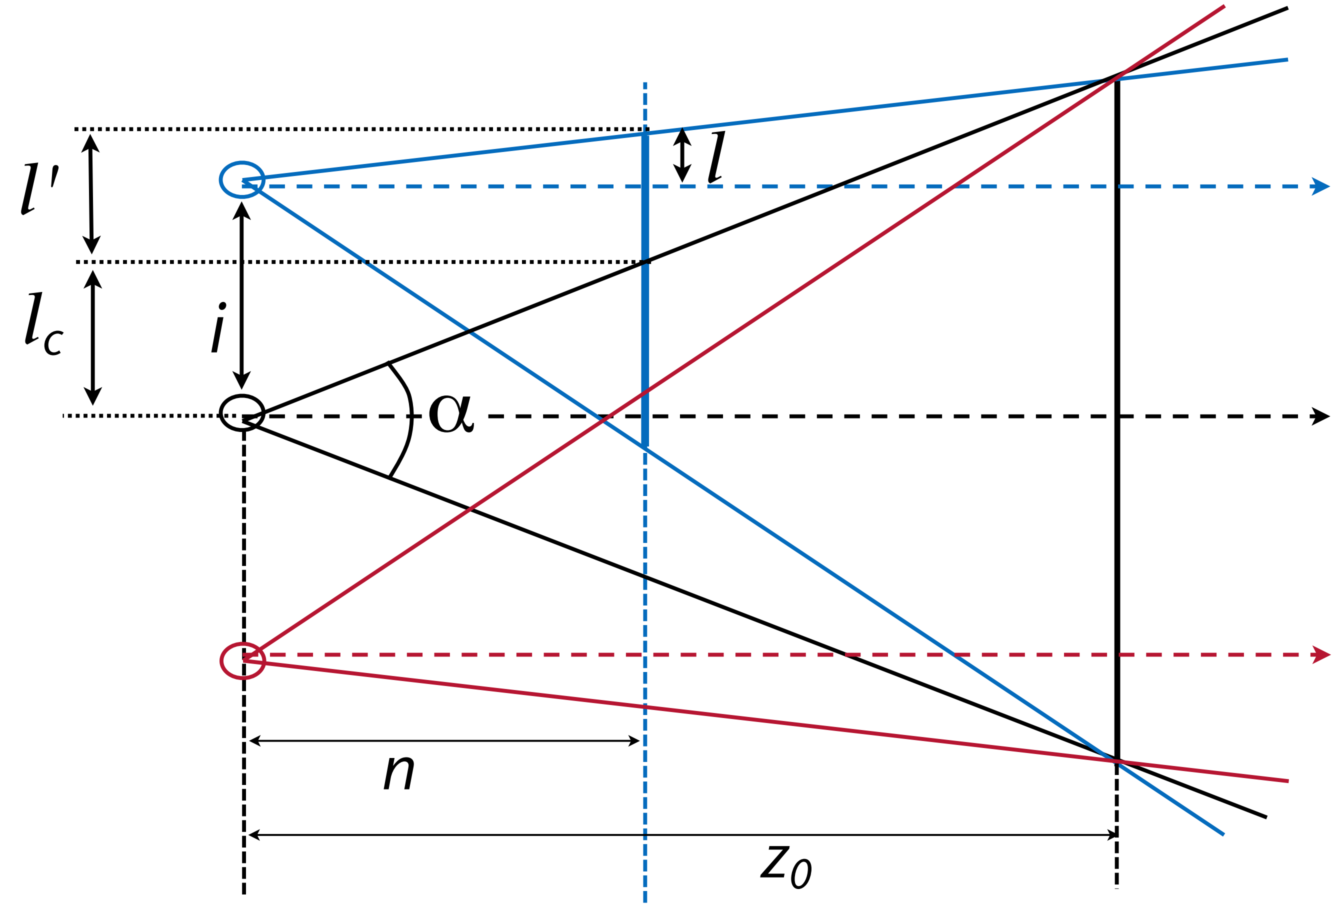

Hypo- and Hyper-Stereo

- In monoscopic filming/display, cameras just have these parameters:
  - Field-of-View, focal length (film), …
- In stereoscopic filming/rendering, (virtual) cameras have *in addition*:
  - Interaxial separation (= IPD)
  - Zero-parallax plane
- Hypo-Stereo: Interaxial < real IPD ⟶ dwarfism effect
- Hyper-Stereo: Interaxial > real IPD ⟶ gigantism effect
- Can make sense for macro/micro scenes

Interaxial Separation between lenses, a.k.a. Stereo Base, a.k.a. Interocular separation, (IPD for human eye)

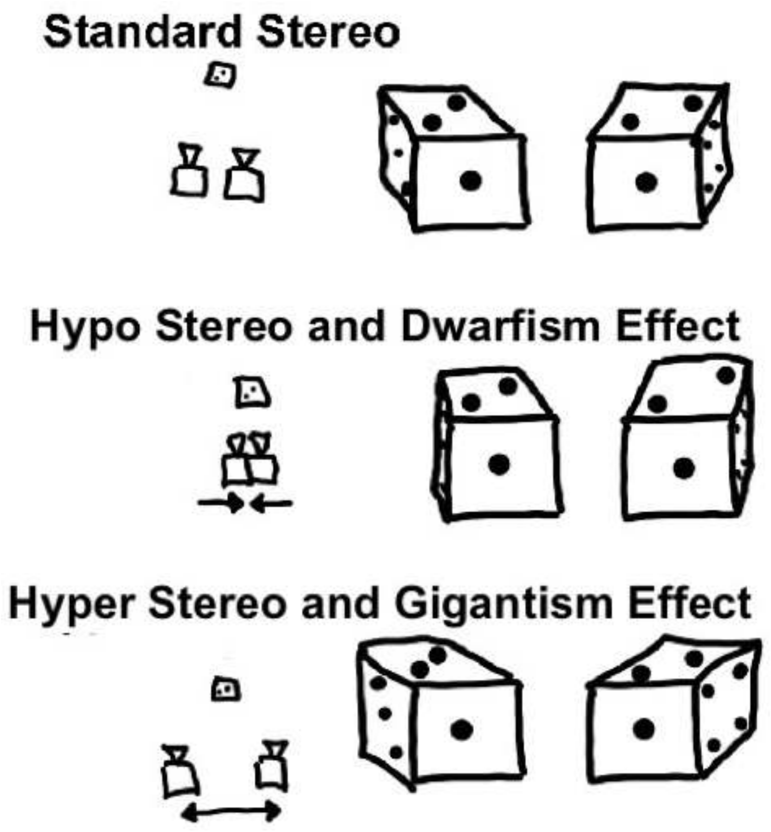

A Warning About Parallax

- Careful not to create too much parallax!
- Assume you created a stereo image for a small desktop display. Then, you run the app on a big screen:

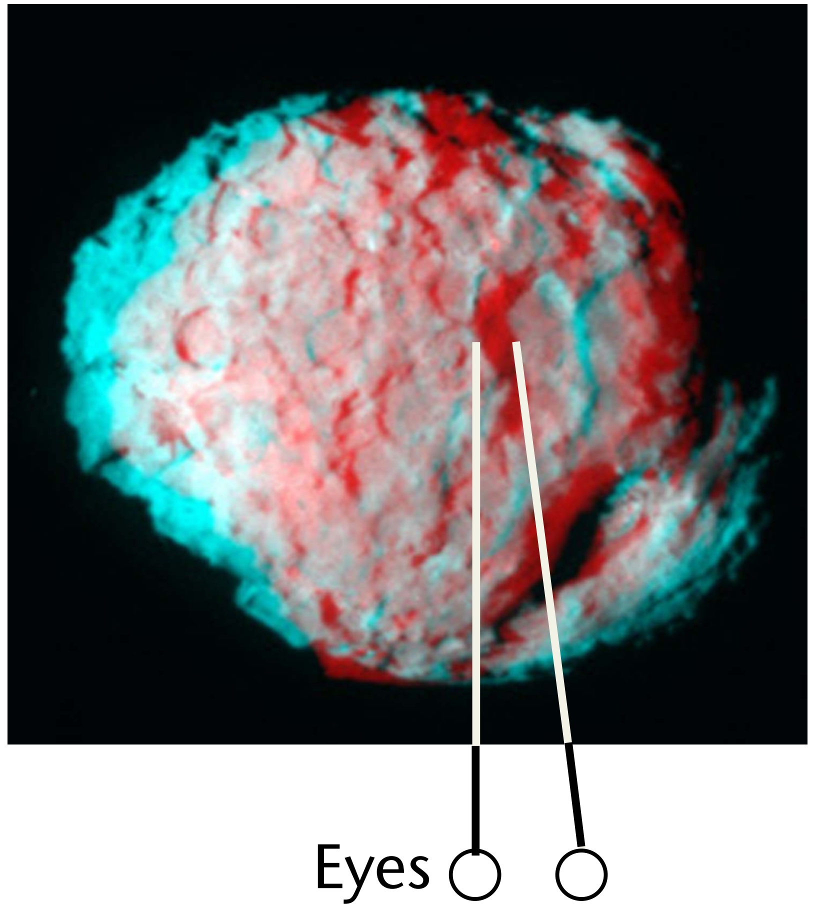

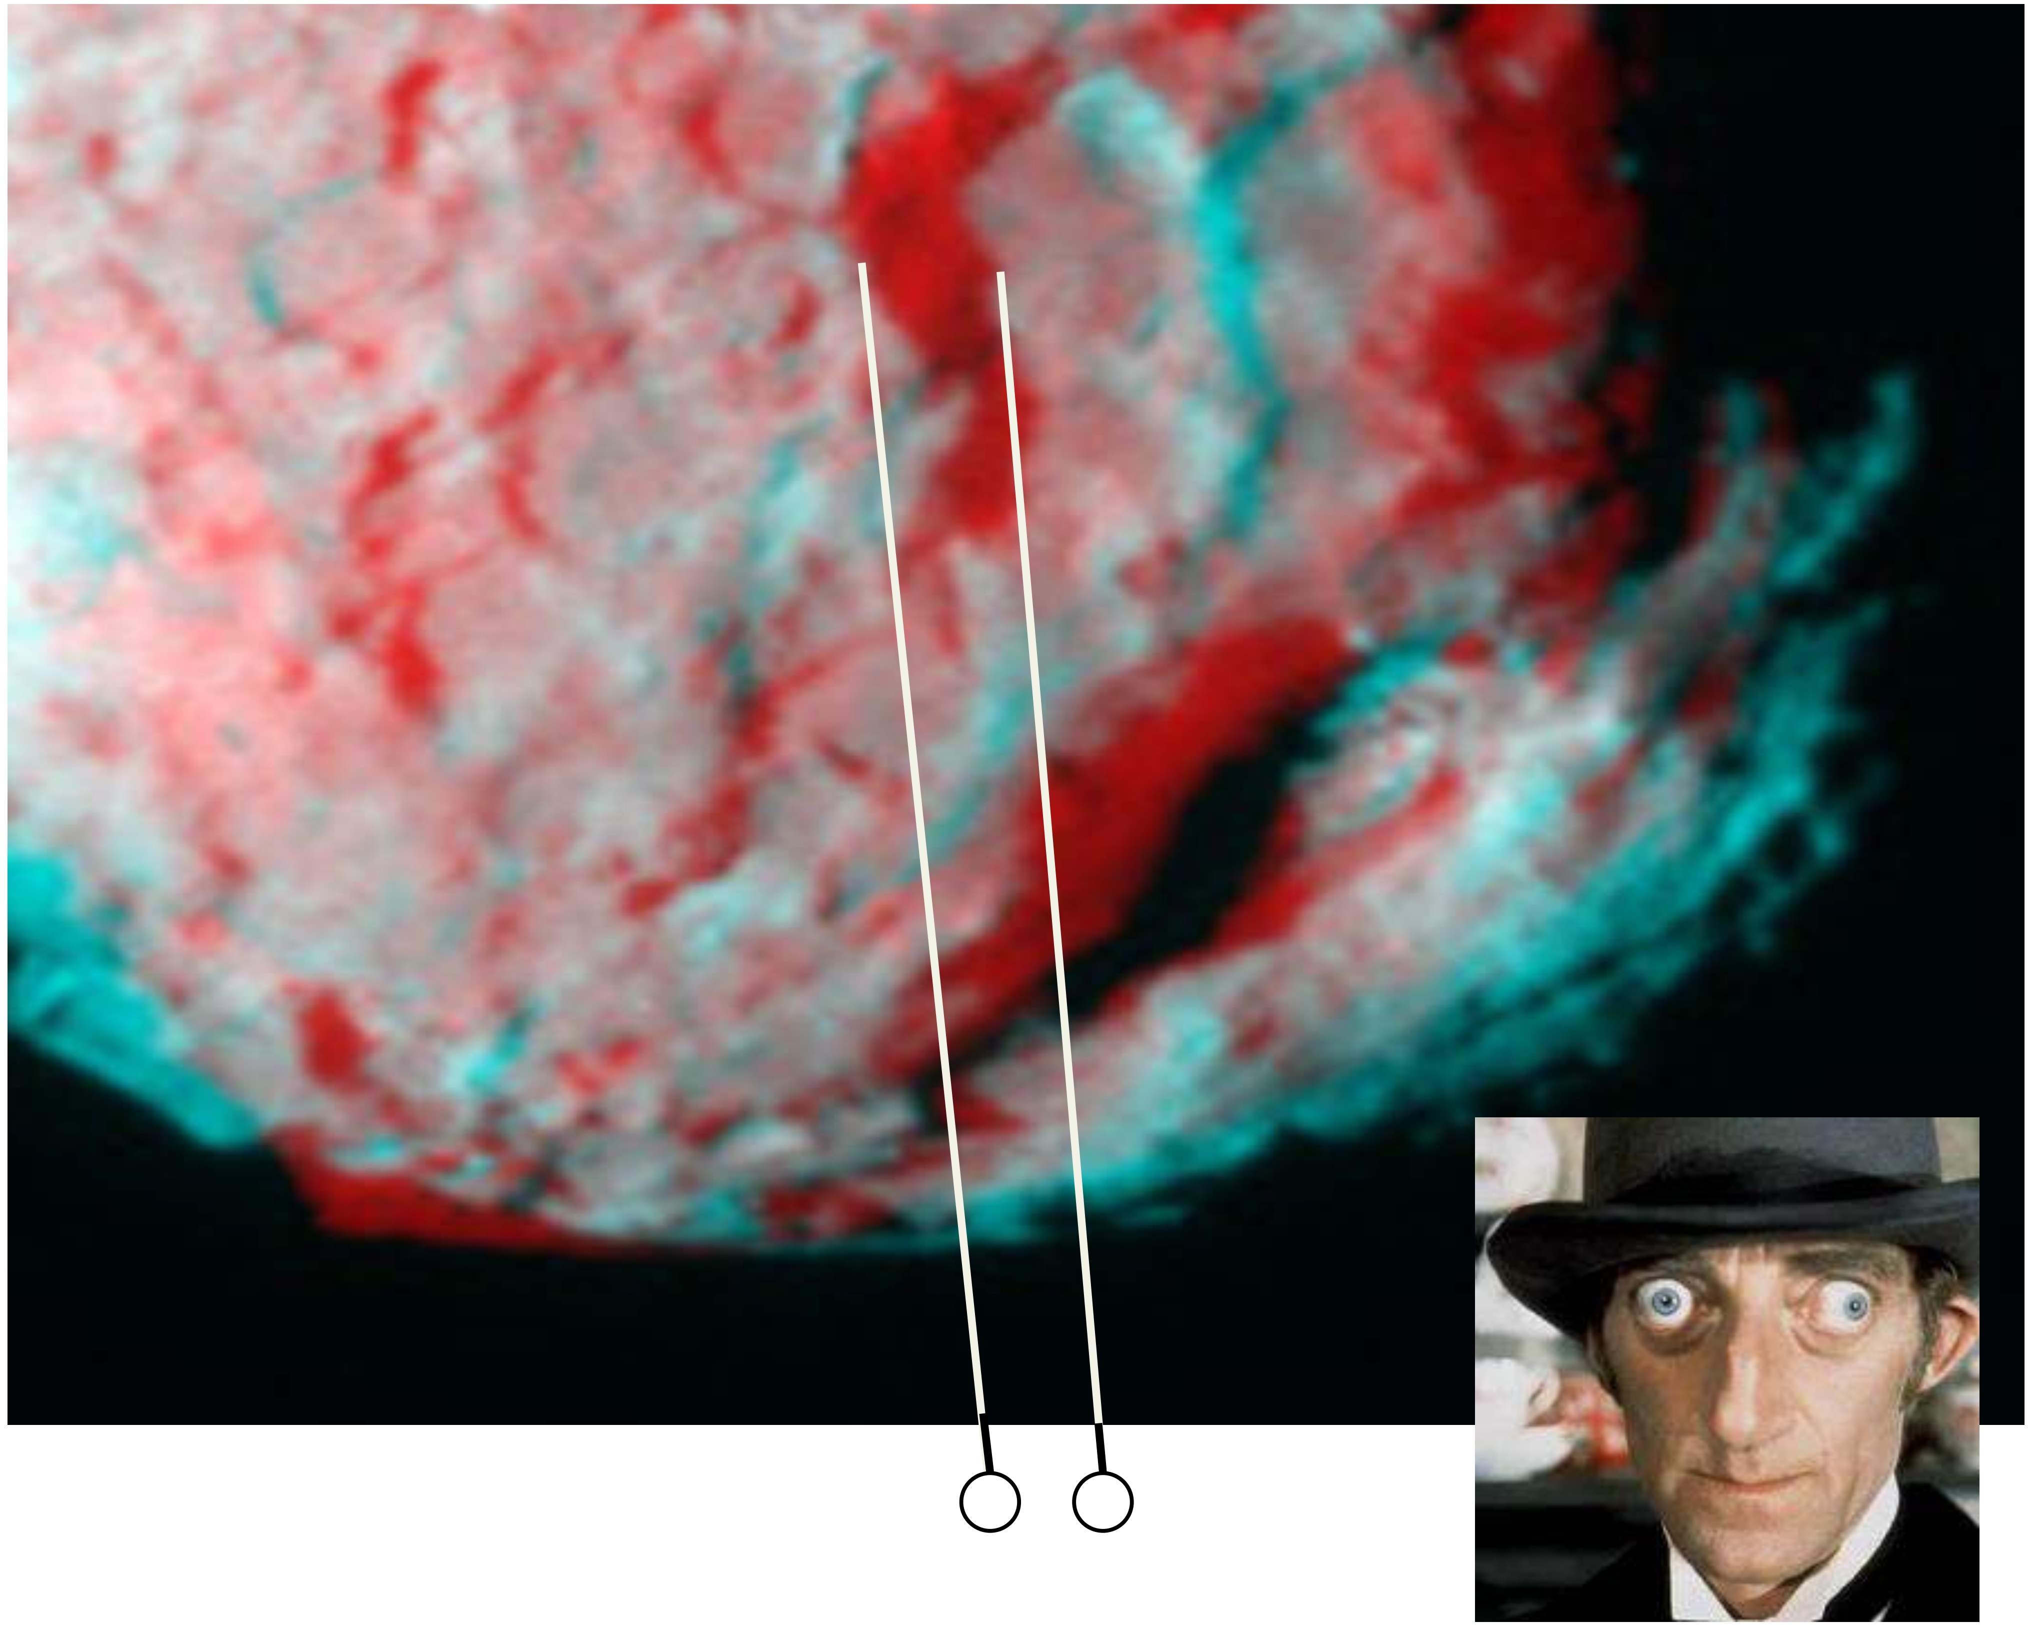

Optimization of Stereo Rendering Performance

- Observation about HVS can be used to optimize rendering performance
  - Objects in the distance can be rendered monoscopically (just once)
  - Only near objects need to be rendered twice
- Approach: 1. near objects in stereo, 2. far objects in mono, 3. composite 4. transparent objects, 5. post-processing

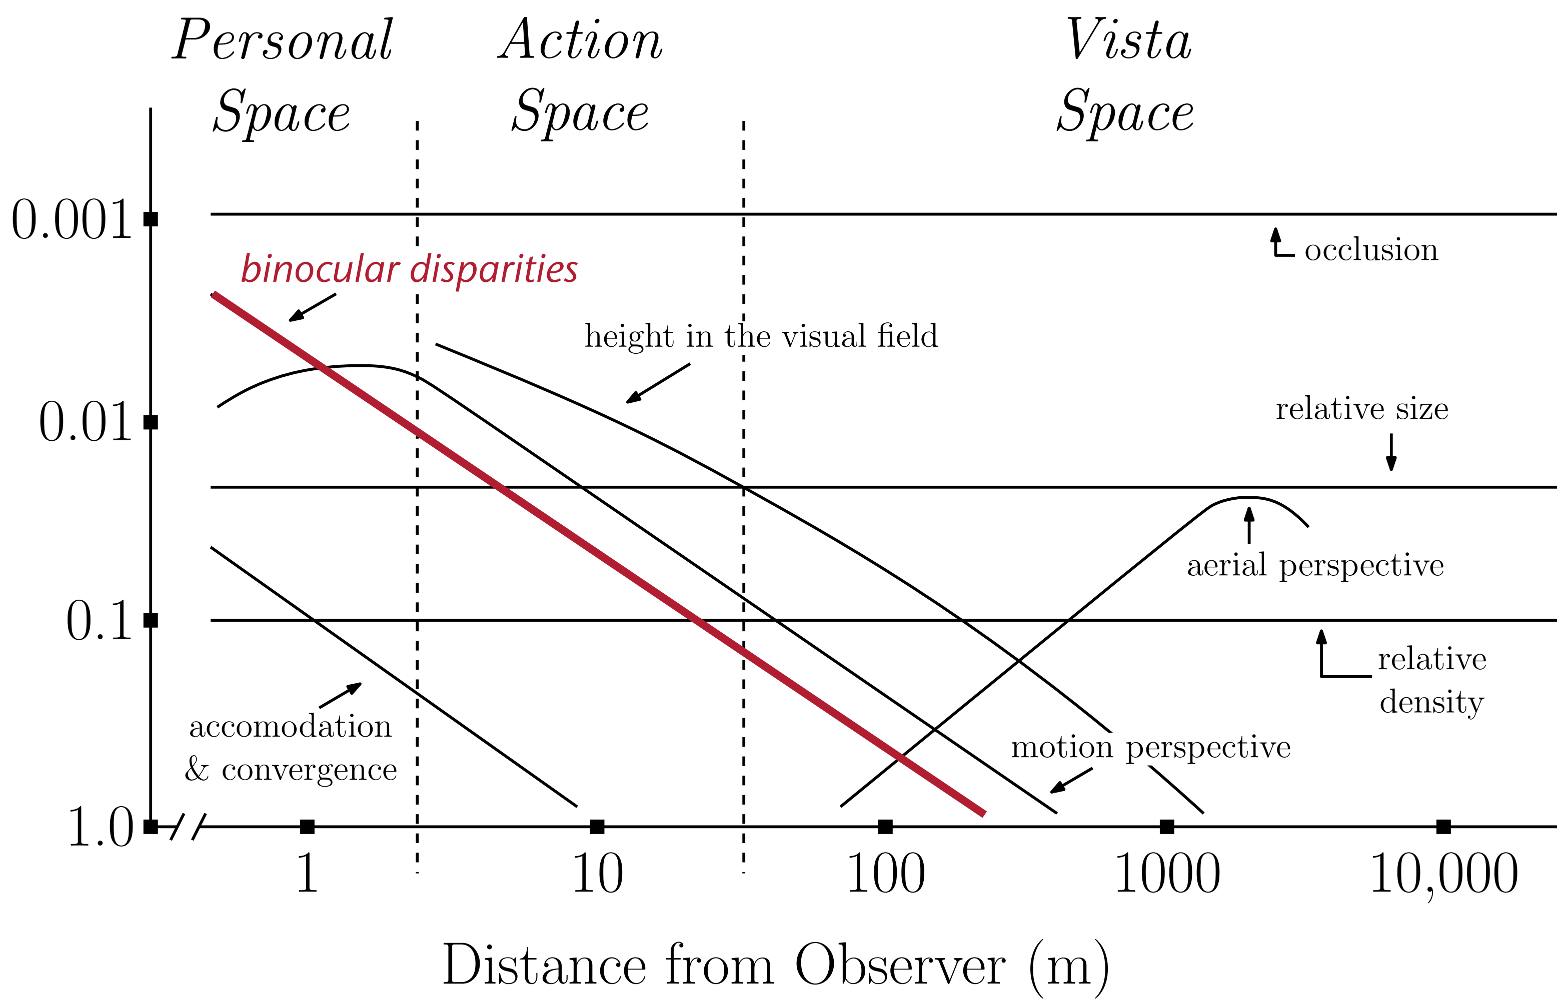

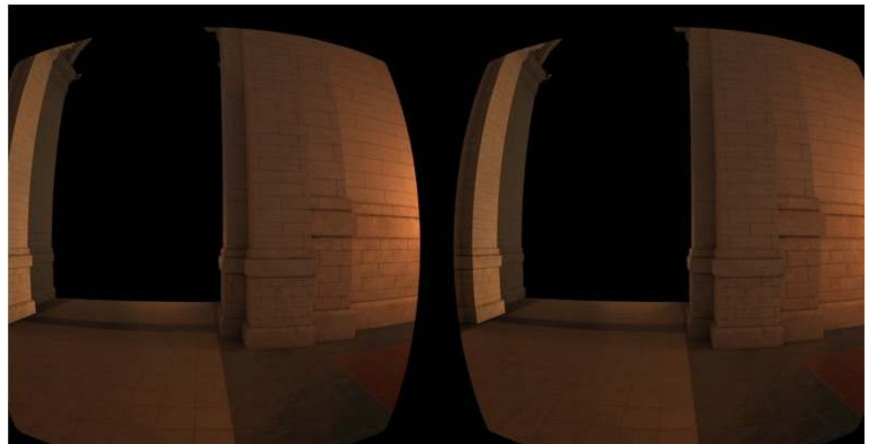

1.

2.

3.

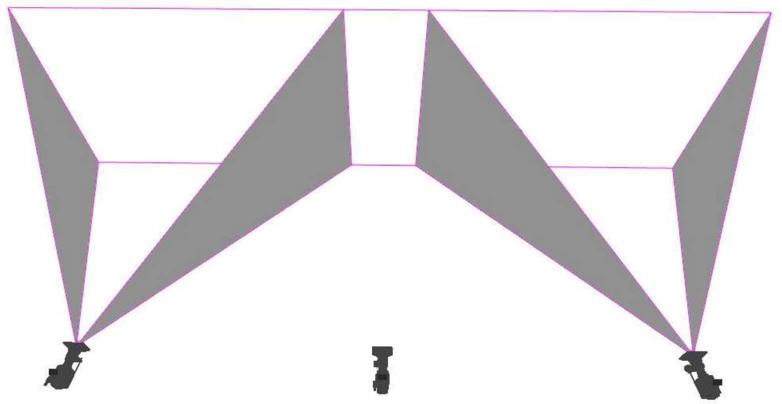

Updating the Virtual Camera for IPT's (e.g., Powerwall)

Imagine a user, standing in front of a (stationary) display showing virtual (stationary) objects.
Imagine the user walking sideways in front of the display.

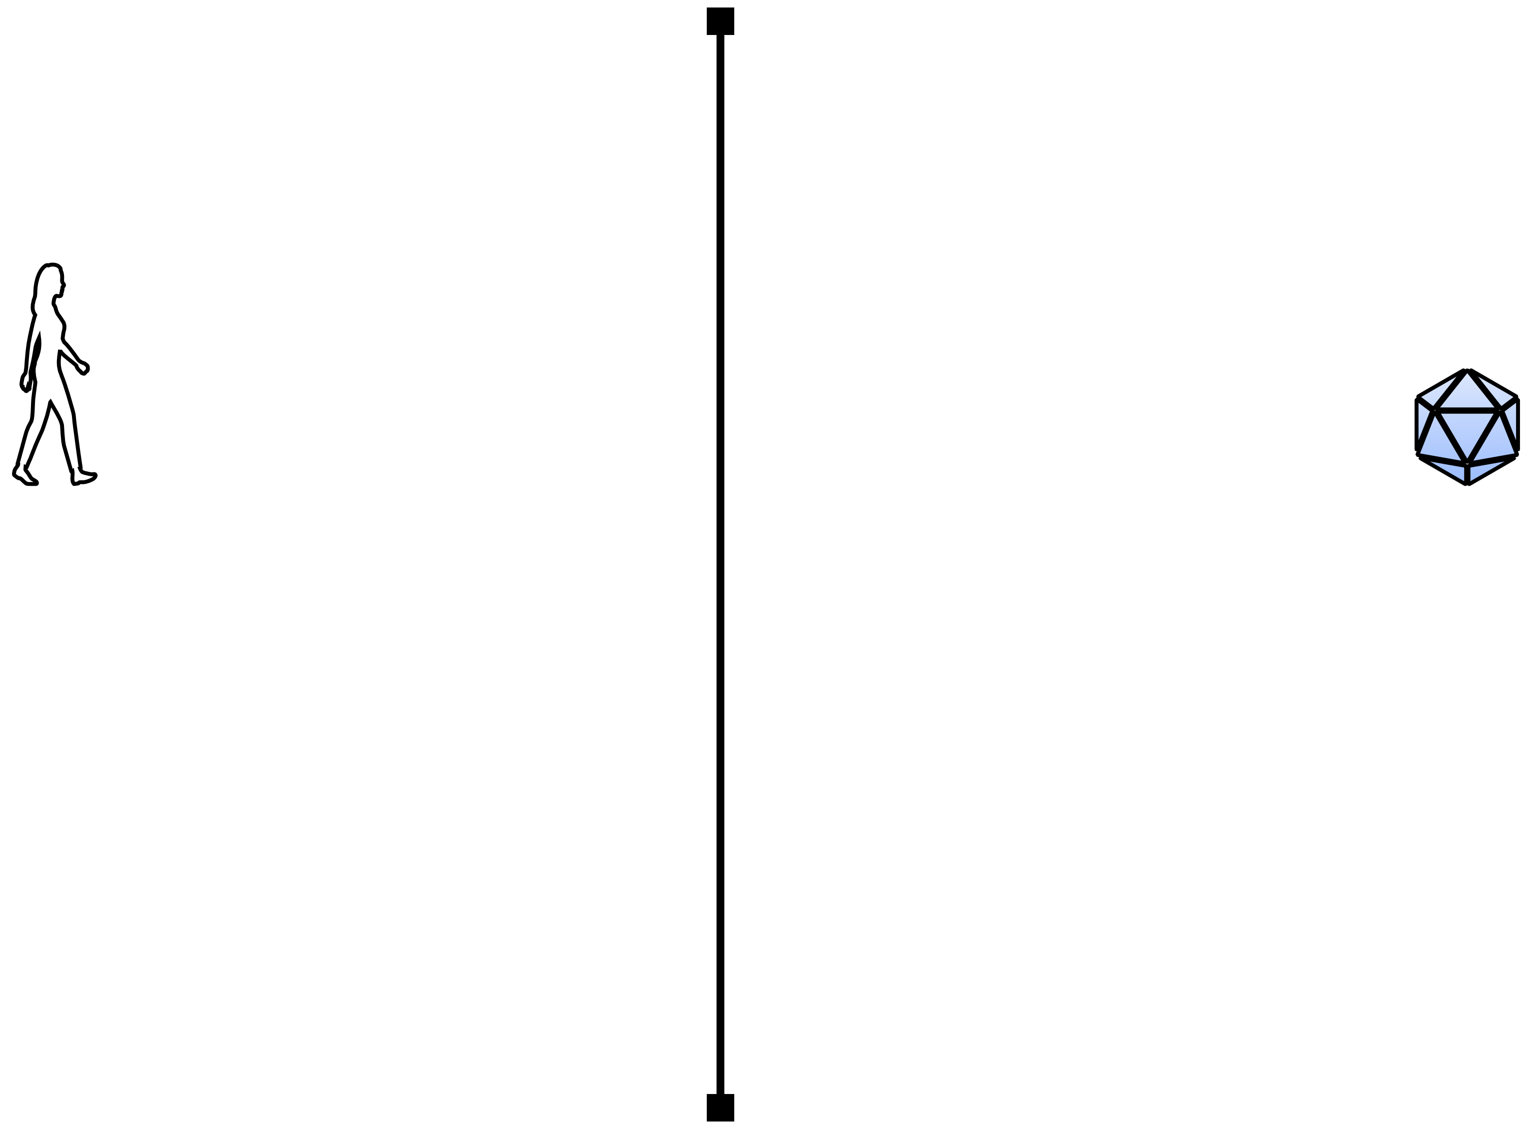

→ Consider the stationary display a window into the virtual world

FYI: Stereo Rendering in Films

- Movie directors use all kinds of "tricks" involving stereo
- Problem: scenes with very large depth range, e.g., hero character at center and close to camera, but background is very far away and also visible
  - Effect with regular stereo rendering: background gets too much parallax $\longrightarrow$ eye strain for viewers
  - One solution: reduce IPD, so that parallax in background is OK
  - Negative consequence: hero character in foreground no longer gets as much parallax $\longrightarrow$ does not look as "round" as before
- Disney Studio's trick:
  - Render foreground objects using one IPD and ZPP distance
  - Render background objects with *another* set of IPD/ZPP

Guidelines for Stereo Rendering

1. Do not make parallax too big! (common error of novices)
   $\pm 1.6° \rightarrow$ parallax $\leq 0.03 \cdot$ (distance to projection wall)
2. Single object → put zero-parallax plane in its center
3. Extended surroundings → 1/3 negative parallax, 2/3 positive parallax
4. Keep objects with negative parallax away from the border of the projection surface

The View Frustums (Frusta) for HMD Setups

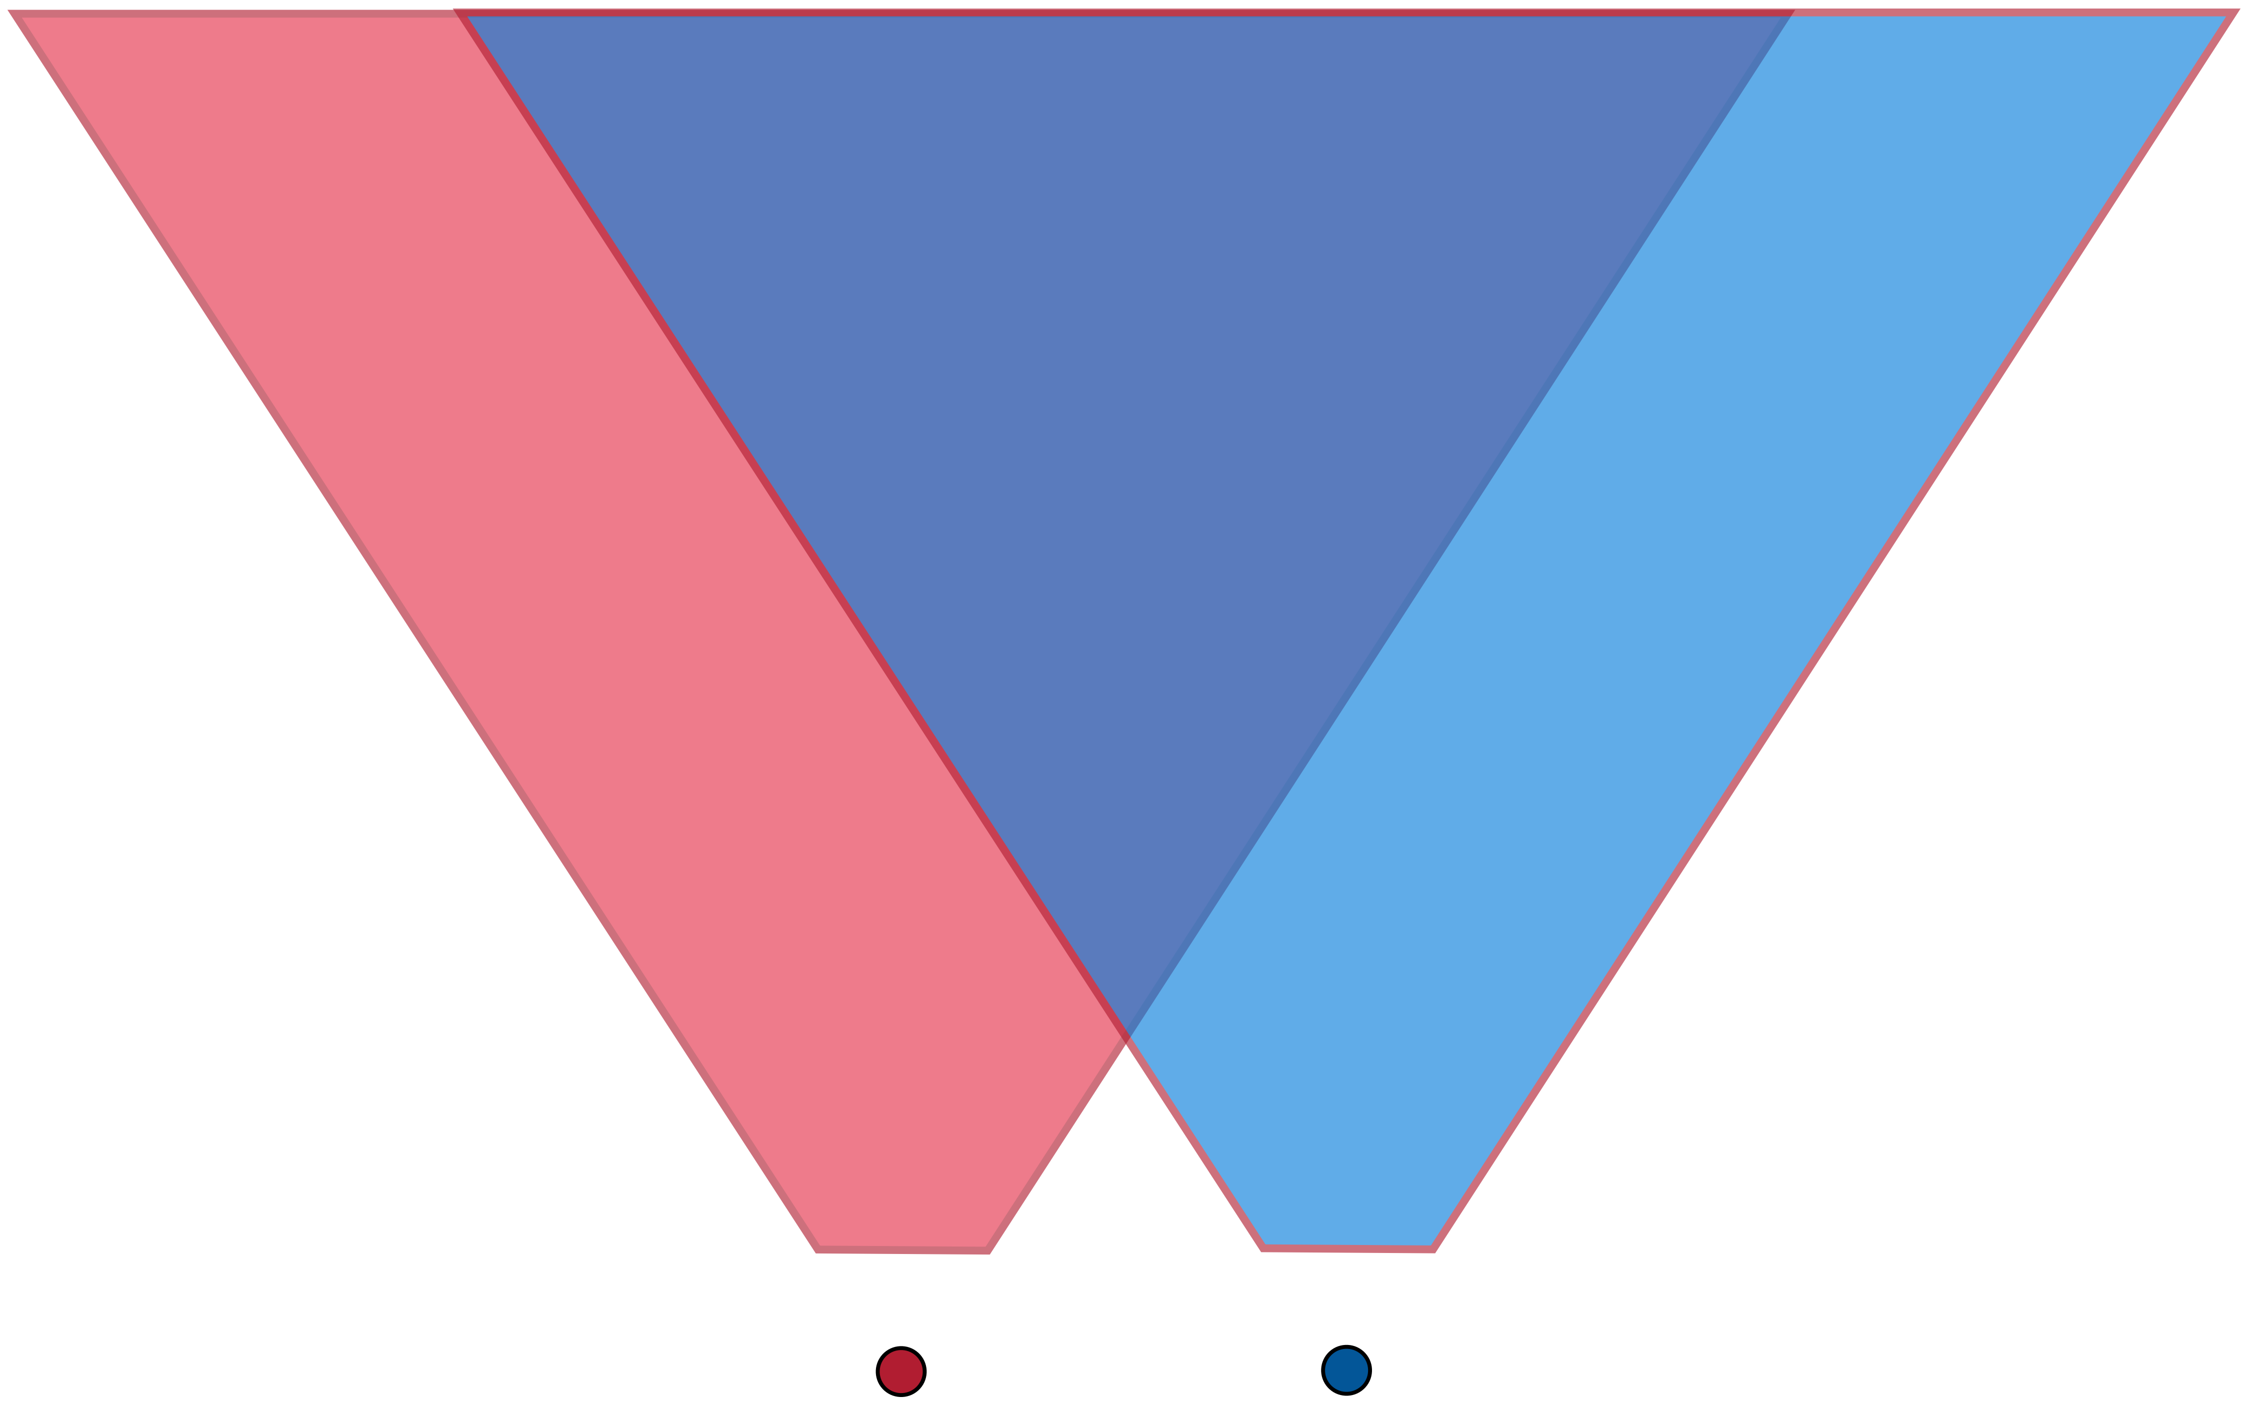

Computing Precise Viewpoints

M^l_e = viewpoint transformation

M_s = current sensor pose, relative to world coordinates

$M_{r \leftarrow s}$ = transformation from head's rotational center to tracked position on user's head

$M_{e \leftarrow r}$ = transformation from "cyclop's eye" to head's rotational center

$T^l \mid T^r$ = translation to left/right eye

Concatenation of all transforms:

$$M^l_e = T_l M_{e \leftarrow r} M_{r \leftarrow s} M_s$$

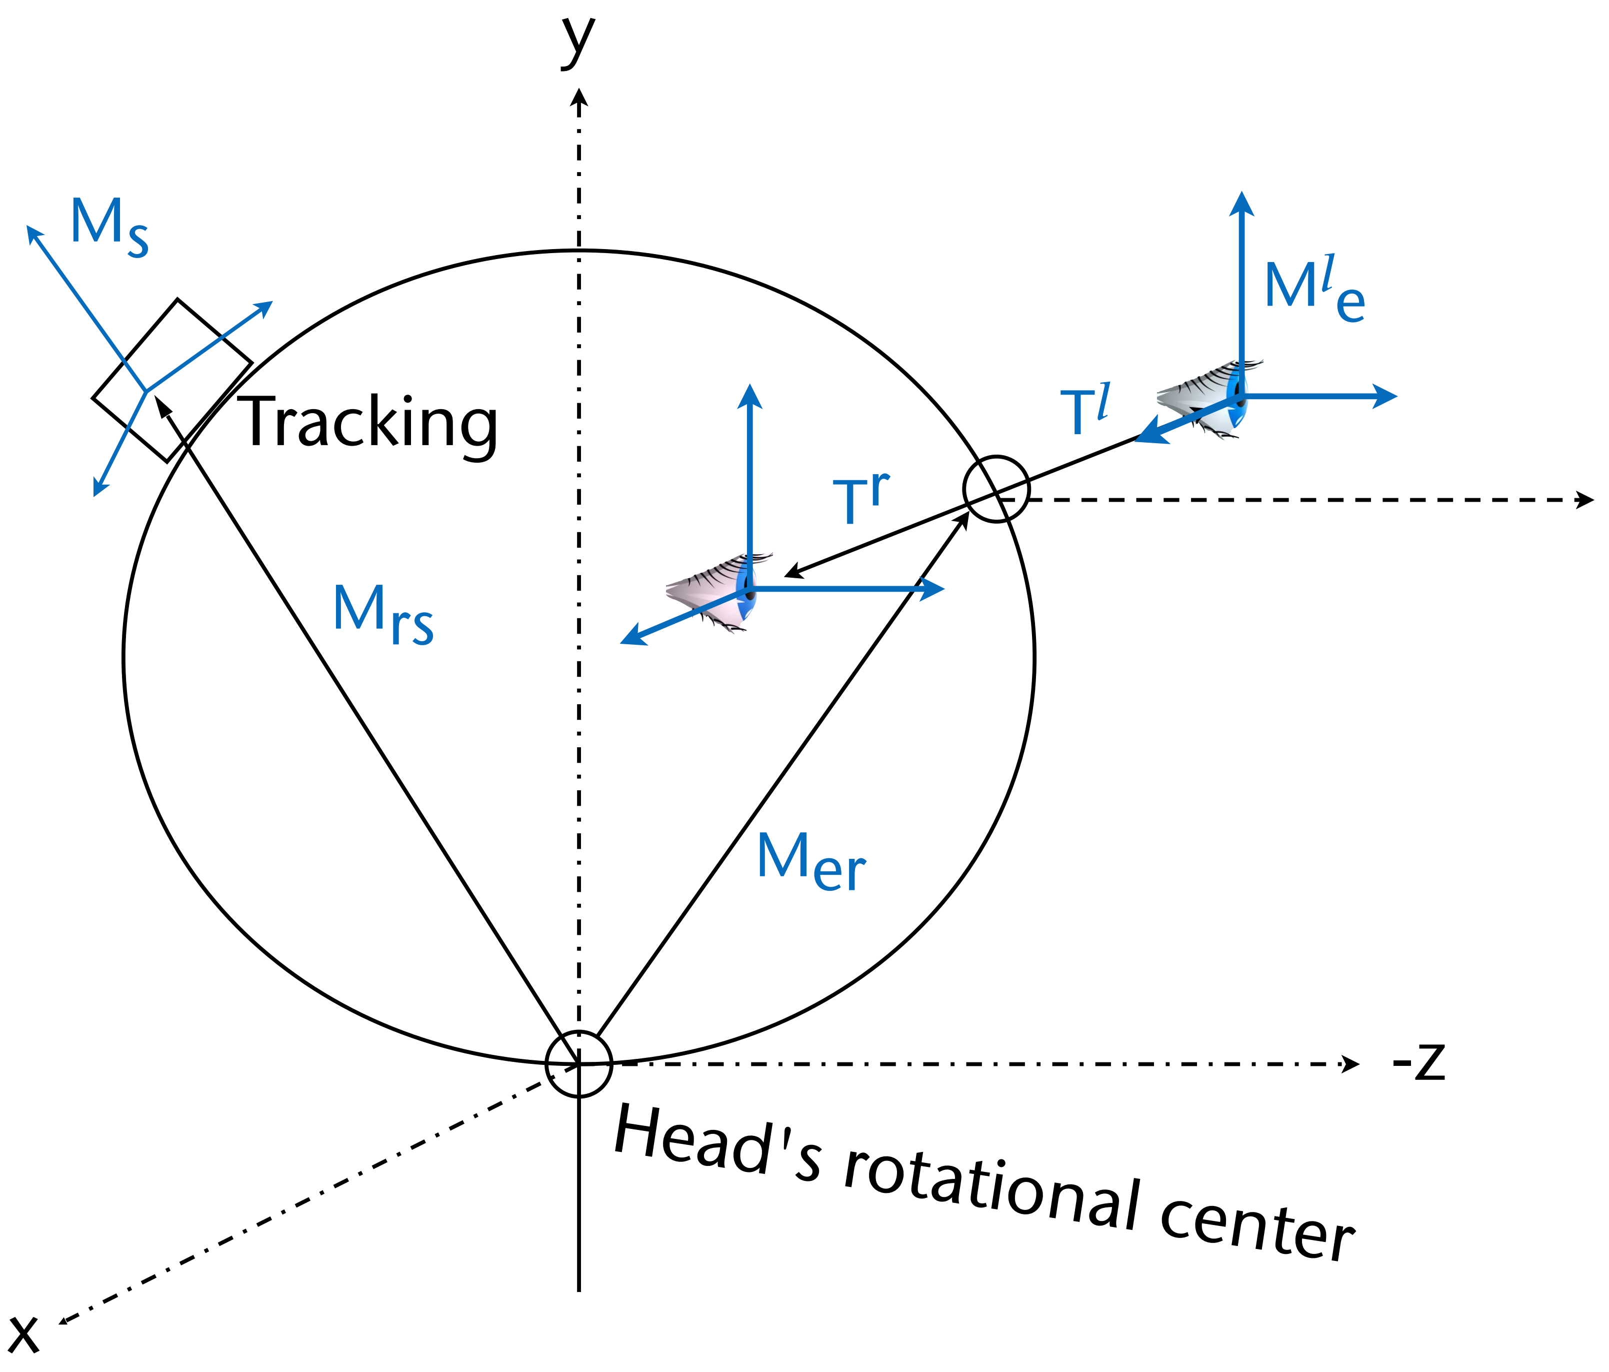

Problems with Stereo Rendering: Depth Aliasing

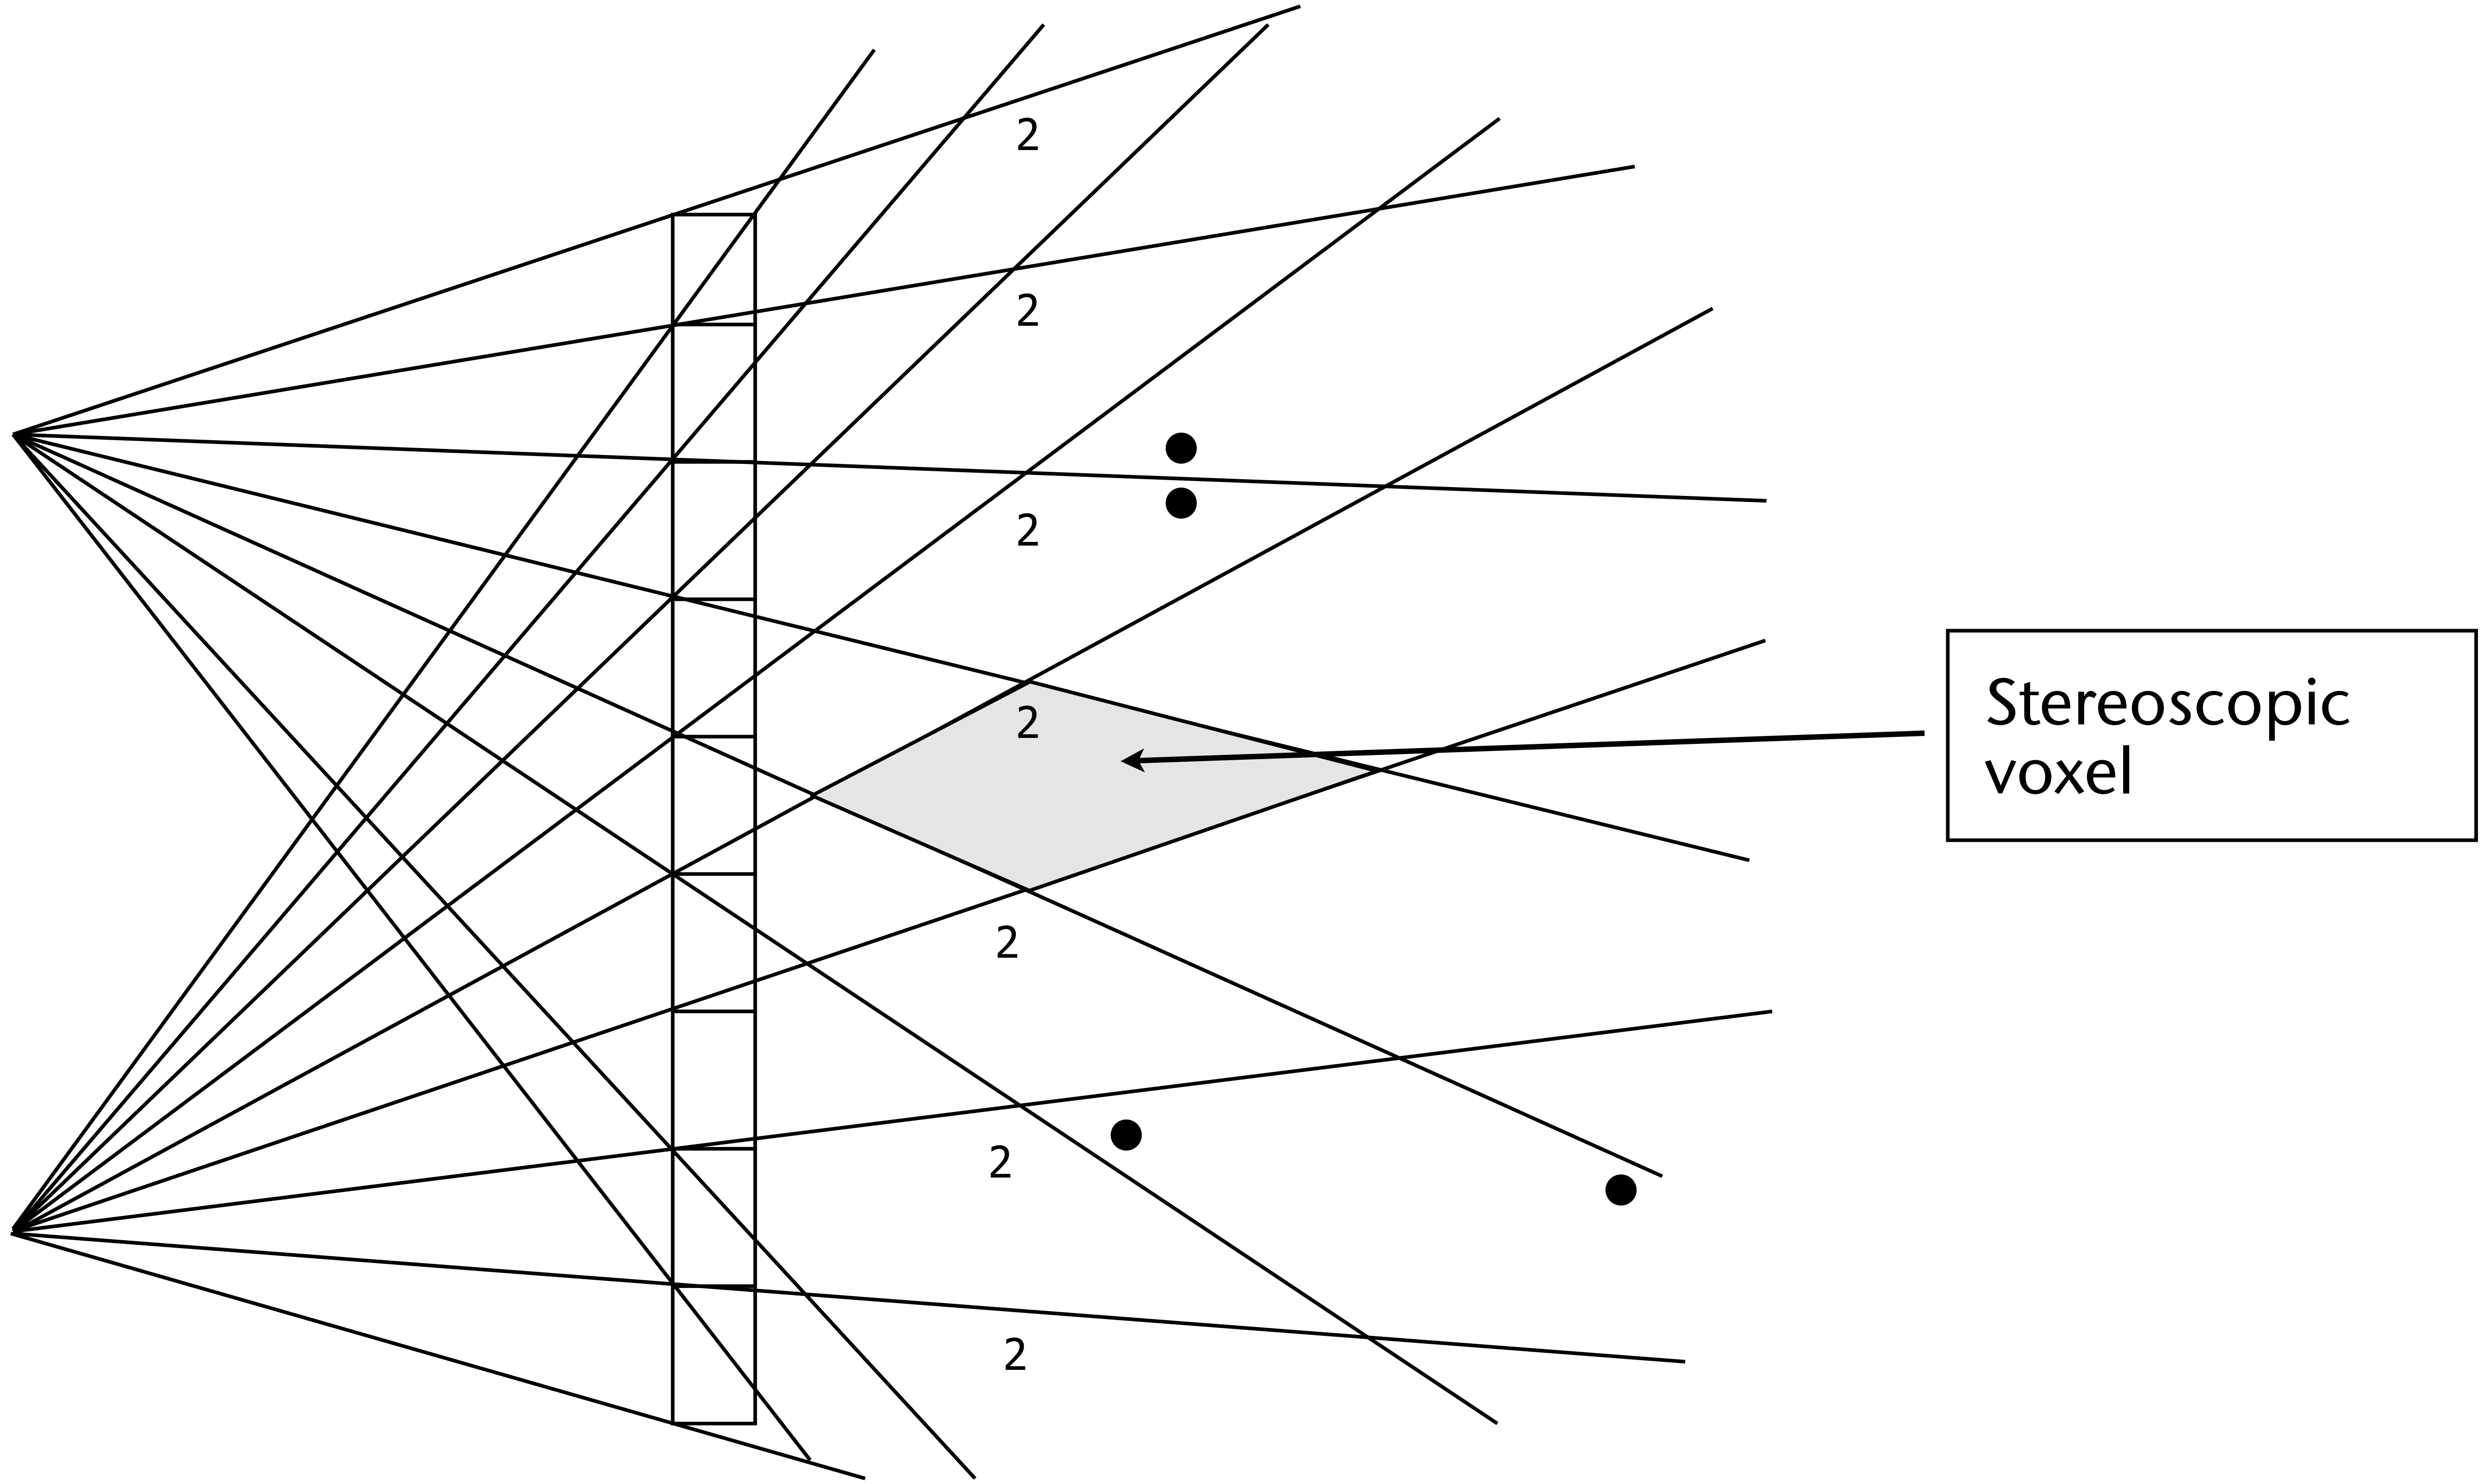

Convergence-Focus Conflict

- Experimental evidence shows: the brain computes a weighted average of multiple depth cues (integration), including focal depth
- With stereoscopic projection displays, our eyes receive inconsistent depth cues
- Effect: in a Cave or Powerwall, ..
  - near objects appear more distant than they are (over-estimation)
  - far objects appear closer than they are (under-estimation)!

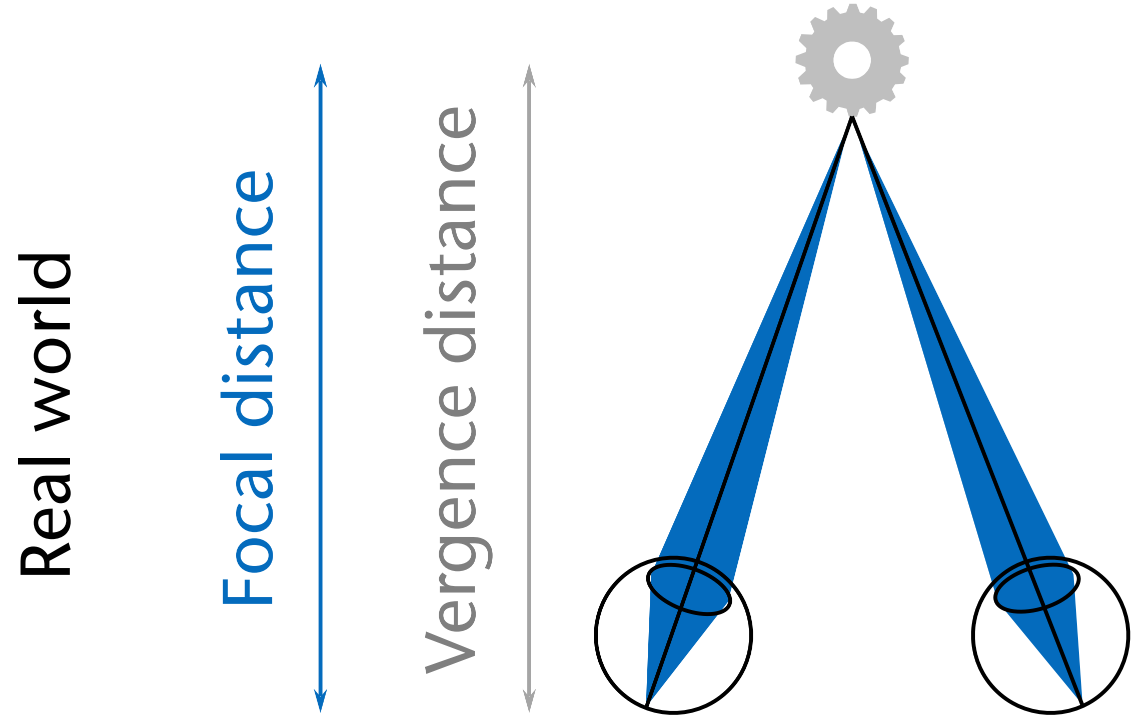

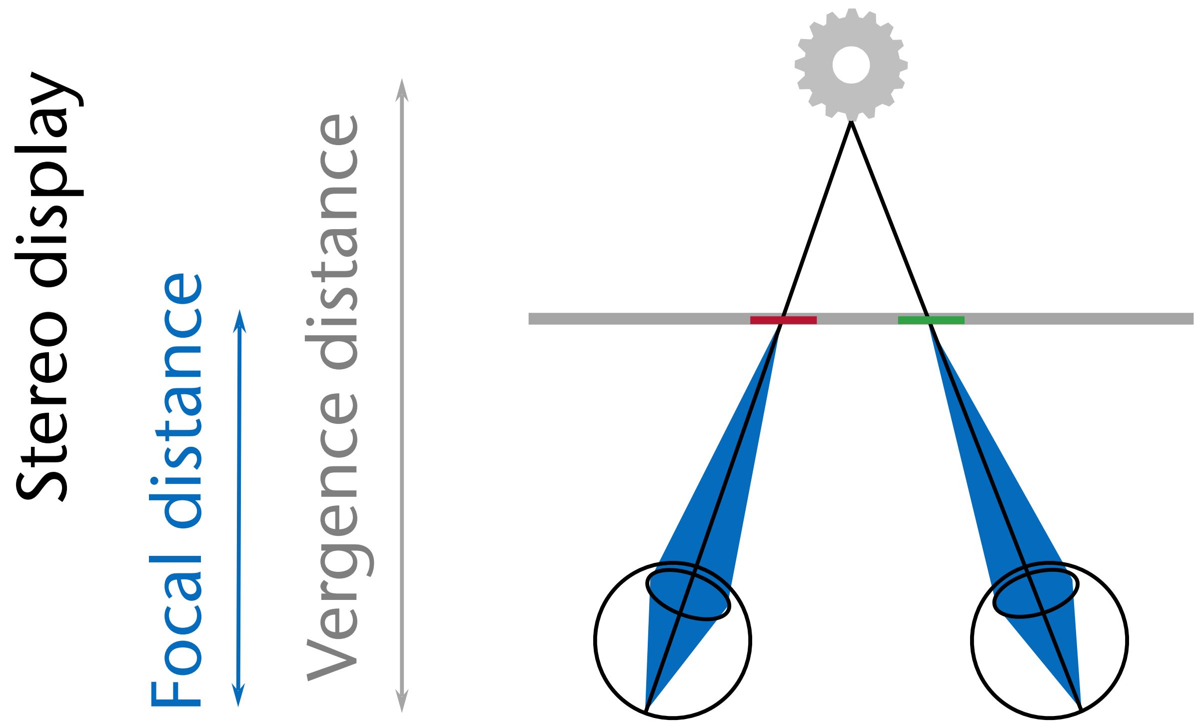

Blur Divergence

- Another depth cue: blur
  - The eye (brain) can estimate (relative) depth from the amount of blur
- If no depth-of-field is being rendered, then our eyes perceive different depth cues:

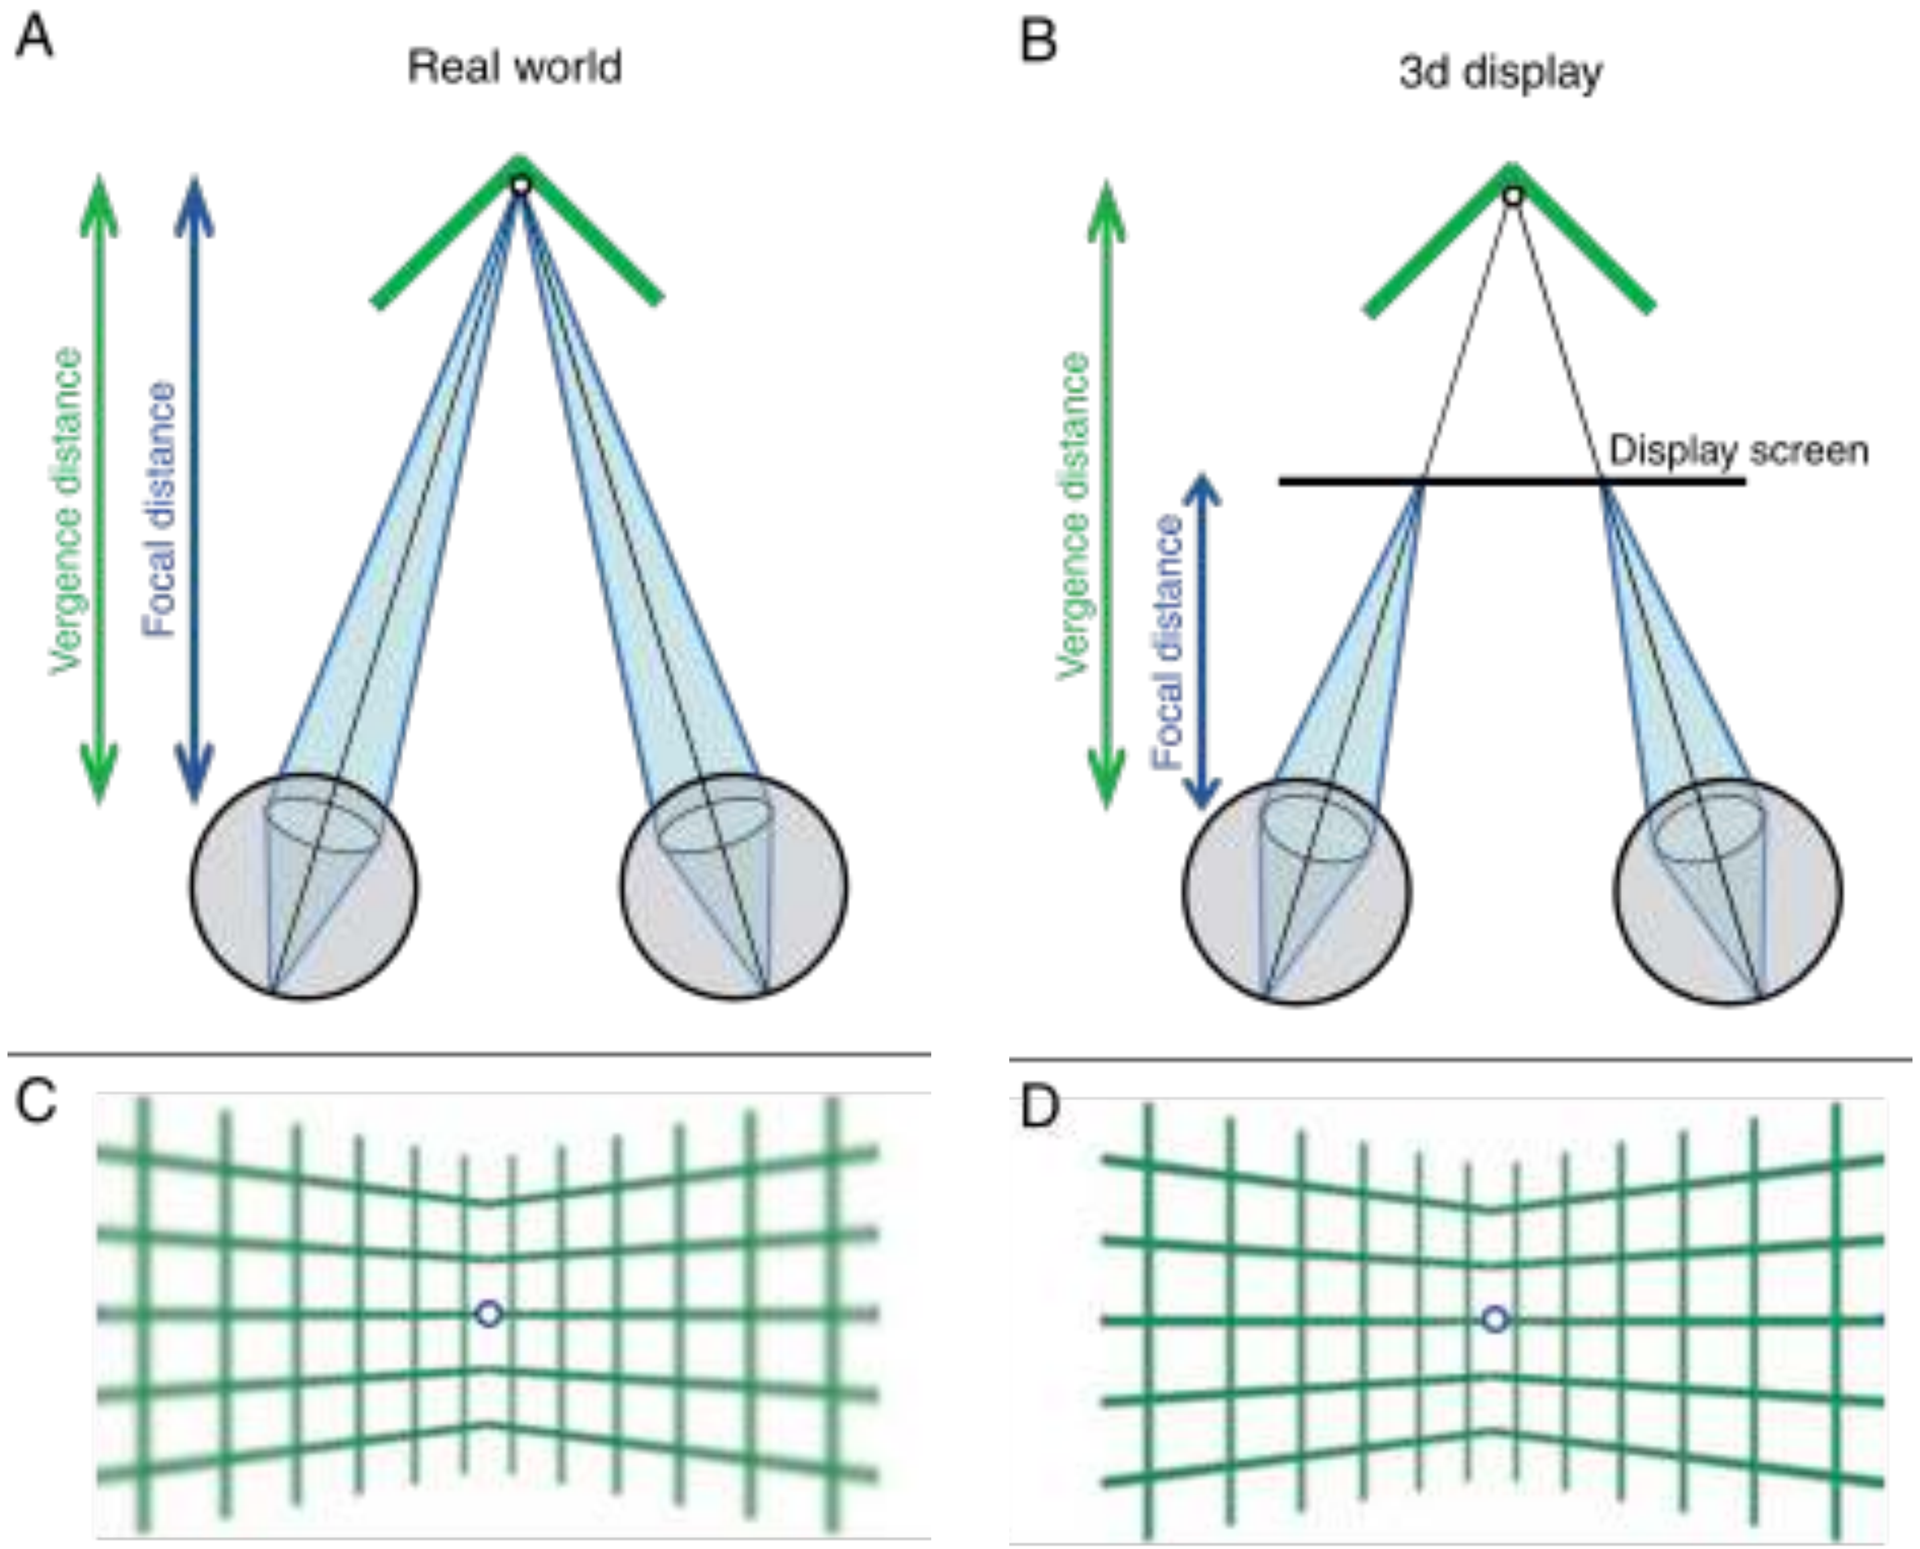

Stereo Window Violation (short *Stereo Violation*)

- Two effects that can occur together:
  1. *Clipping*
  2. *Negative parallax* from stereoscopic image
- Problem:
  - Object is clipped, although apparently *in front of* the projection surface!
  - Consequence: conflicting depth cues ⟶ stereo violation (a.k.a. window violation)
- Example: lower left corner of the anaglyph mars image

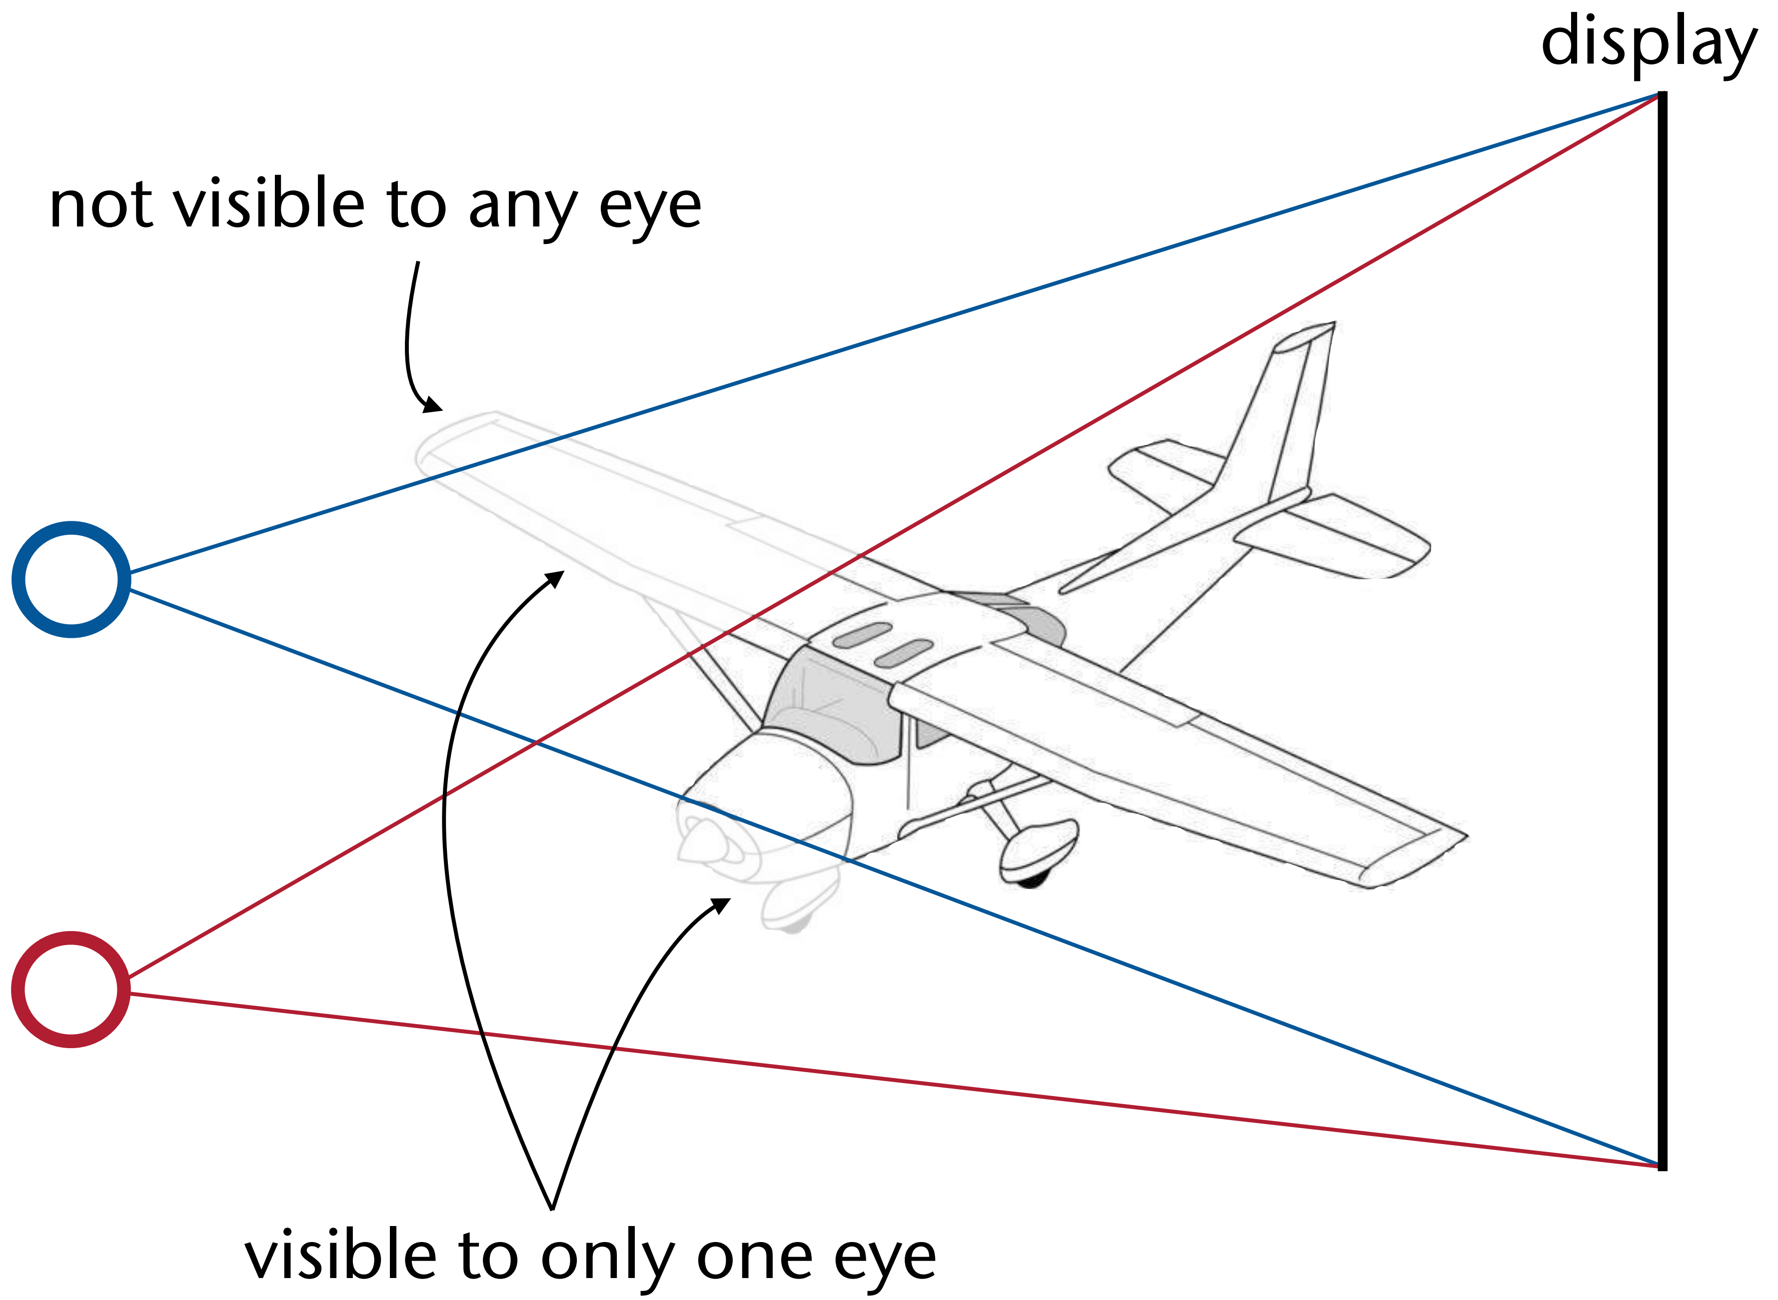

Stereoscopic Rendering is (Usually) a "One Man/Woman Show"

- Why is a single pair of stereoscopic images correct only for 1 viewpoint? (More or less distortions for all the others!)
- One of the problems: images (e.g., on a powerwall) shift and deform for the un-tracked user when the tracked user moves
- Similar problem, if user tracking is incorrect

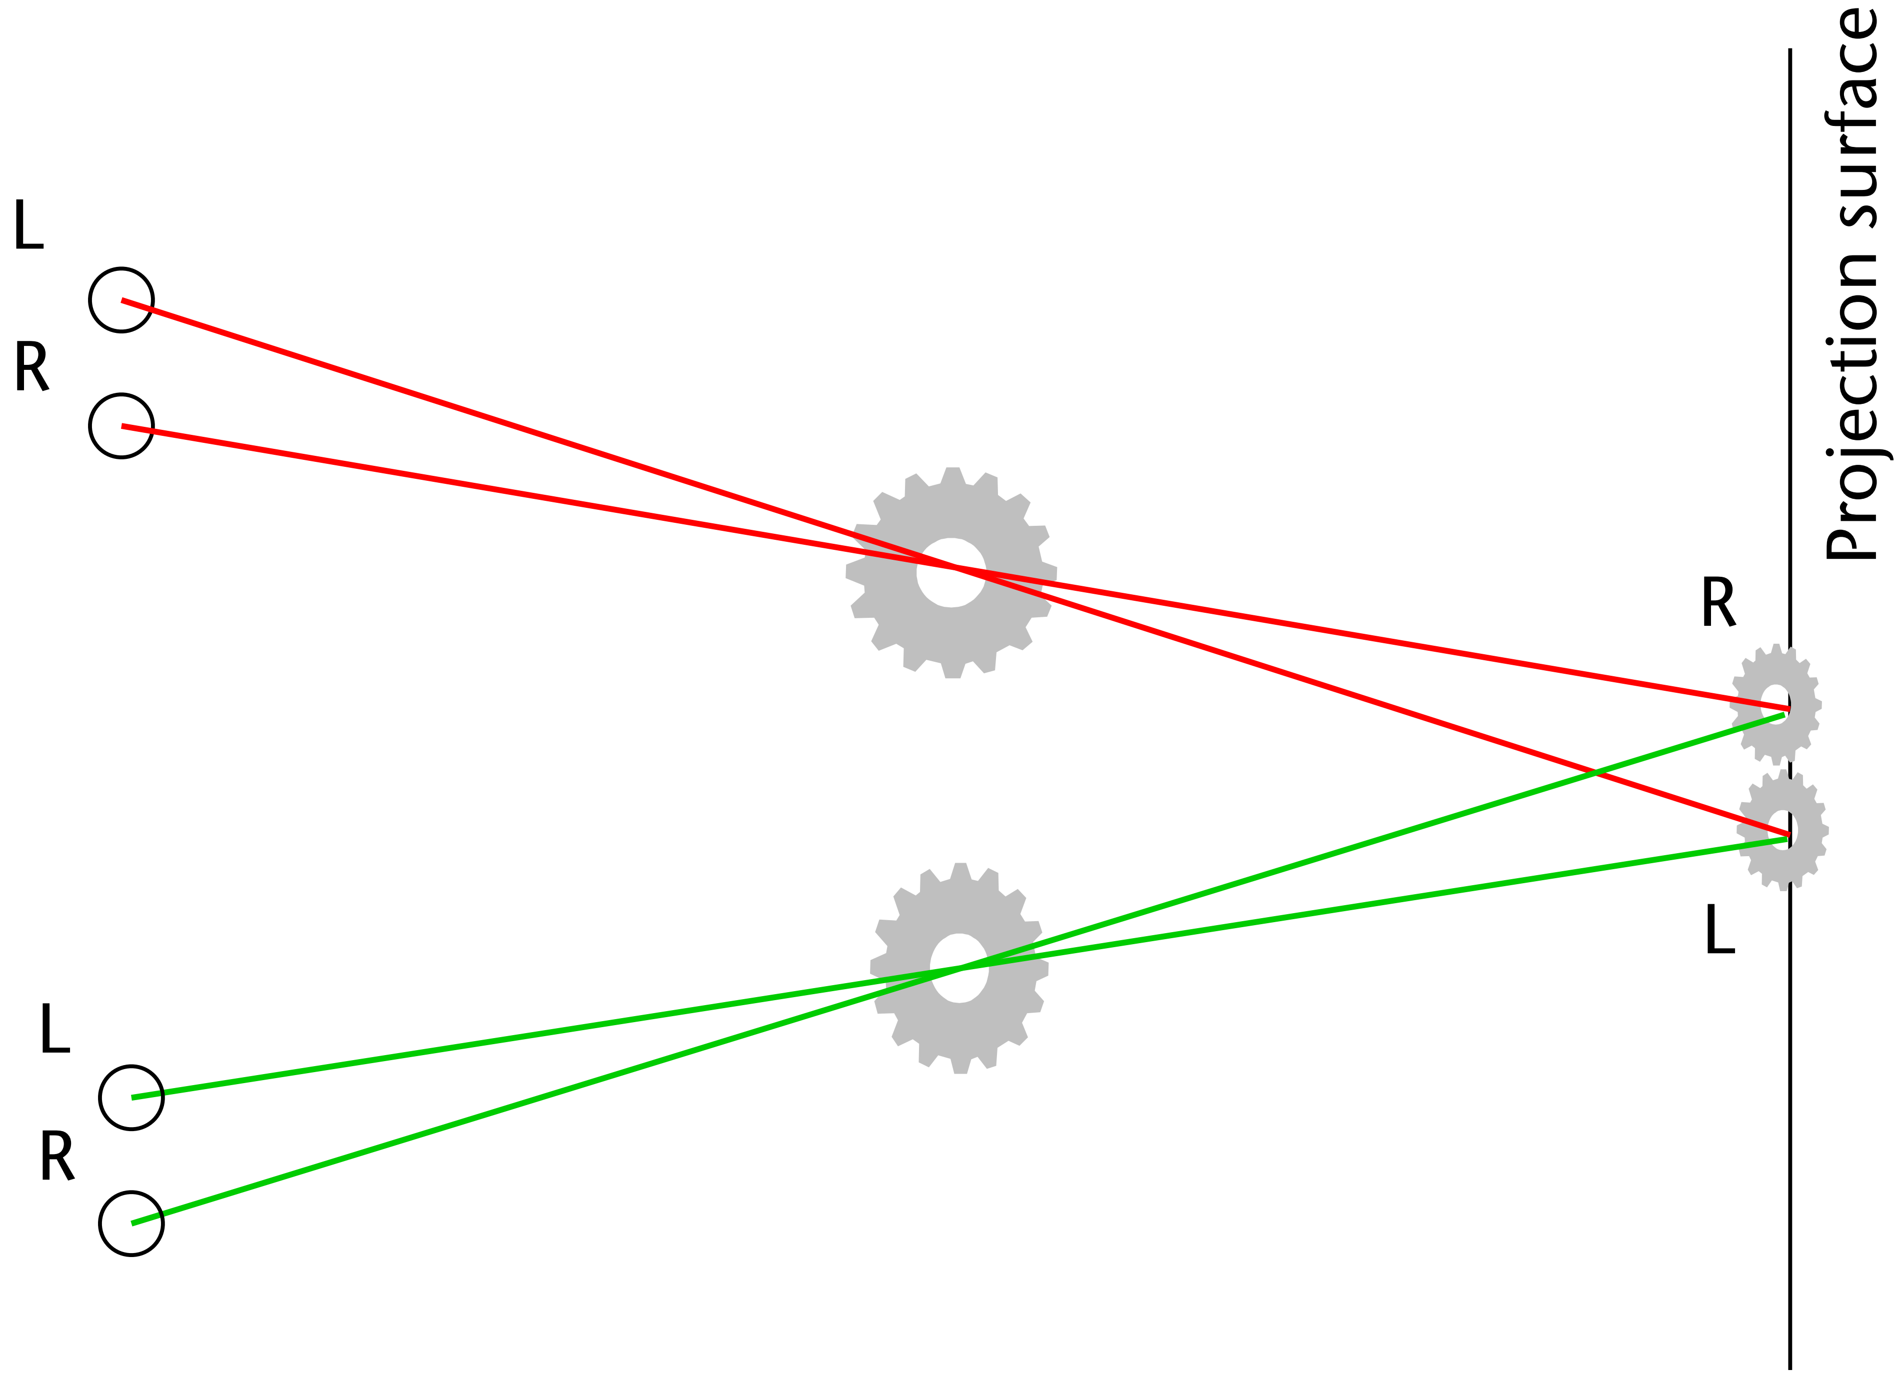

Effect in the Cave

If user's eye matches virtual camera perfectly

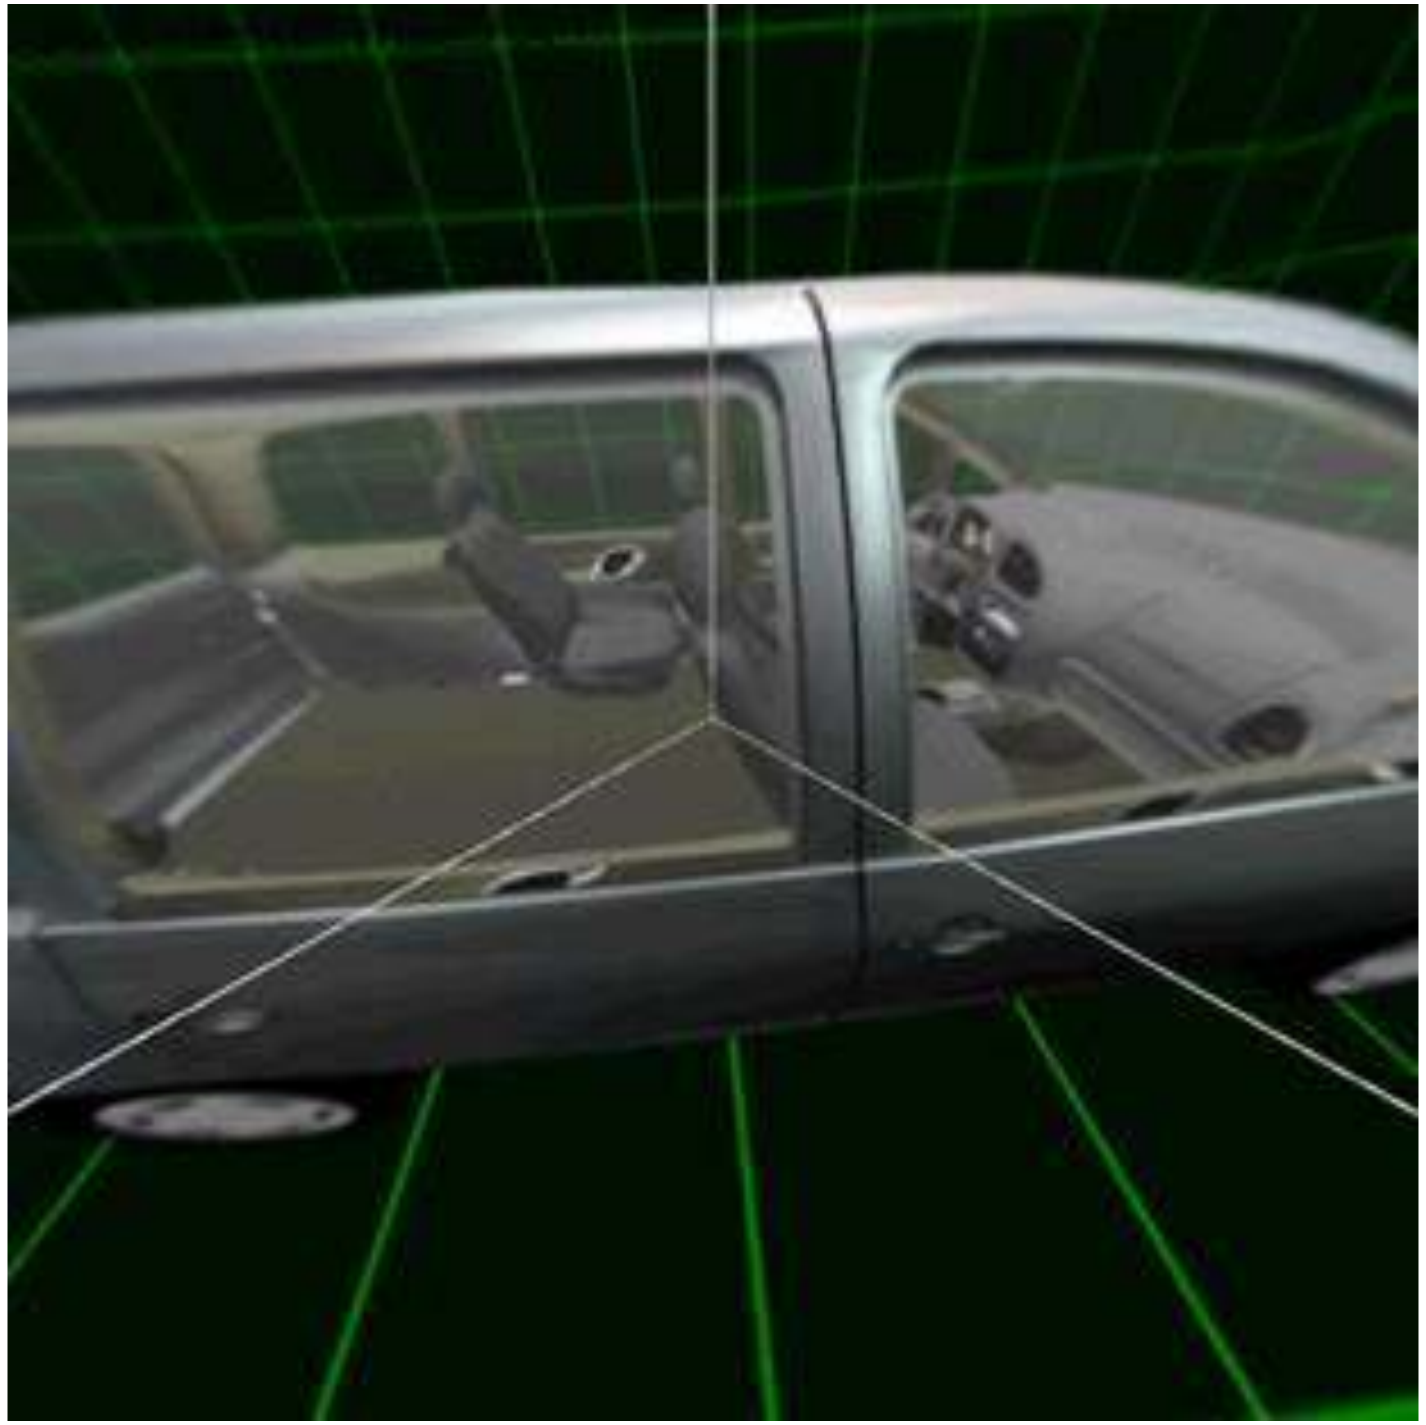

If user's eye is different from virtual camera

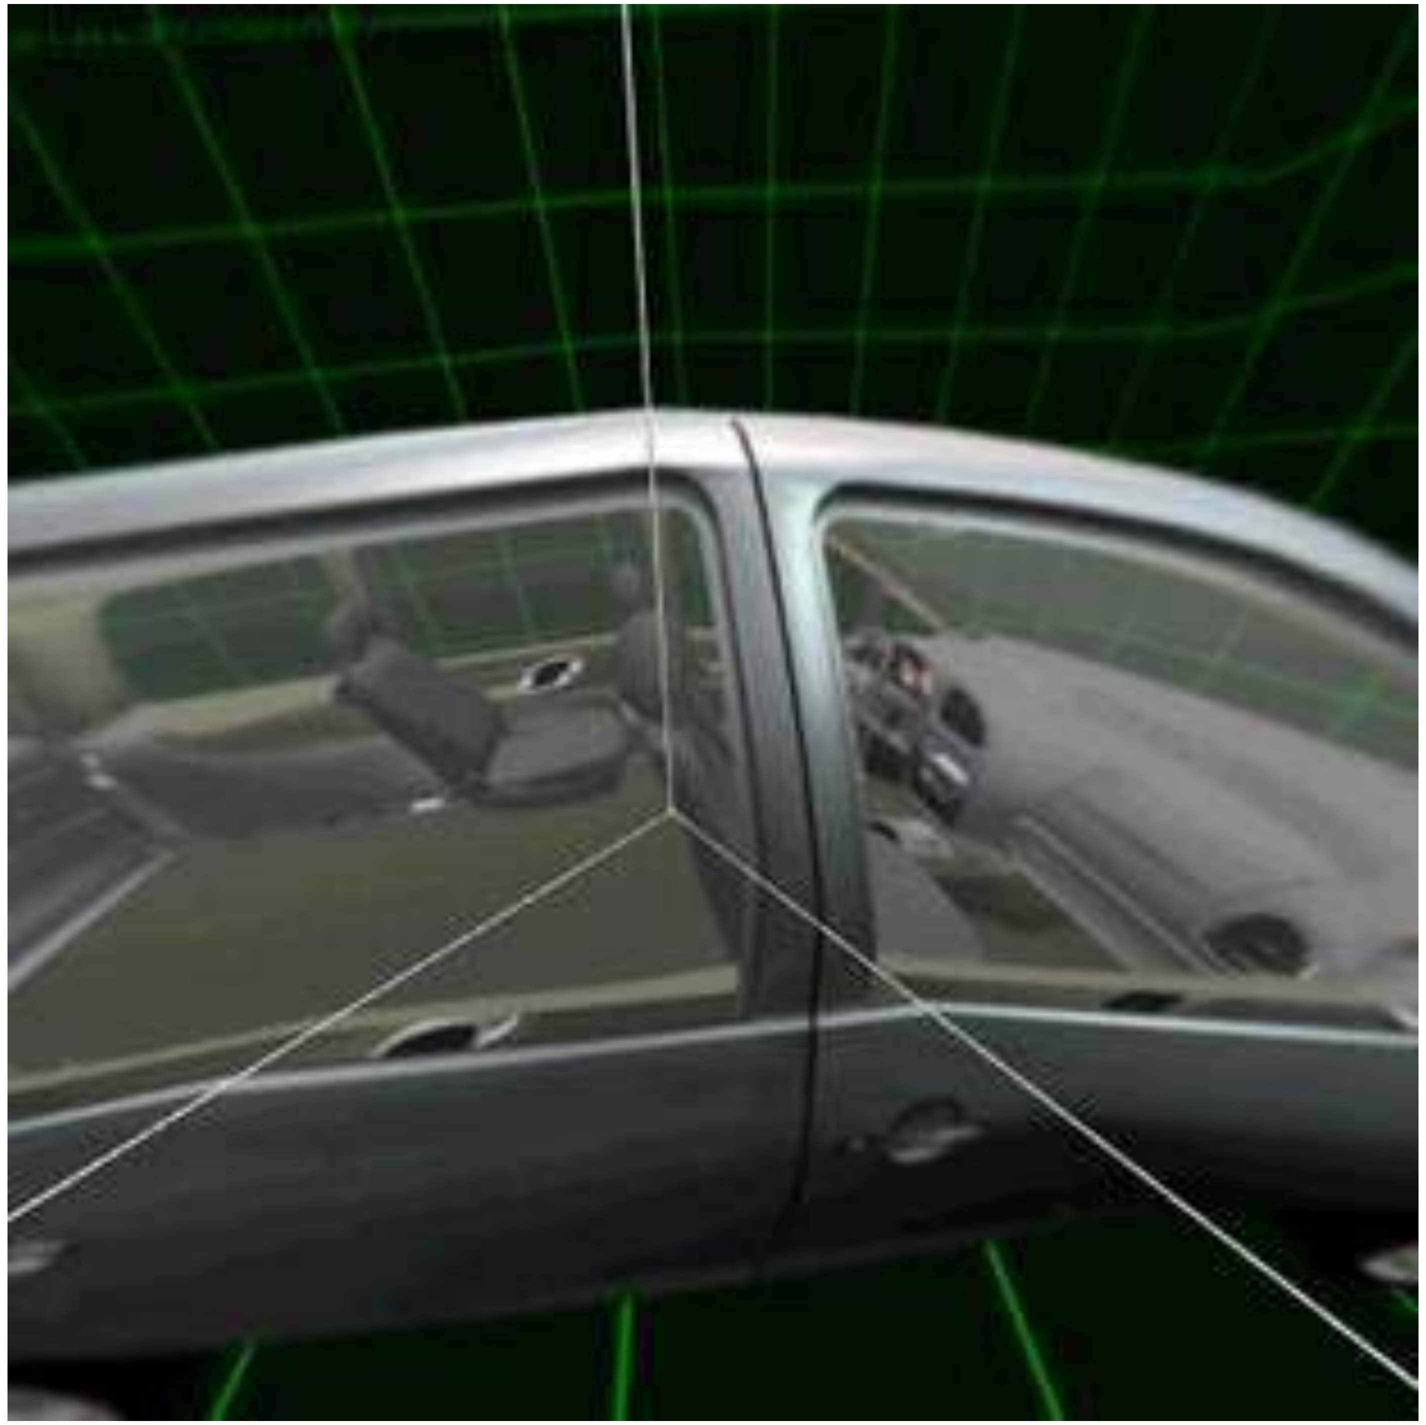

An Optical Illusion Exploiting the Same Effect

Similar Problems Exists in Projection-Based AR

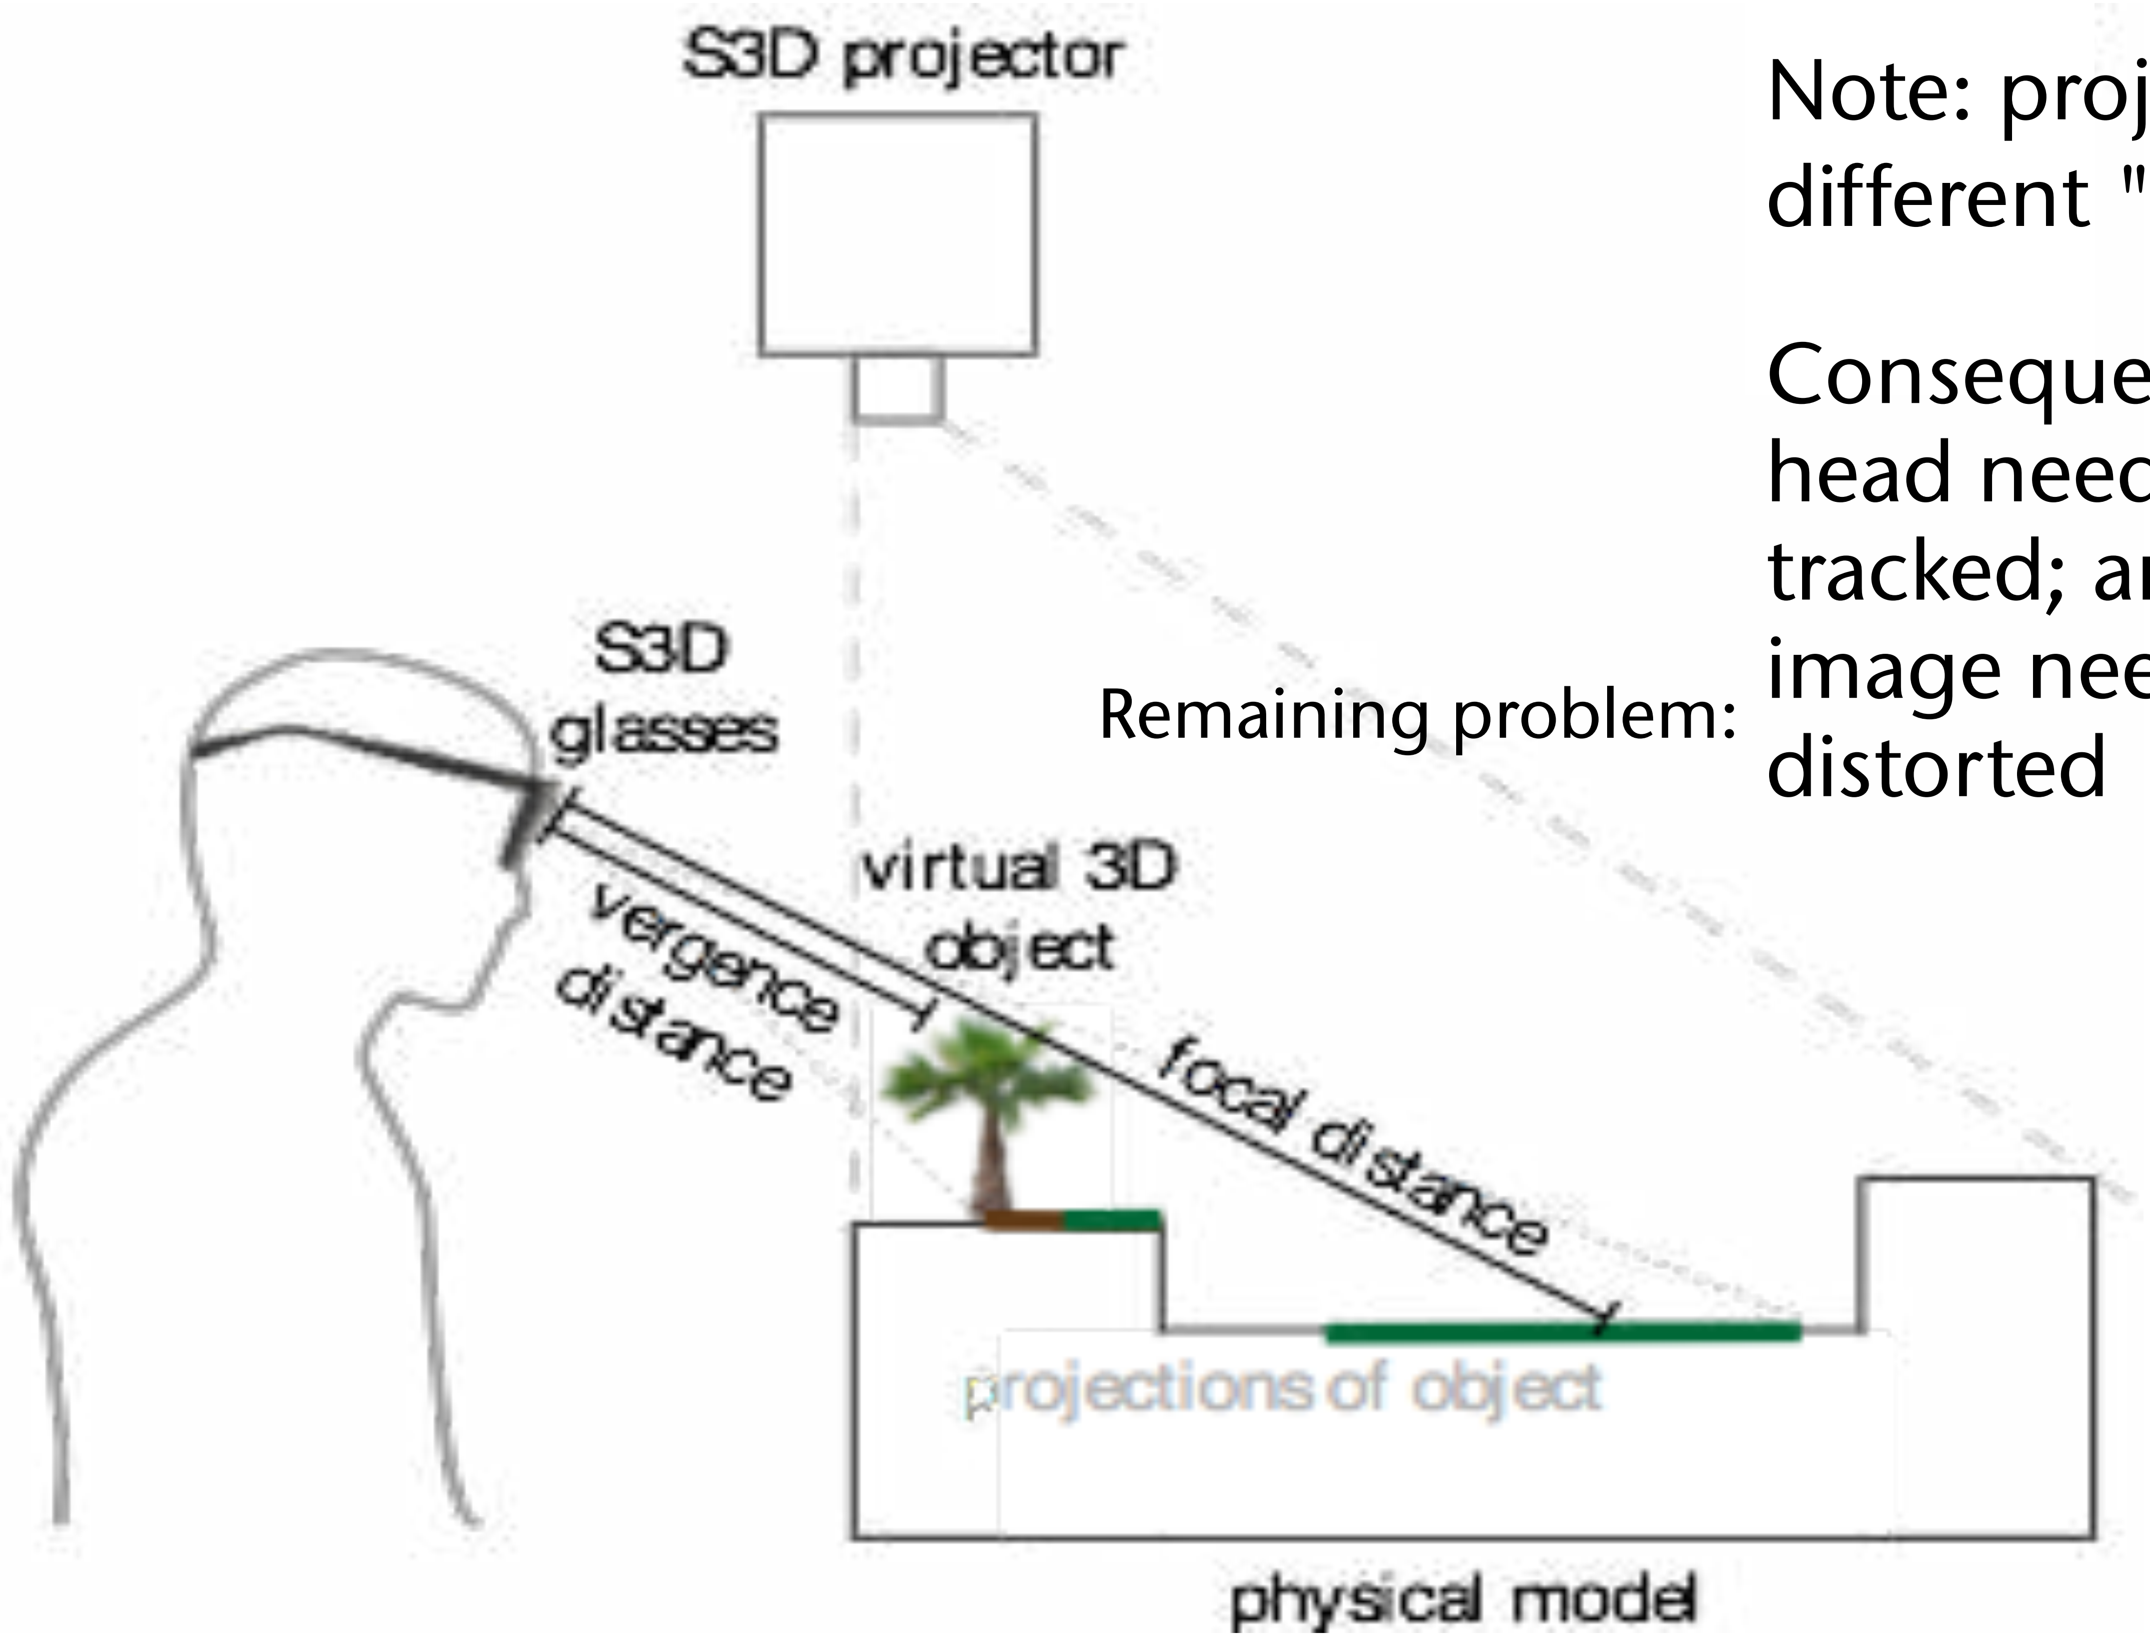

Remaining problem:

Note: projector and user have different "viewpoints"!

Consequence: user's head needs to be tracked; and projected image needs to be pre-distorted

Coherent Virtual Workspaces

- Assume this situation: one stereo display wall, several users in front of it
- Problem with single-user projection (stereo or mono) and multiple users: only the viewpoint of *tracked* (primary) user is correct, only she will see correct images!
- One of the problems: communication using pointing fails

Correct for primary user *only*

Correct for *both* users

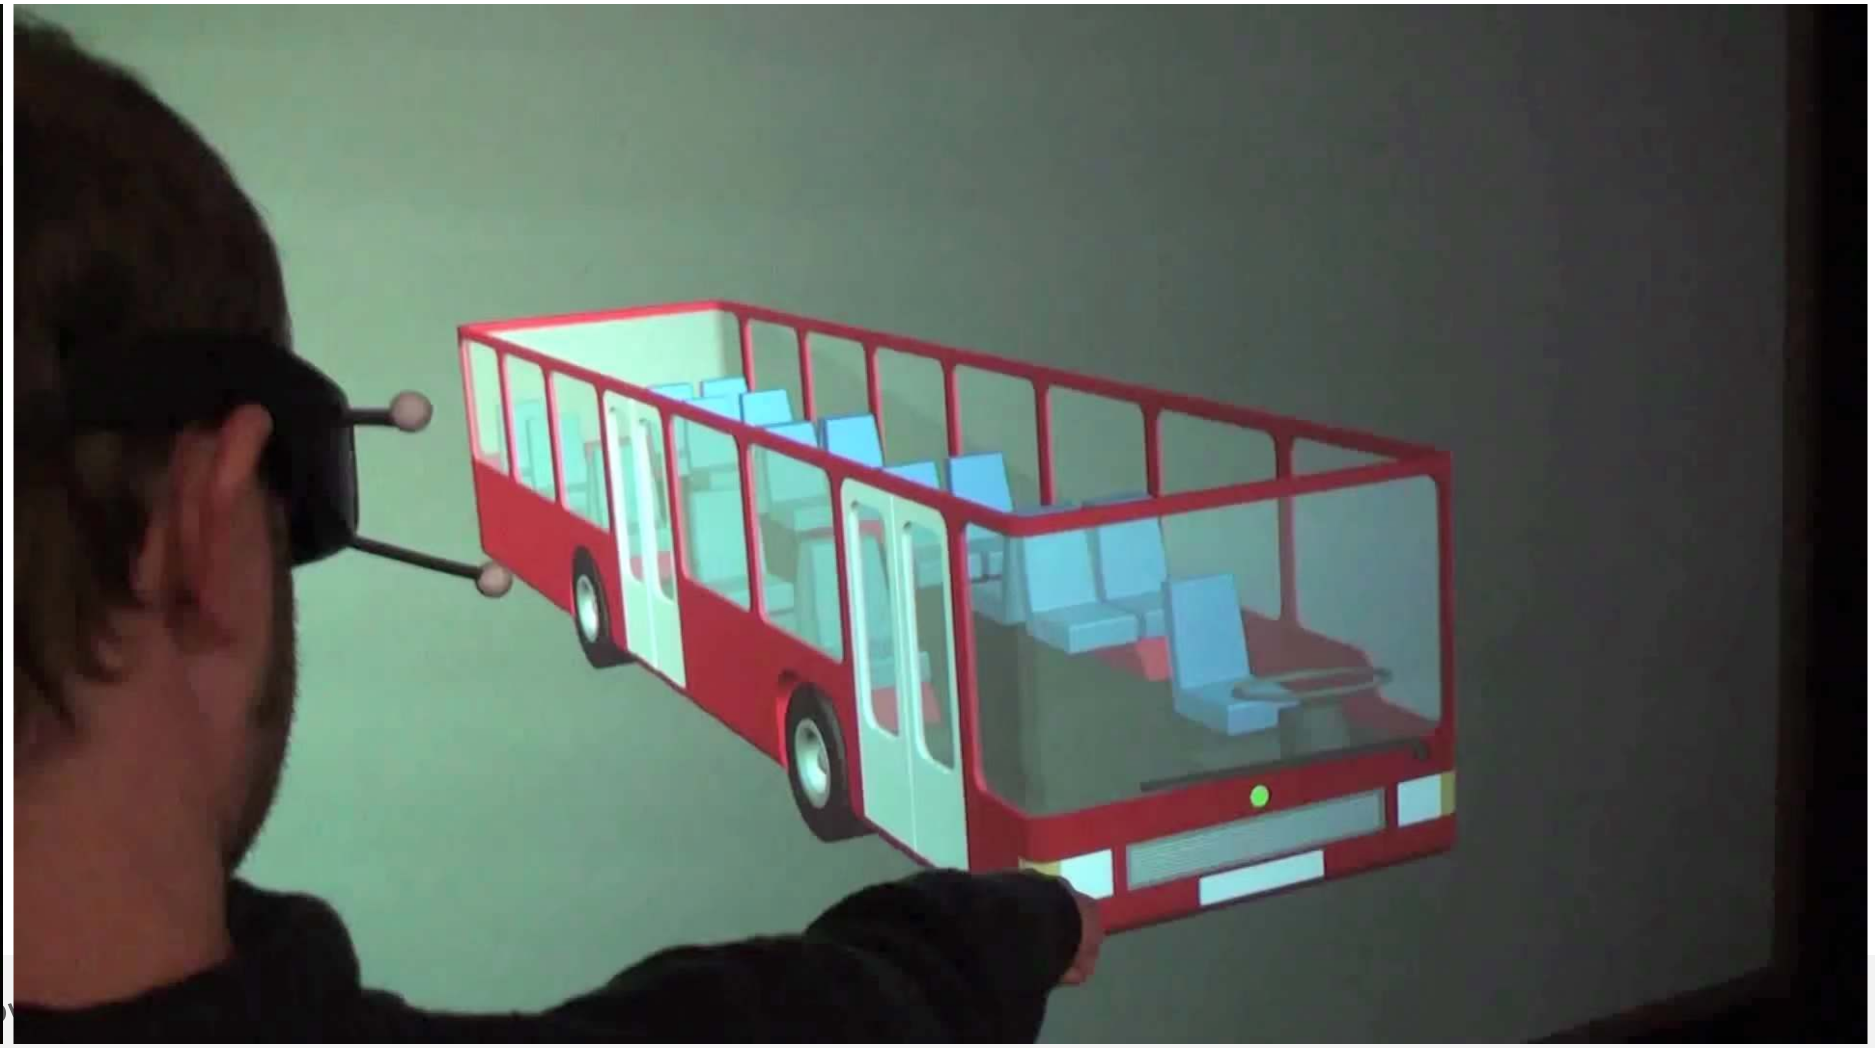

Benefit of Correct Projection for All Users

- With *perspectively correct* projections for *all co-located* users, the shared 3D space will become coherent for all users

- Benefit: direct communication (including *pointing!*) in co-located collaborative virtual environment (CVE) becomes possible
  - Note: 80% of all human communication is non-verbal (!)

Solution: Correct (Stereo) Projection for Multiple Users

- Probably only possible for a small number of users
- *Temporally multiplexed* (shutter glasses):
  - Framerate for multi-user stereo = framerate for mono × 2 × #users
  - Light intensity reaching each eye gets is extremely low
- Infitec for several users:
  - Each user gets glasses with slightly shifted comb filters
  - With *n* users we need 2*n* different comb filters → extremely narrow bands, 2*n* projectors needed
  - Same problem with light intensity
- *Spatially multiplexed*
- Combination of the above

Example Hardware Setup

FYI

[Fröhlich, 2011]

- Combination of active and passive stereo, plus ingenious utilization of field-sequential projectors
- Recap from CG1:
  field-sequential RGB with DLPs

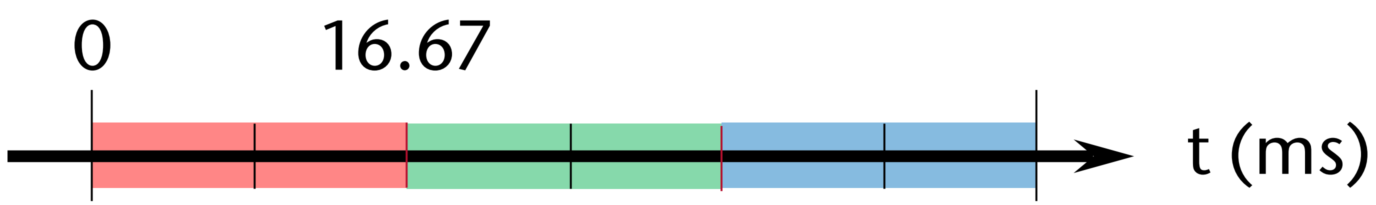

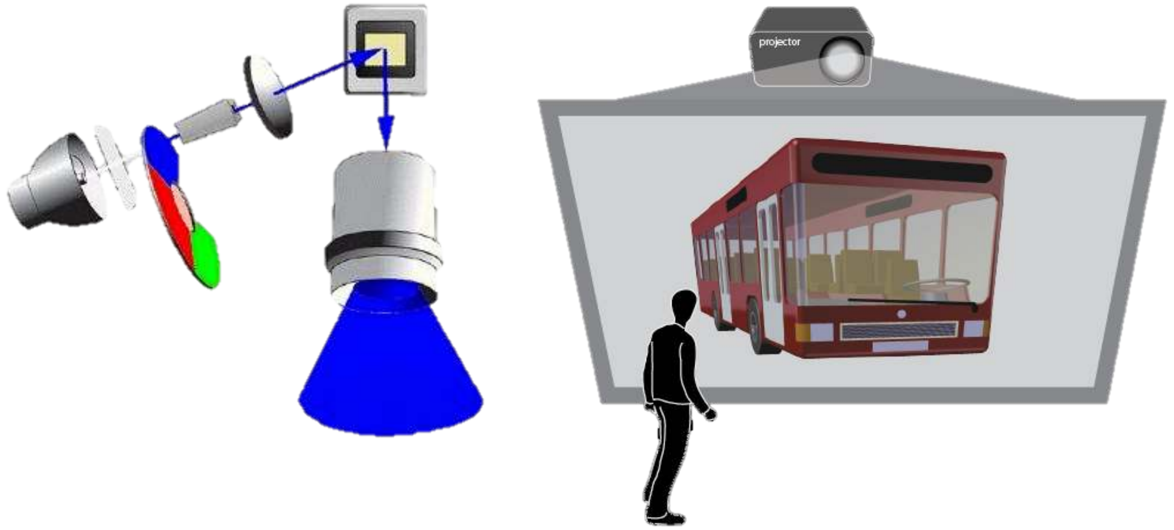

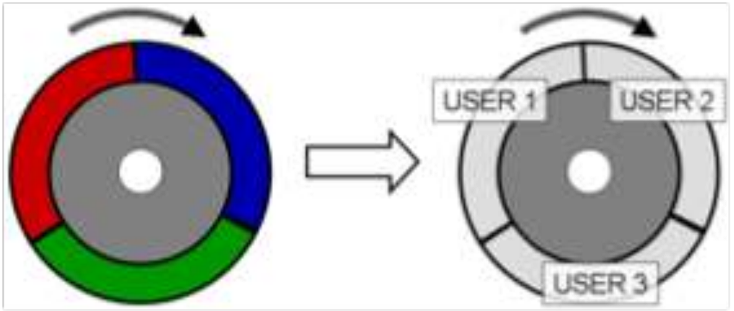

1. Modification: remove color wheel
2. Modification: each user gets shutter glasses that additionally has left/right polarization filters
   - Must be fast enough to prevent cross-talk!

The Hardware in Principle

FYI

- 6 stereo video streams are generated by 6 graphics cards in 1 PC
- Distribution of the video streams to 6 projectors via multiplexers

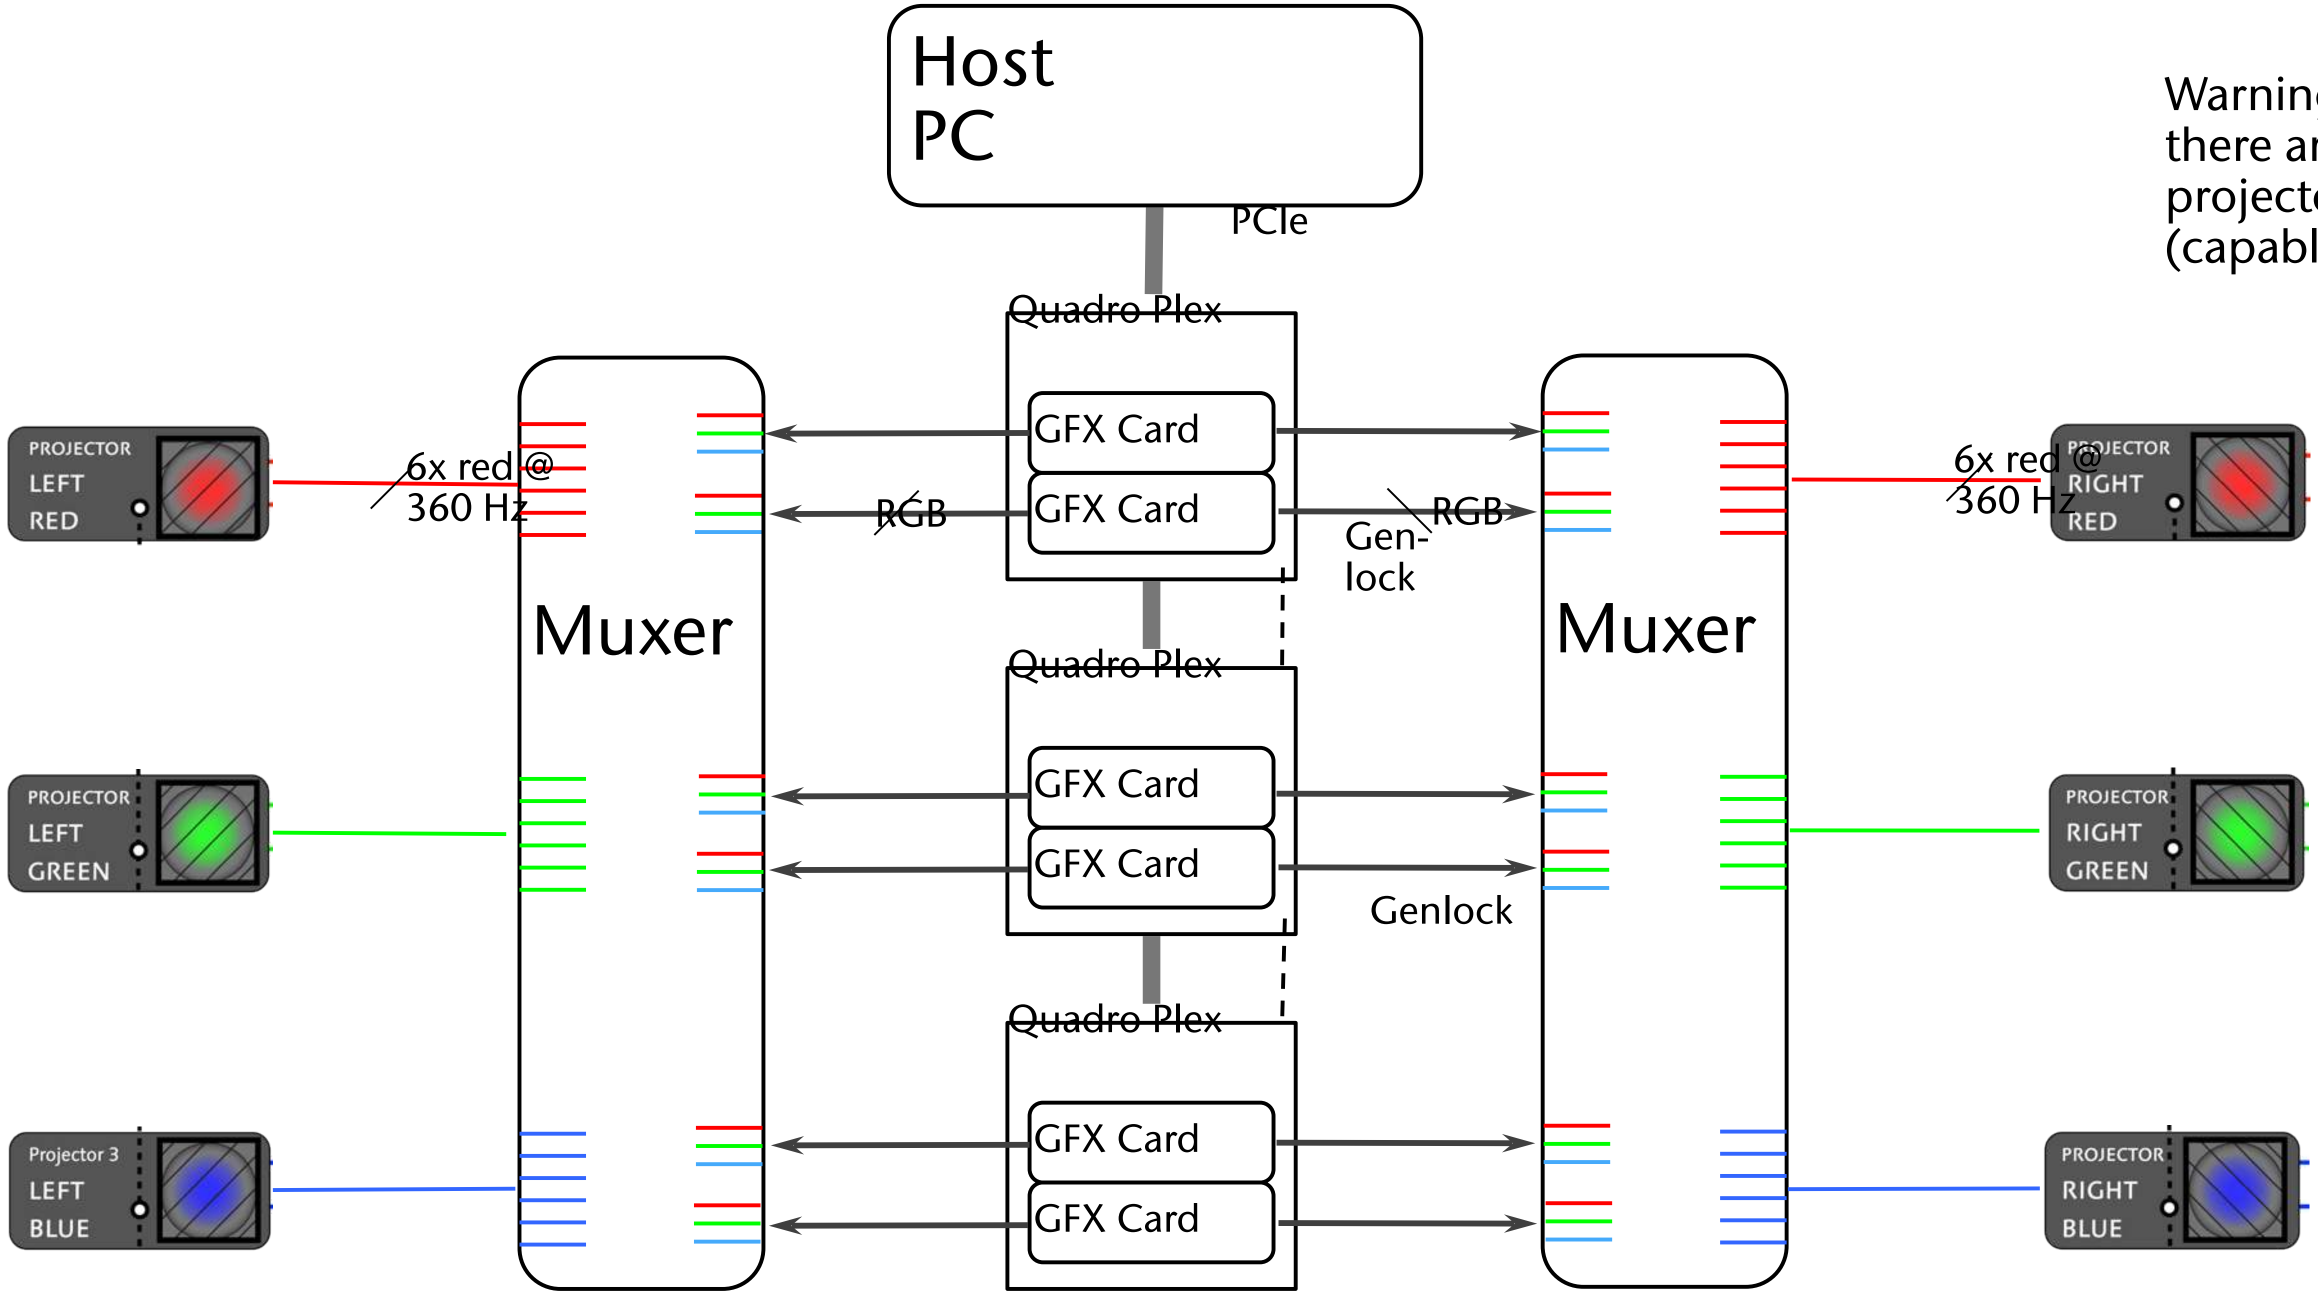

Warning: in the real world, there are no such projectors available! (capable of 360 Hz inputs)

- Timing:

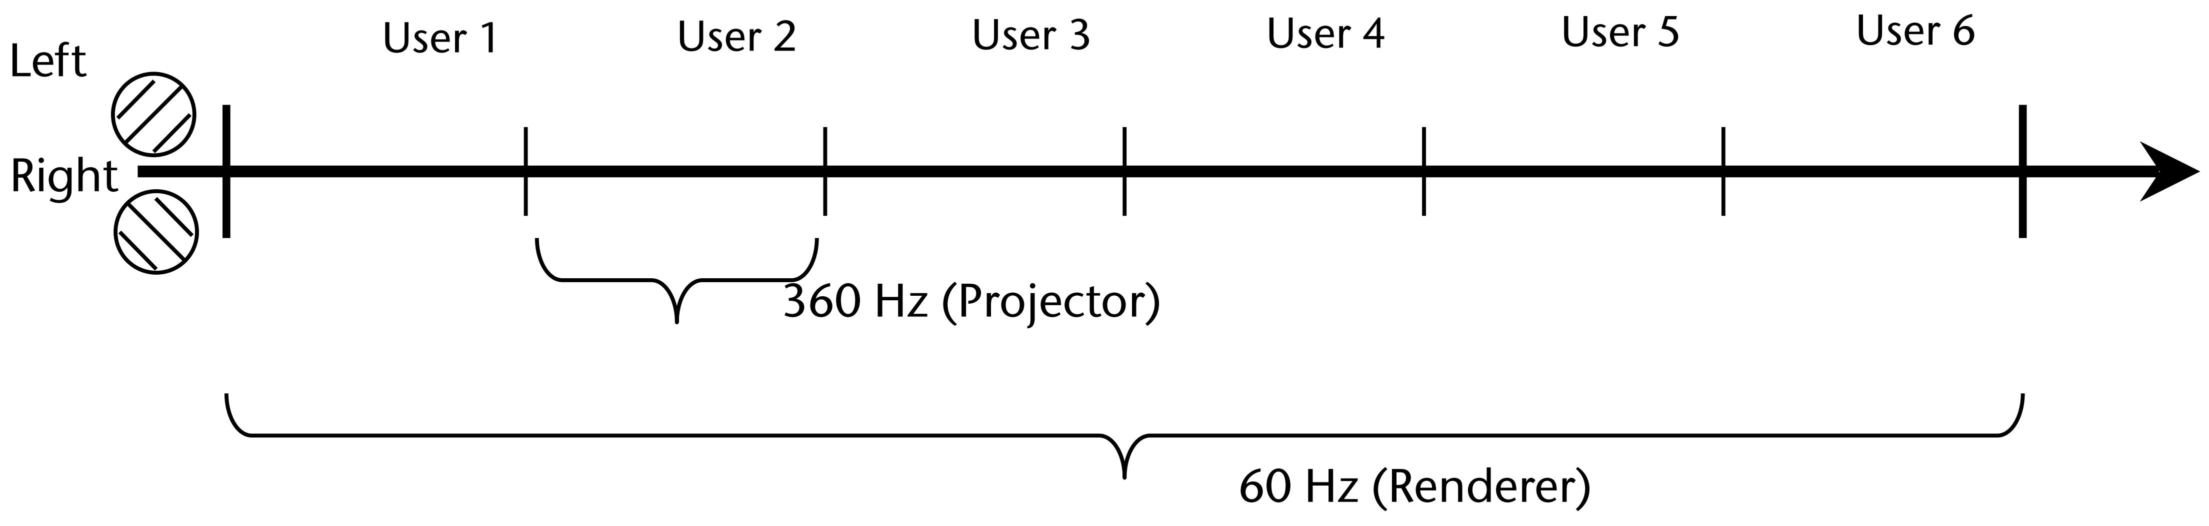
Left
Right
User 1
User 2
User 3
User 4
User 5
User 6
360 Hz (Projector)
60 Hz (Renderer)

Demo Application

Uni Weimar,
Prof. Bernd Fröhlich

Workspace Awareness in VR-CSCW (CVE's)

- Workspace Awareness = "up-to-the-moment understanding of the other person's interaction with the shared workspace" [Gutwin & Greenberg, 2002]
- Factors / questions:
  - Who is participating / interacting? (People)
  - What are they doing ? What will they be doing next? (Actions / Intentions)
  - What can they see? Where can they have effects? (Perception / Influence)

An Interaction Issue with Multi-User VR-CSCW

- Navigation: the "navigator" controls the path for all users (and he sees only his own viewpoint!)
- Problem: the other users' viewpoint goes through walls
- Solutions:
  - Adjust the paths of the other users automatically to bring them closer to the navigator's viewpoint
  - Fade away obstacles in the path of each user

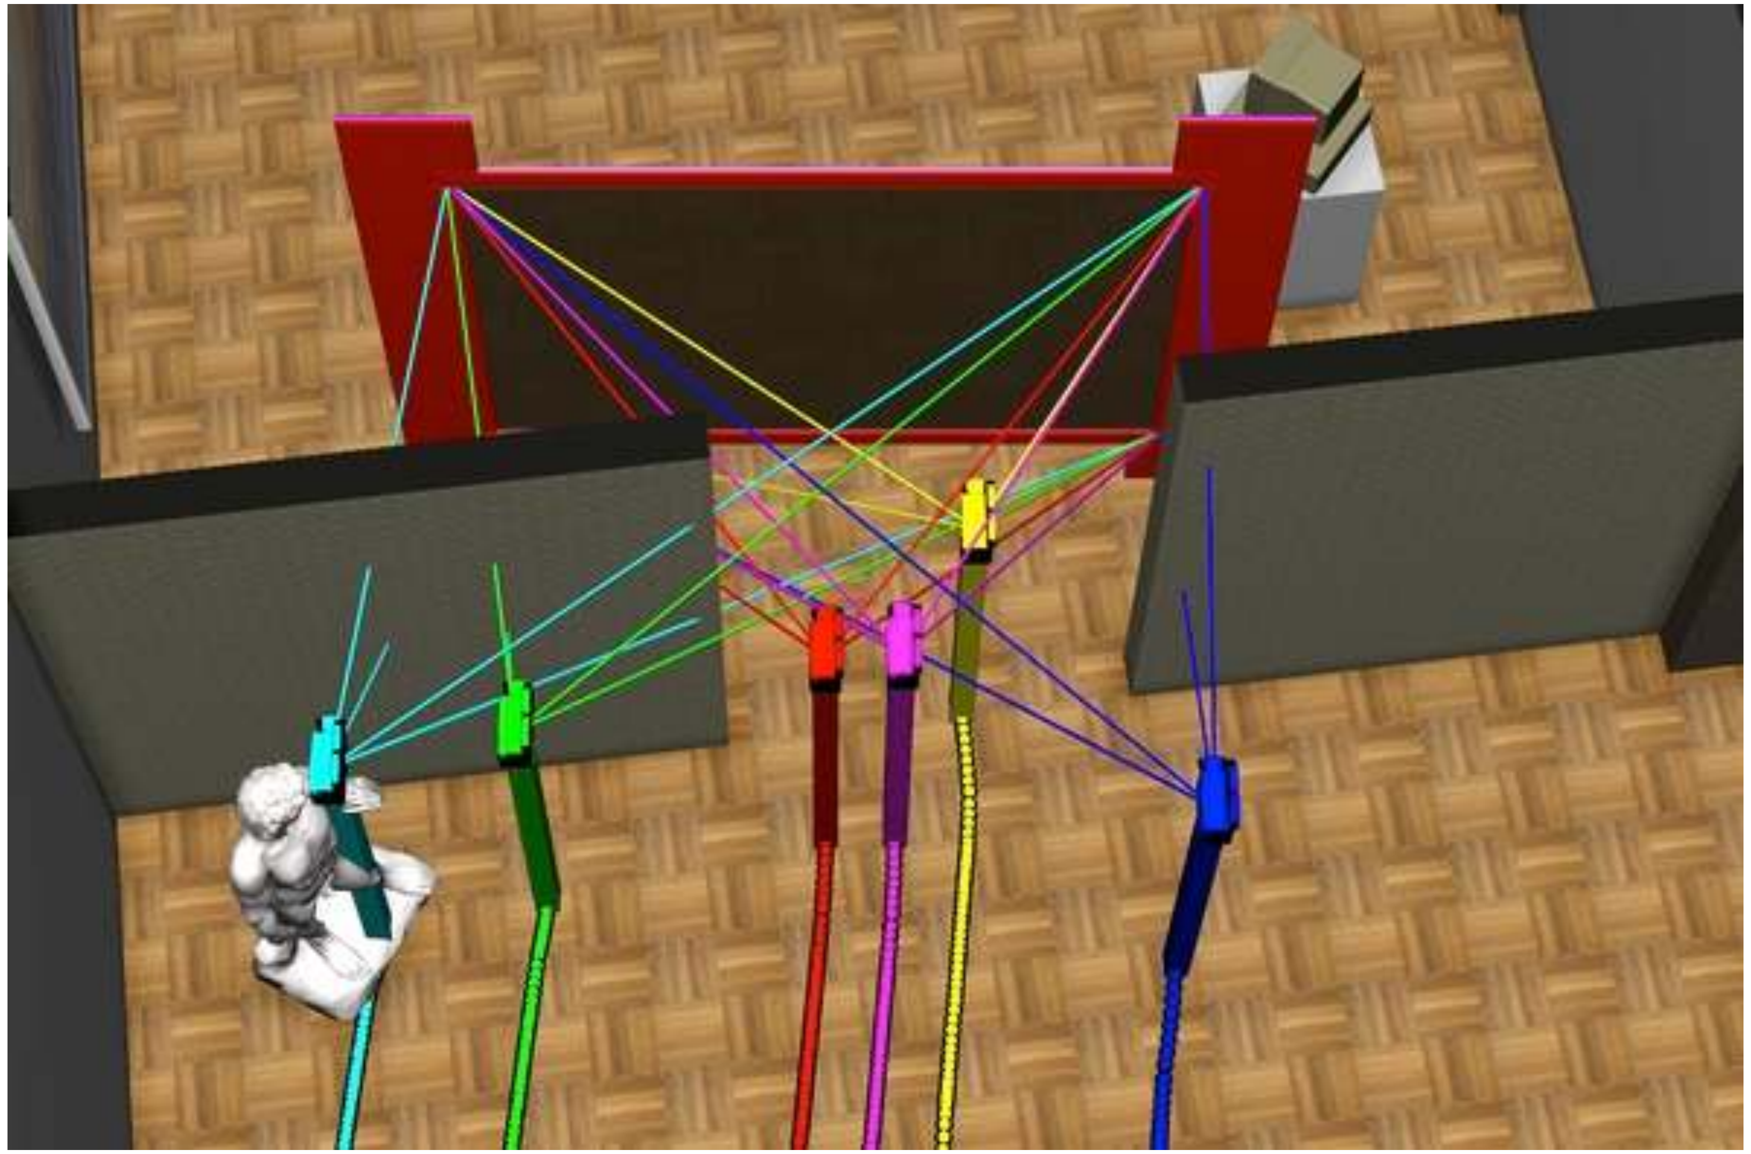

Demo

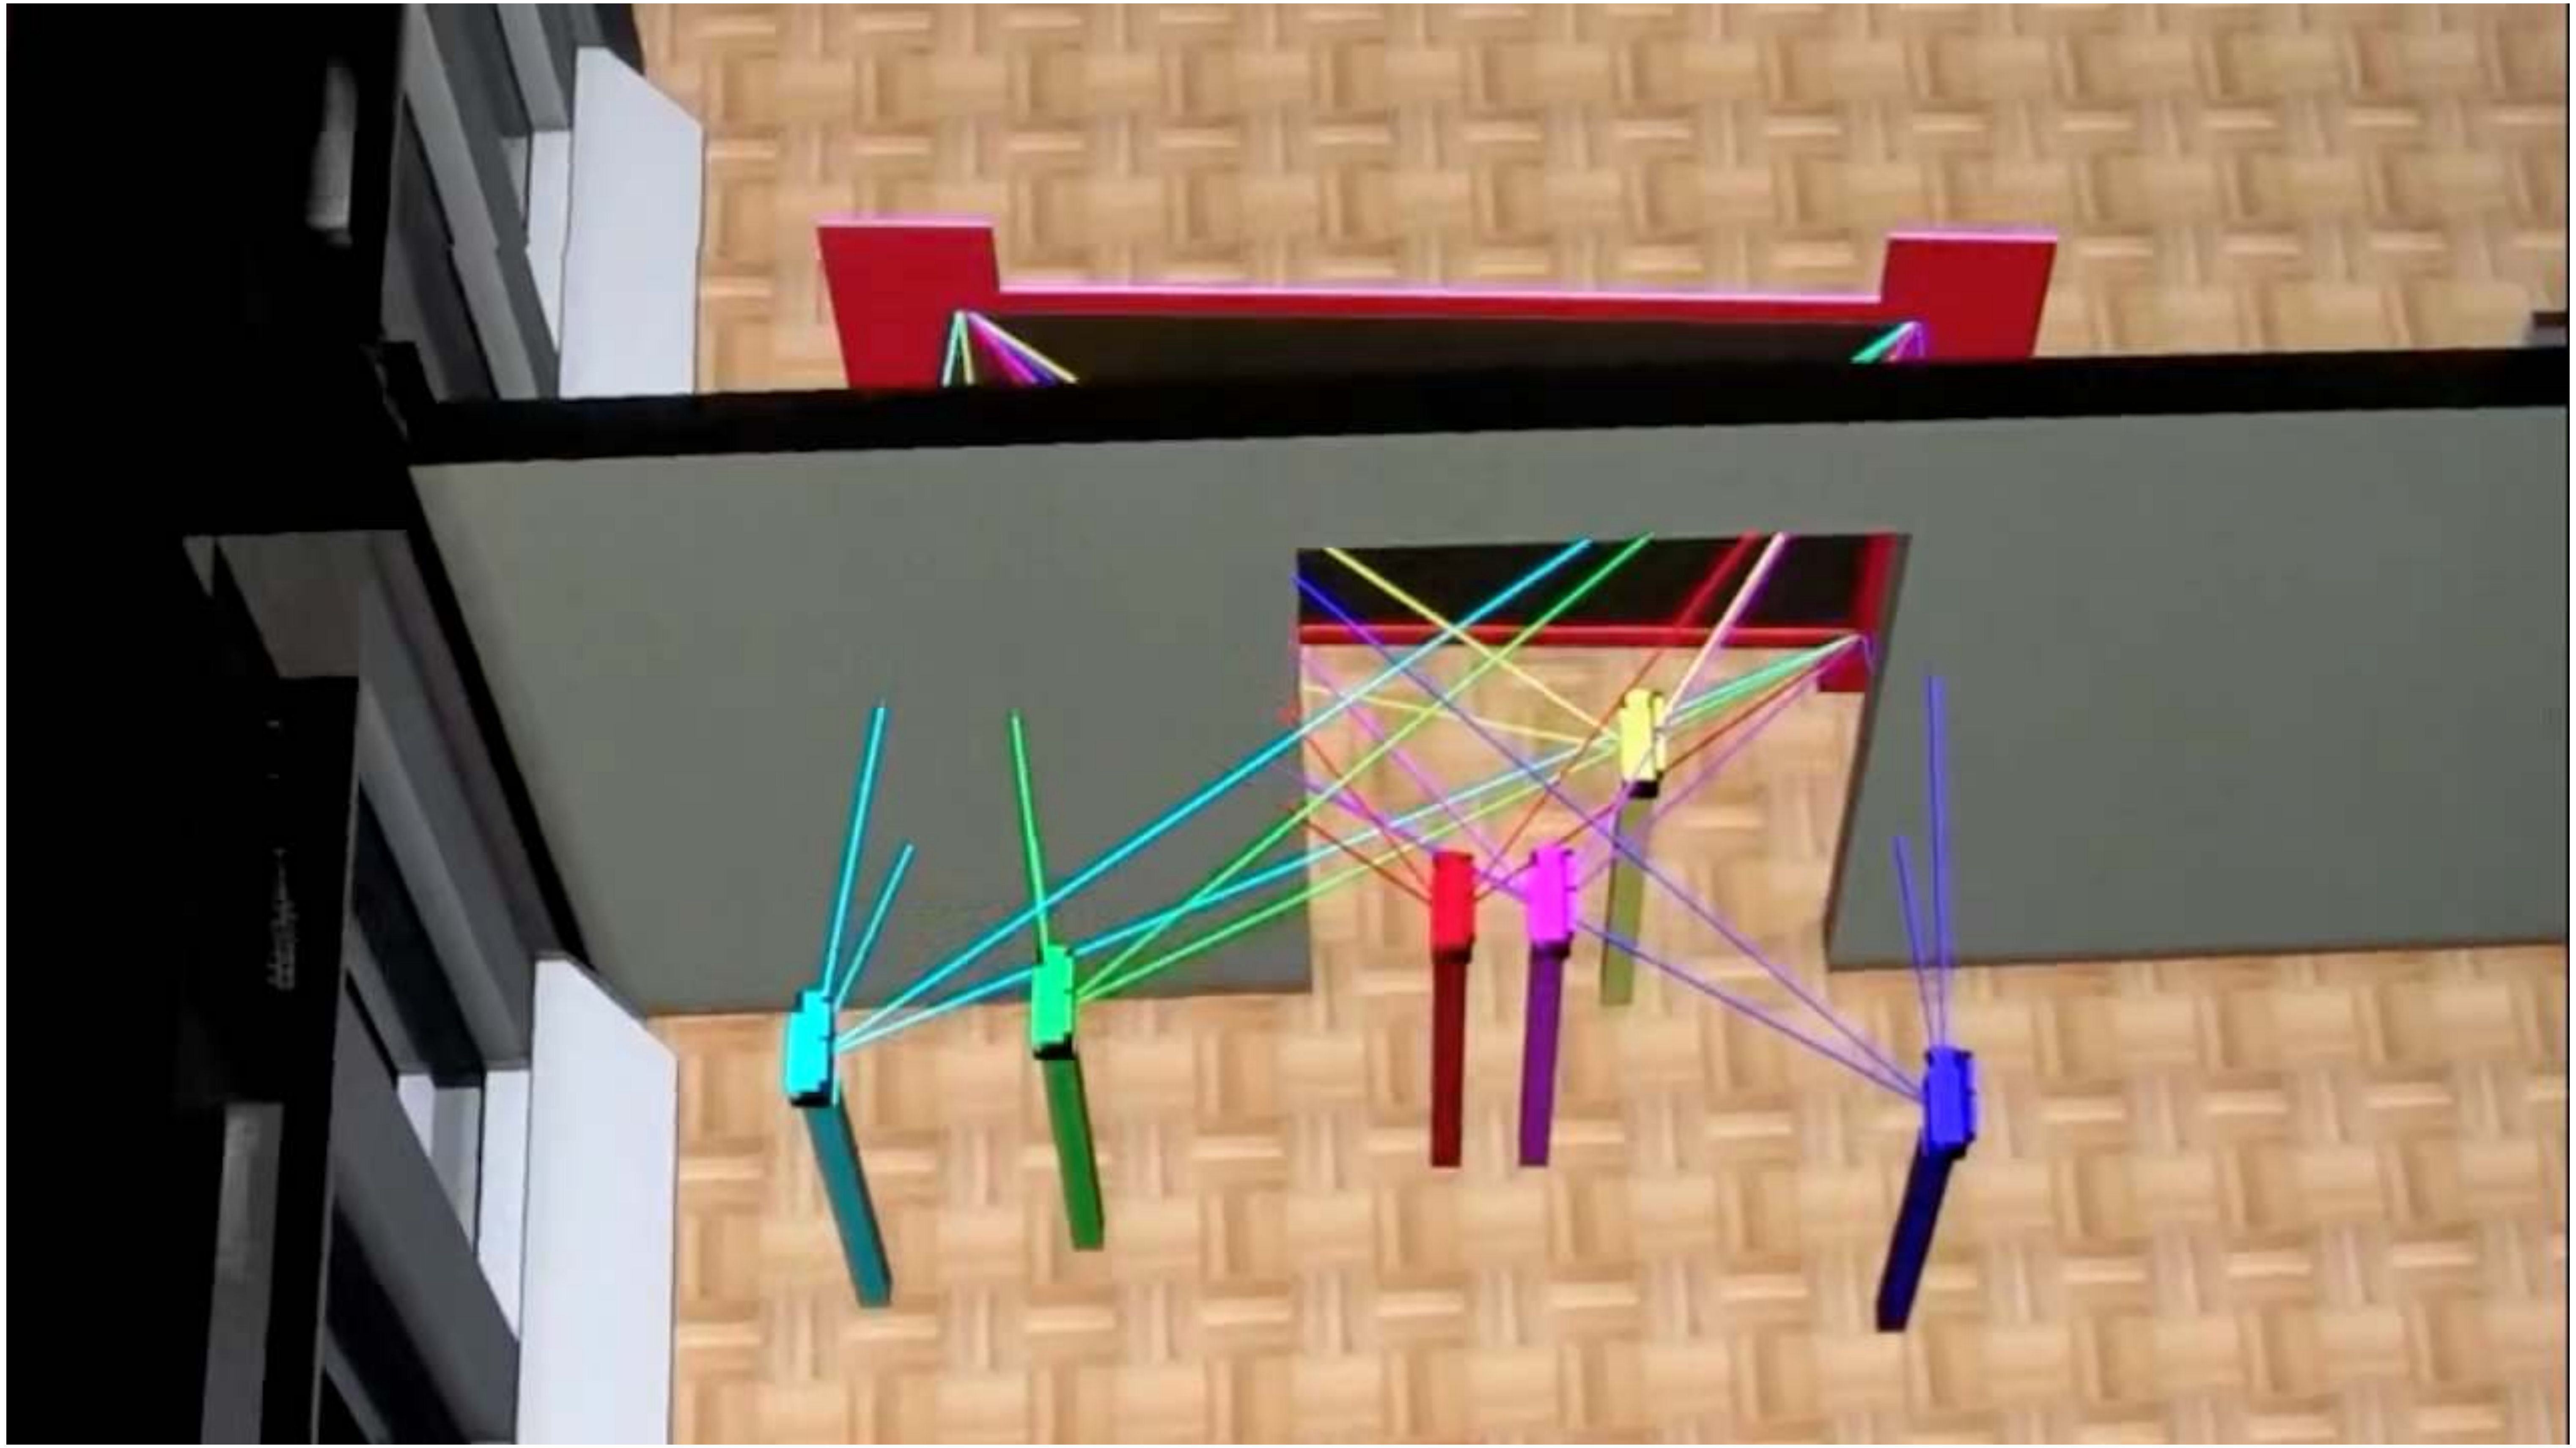

Rendering on HMDs

- Optics in HMDs usually cause some amount of distortion
  - Especially the Oculus Rift
- Idea: pre-distortion (using multi-pass and texturing or shaders)

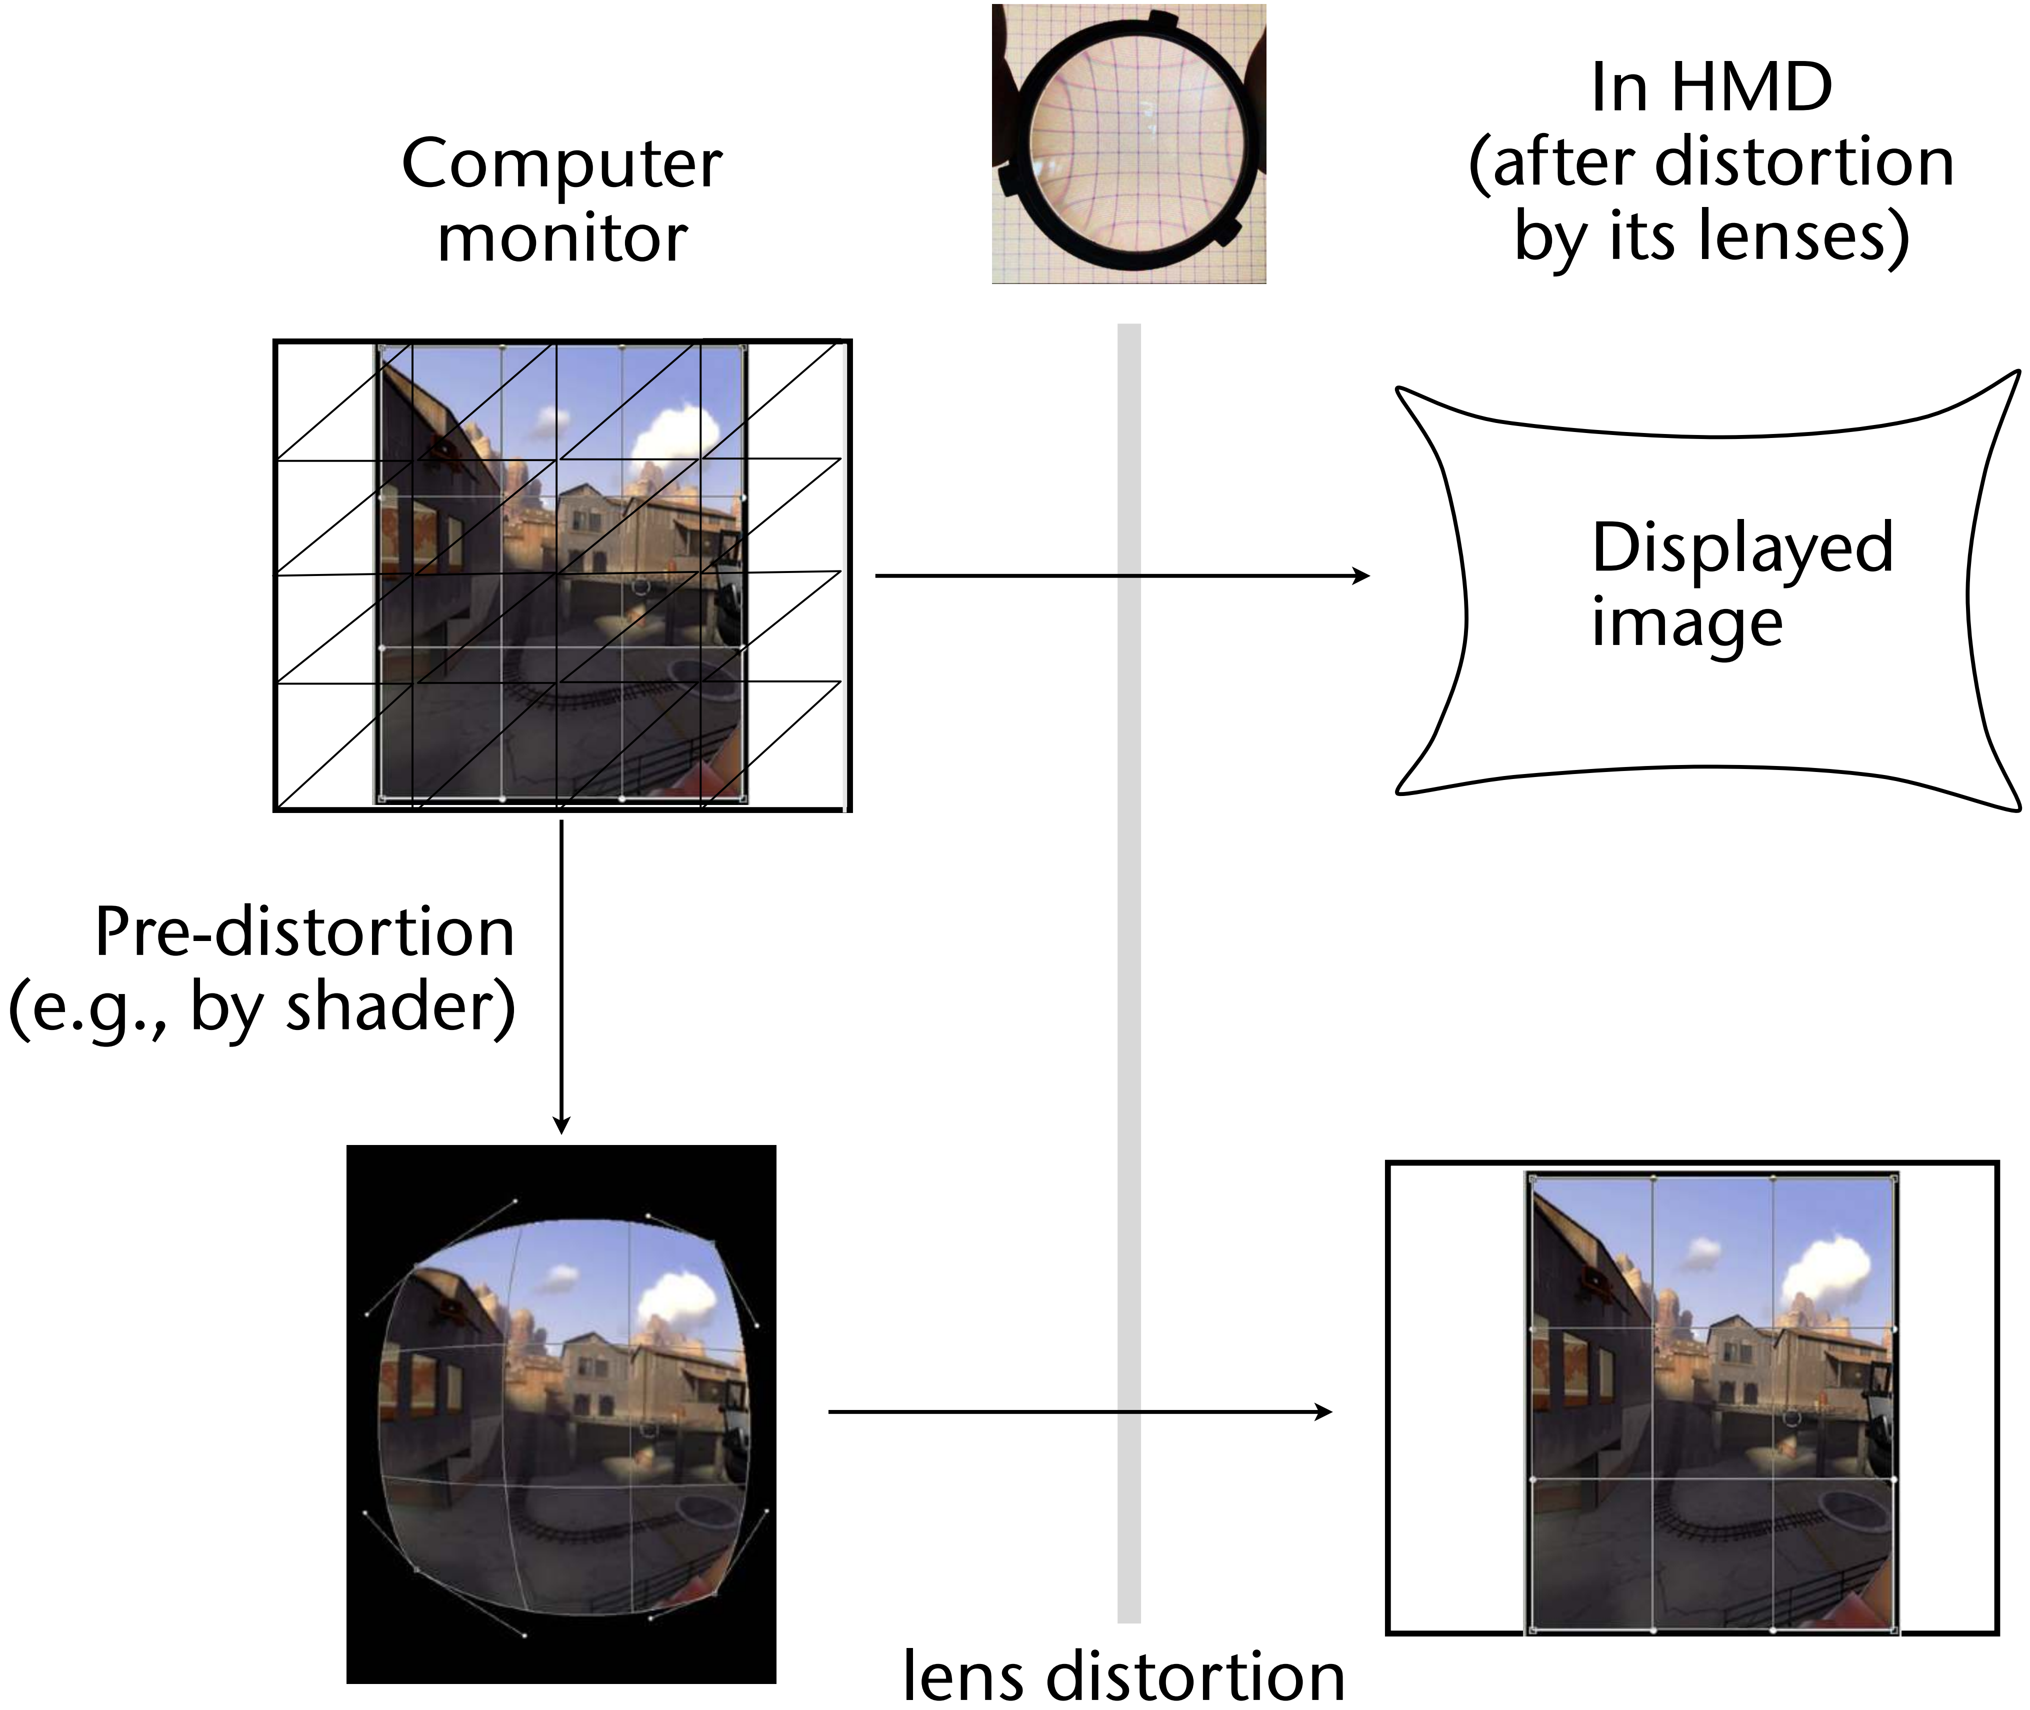

One of the Hard Requirements for VR / AR

- Images must appear fixed in space, no matter how fast a user moves

- This is hard even for a still environment!
  - Reason 1: latency (later)
  - Reason 2: display persistence (in the following)

Recap: the Graphics Backend Hardware

- Current displays are always raster displays:

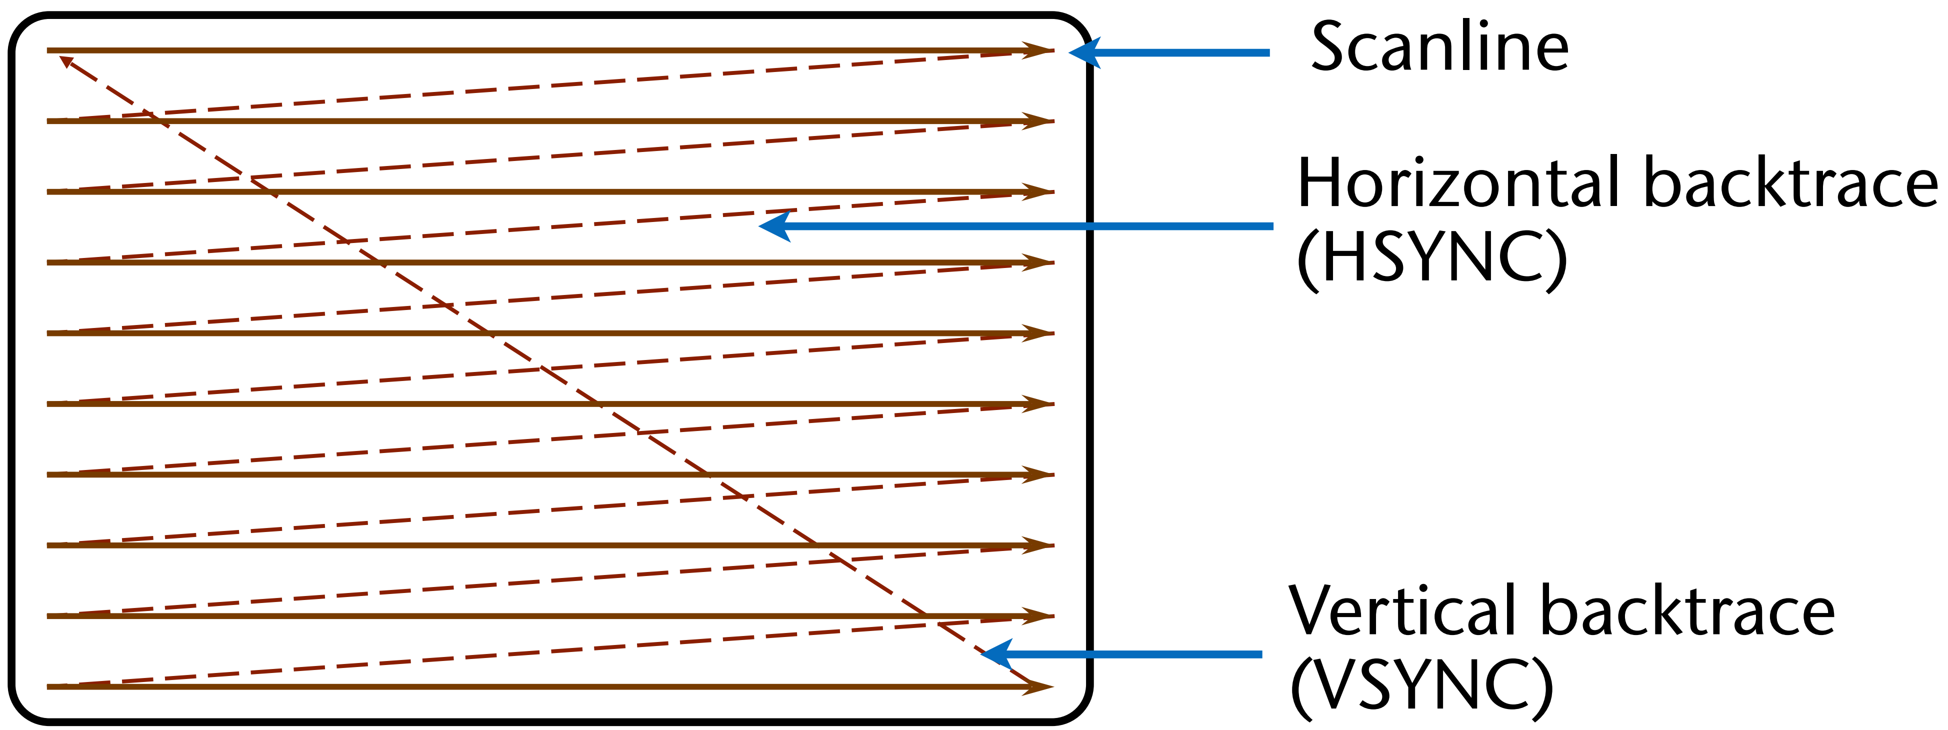

- Double buffering to prevent flickering:
  - Best point to swap the buffers?

  → VSYNC

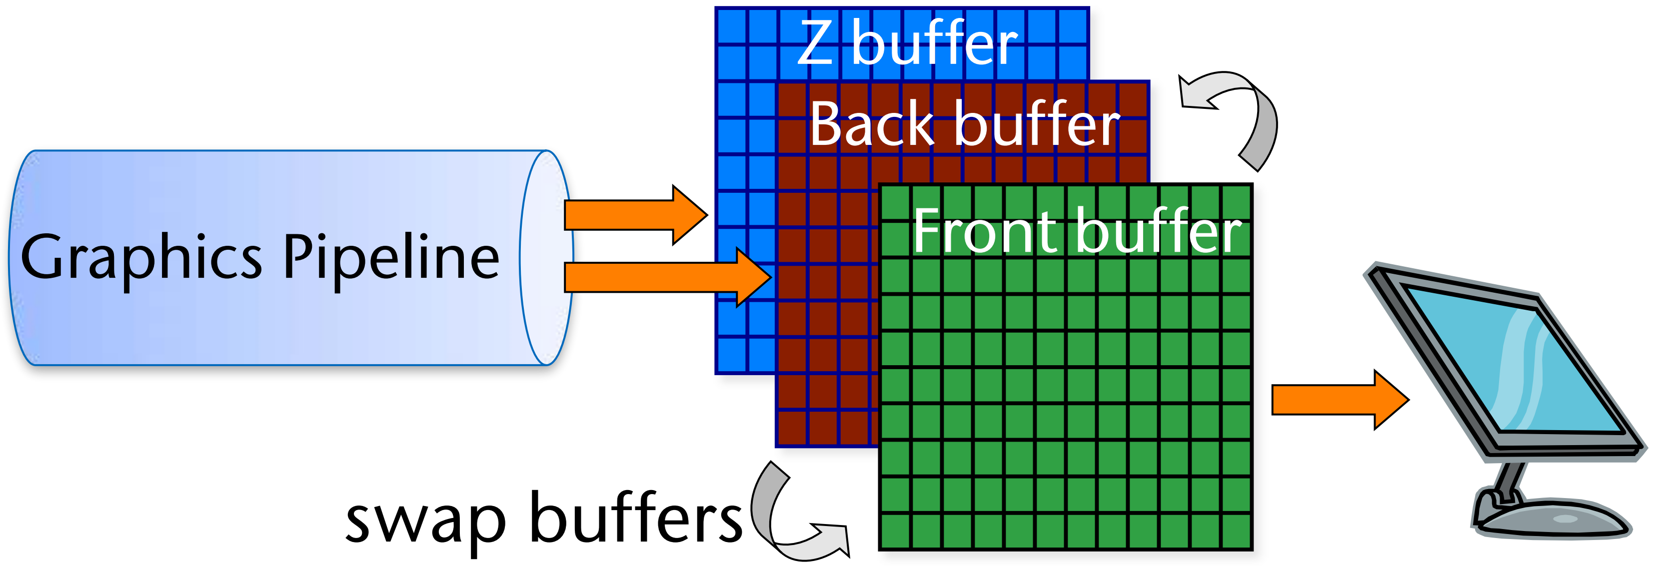

Low Persistence vs. Full Persistence

- Definition: persistence (in displays) = length of time that a pixel on a display remains emitting light after it has been switched on / energized
- Persistence in
  - CRT's: phosphor gets energized by electron beam, illuminates, then decays → low persistence / short persistence
  - LCD's & LED's: pixel can be illuminated virtually infinitely (need to be turned off) → full persistence / long persistence

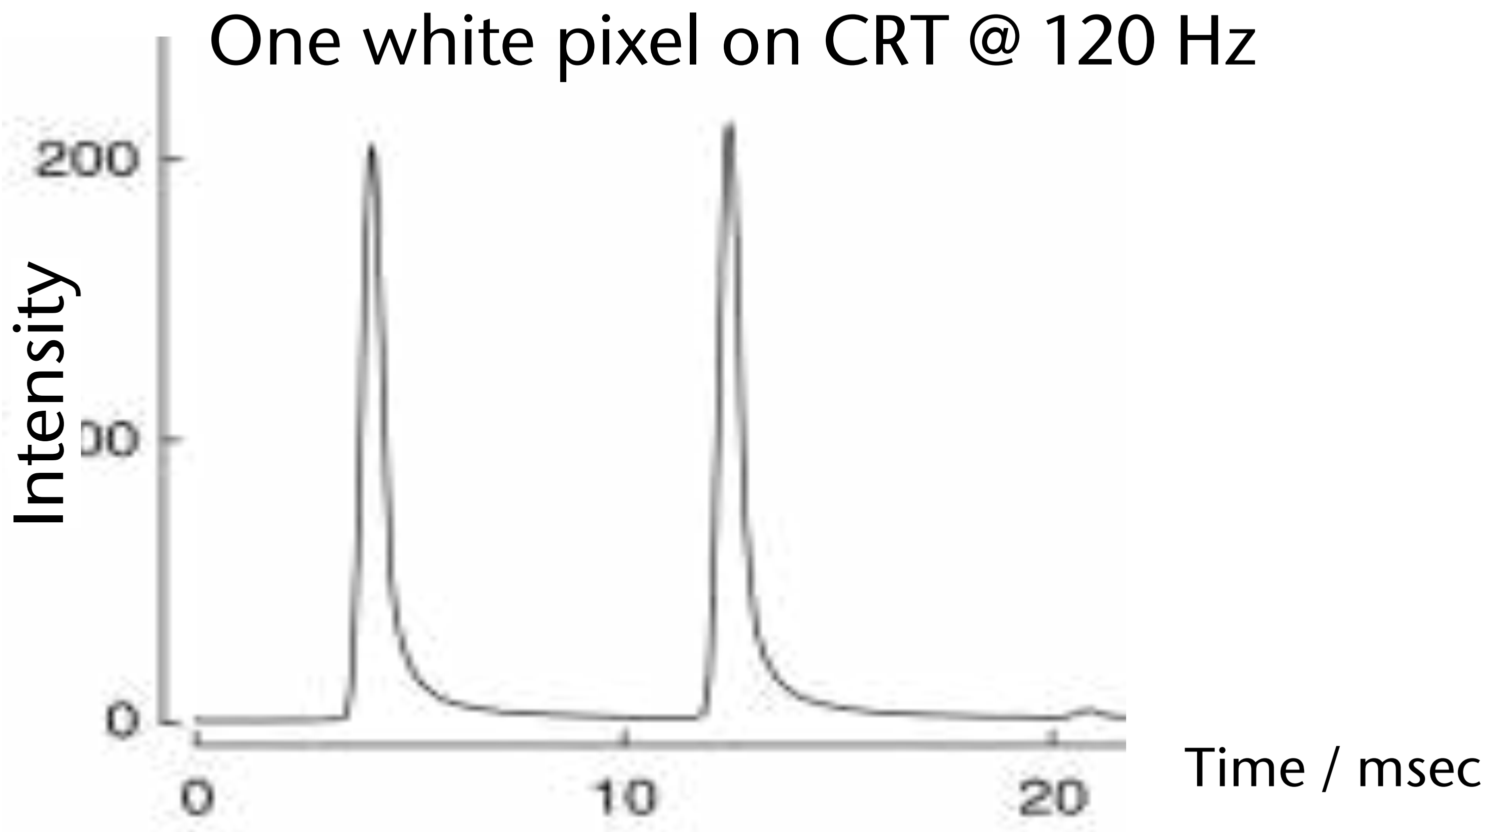

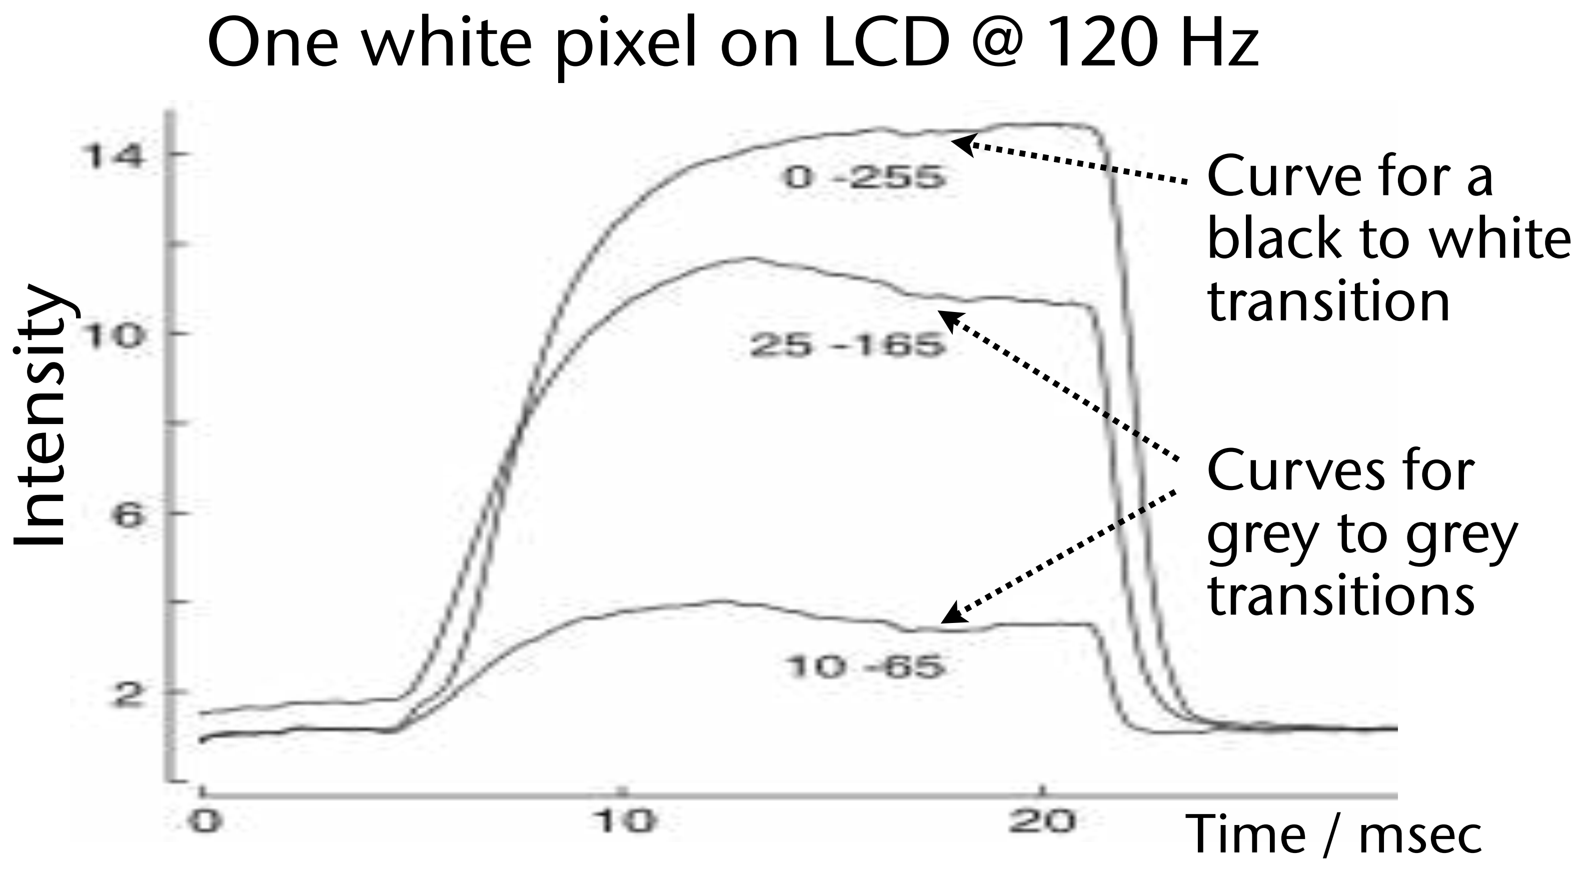

High-speed video in slow-motion, comparing an LCD and a CRT display

A Simple and Quick Experiment

- Grab a window with high-contrast borders with the mouse and drag it left and right with medium speed; with your eyes, follow one of the vertical borders (a.k.a. eye tracking)
- What (shape) do you see?

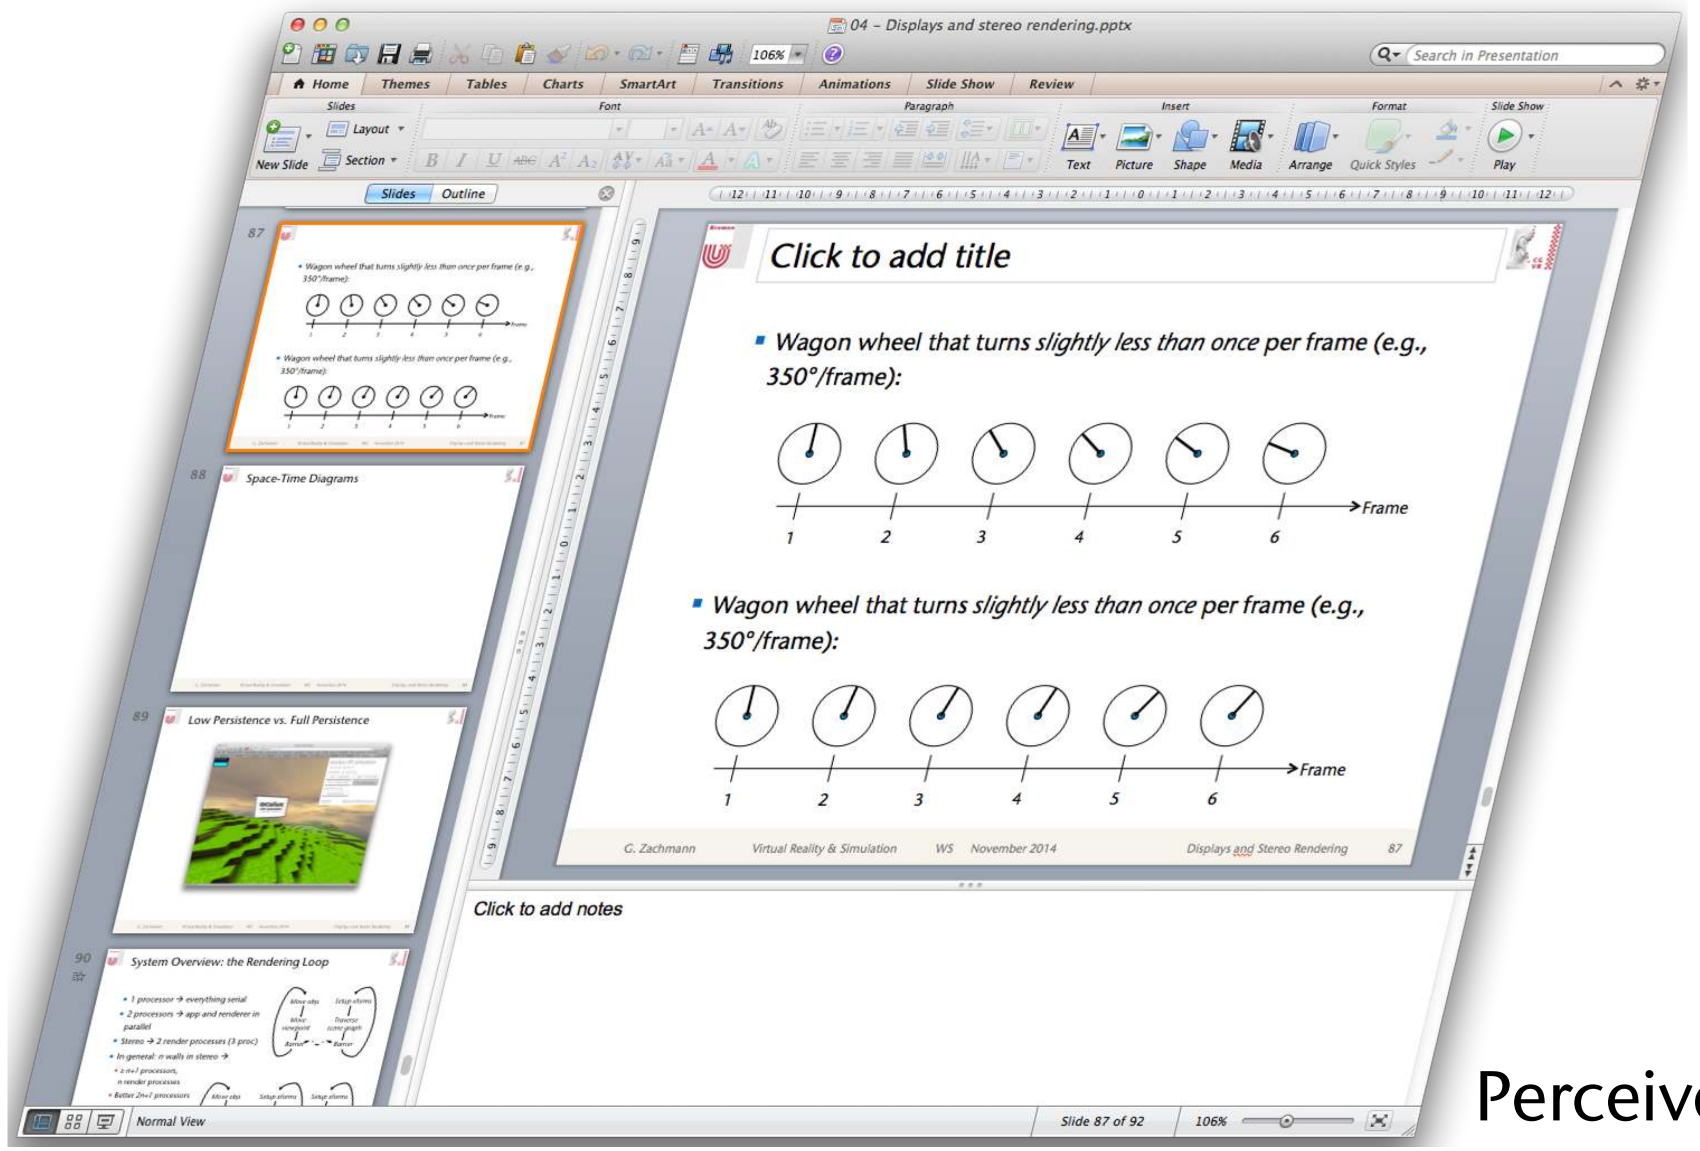

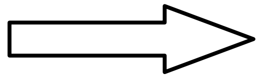

Direction of motion of the window

Perceived shape is exaggerated

Explanation

- Assumptions, for sake of simplicity:
  - Monitor is a CRT (for LCD, the argument works, too, but a bit more complicated, at least with full-persistence)
  - Graphics hardware waits for VSYNC before scanning out framebuffer
- The eye's fixation direction moves with constant speed across display
- Because scanlines are displayed one after another, pixels with same *x* coordinate on screen are projected onto positions on the *retina* with *different retinal x position*!

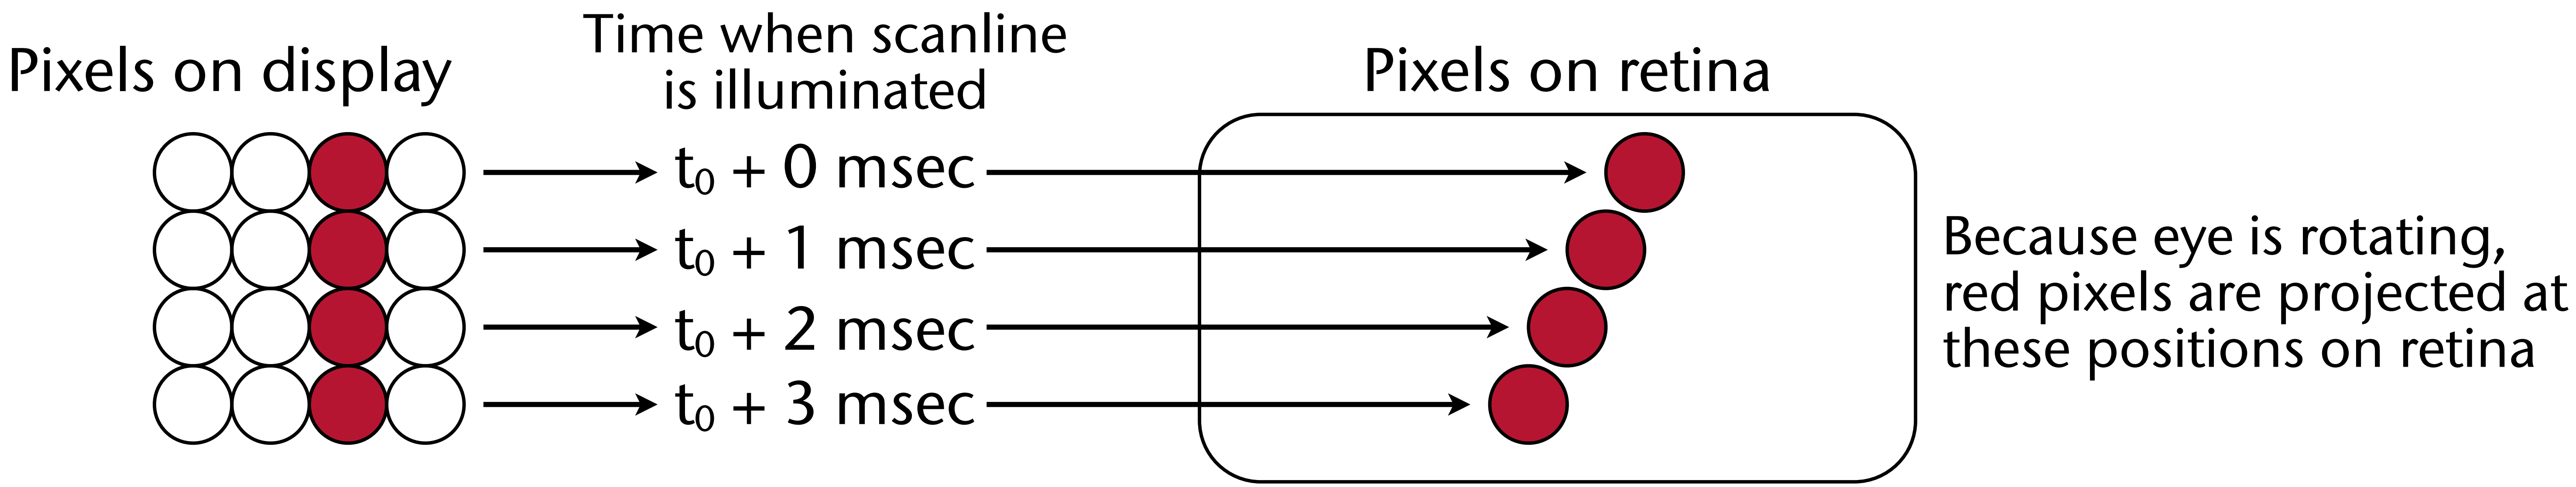

- If your graphics hardware does *not* wait for VSYNC, you might see something like this: *tearing & shearing*

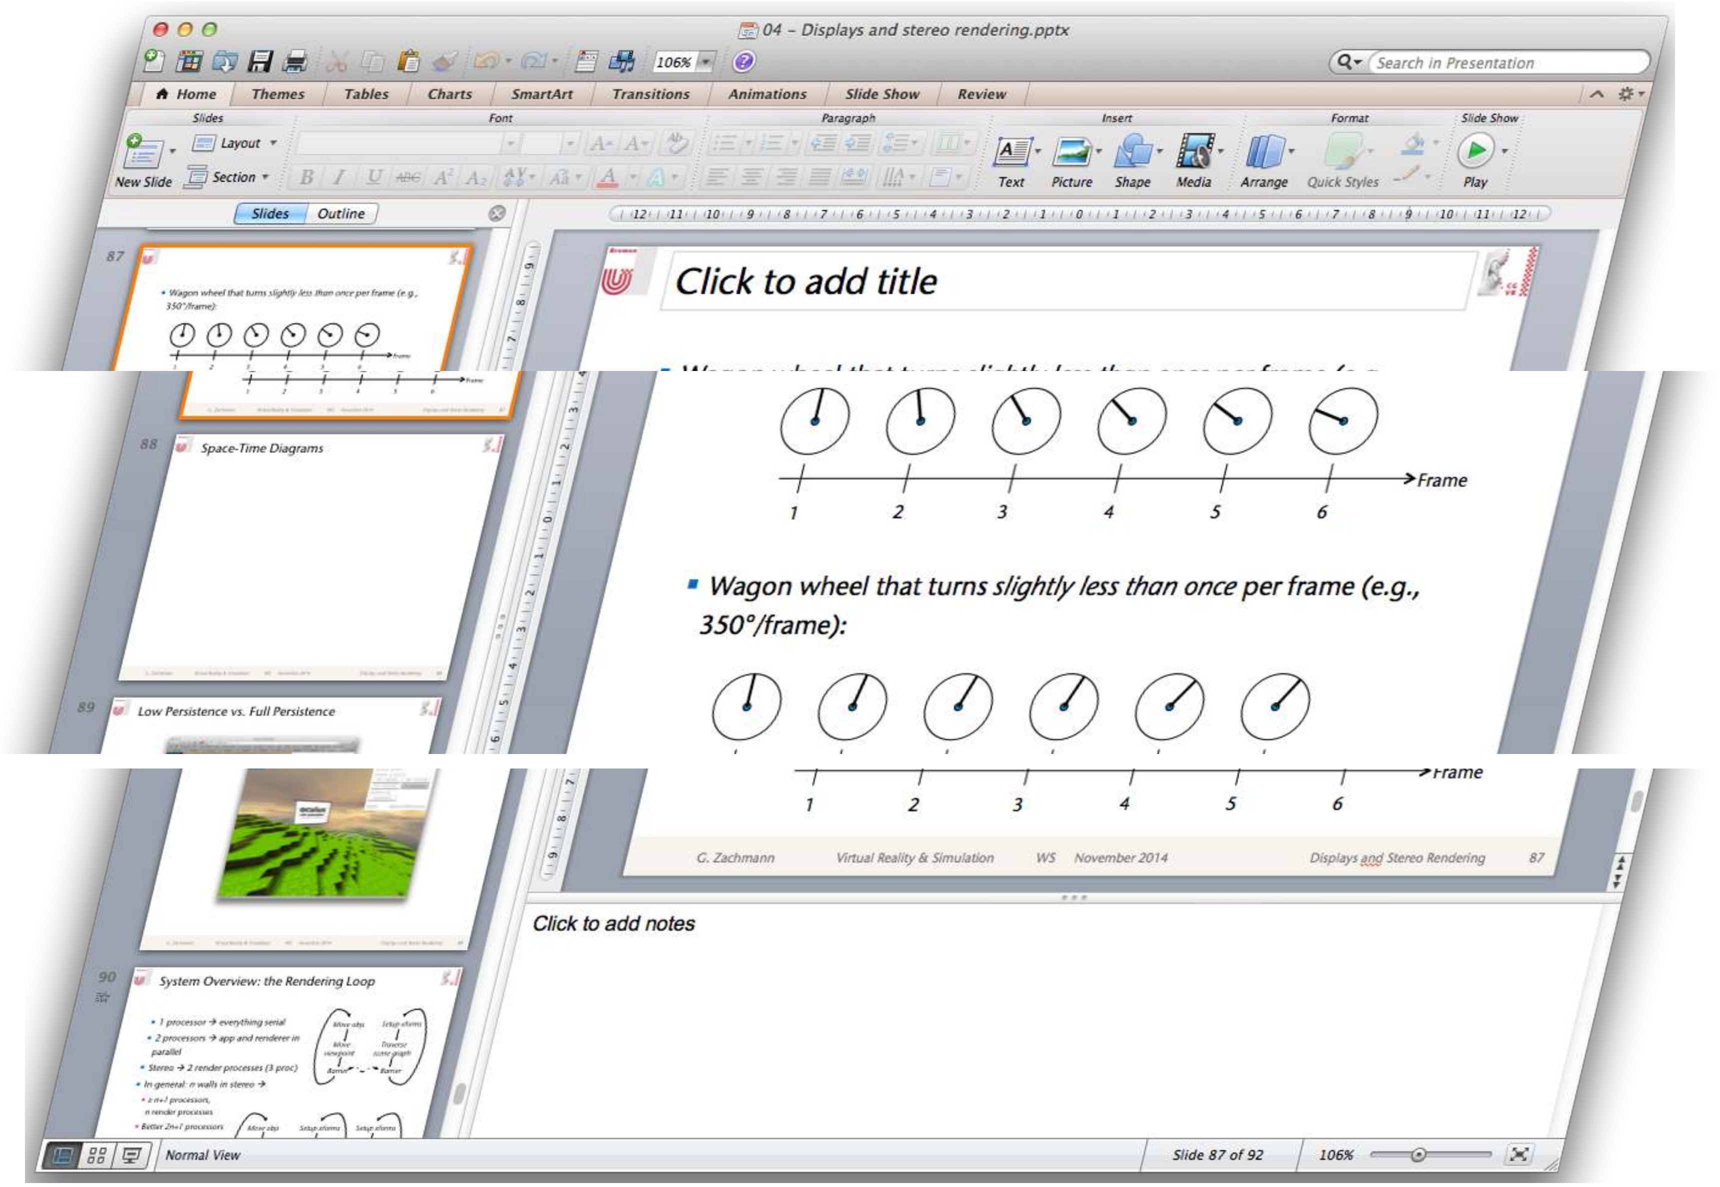

Temporal Aliasing (aka. Wagon-Wheel Effect)

- Wagon wheel with a rotation that is *slow relative* to the FPS:

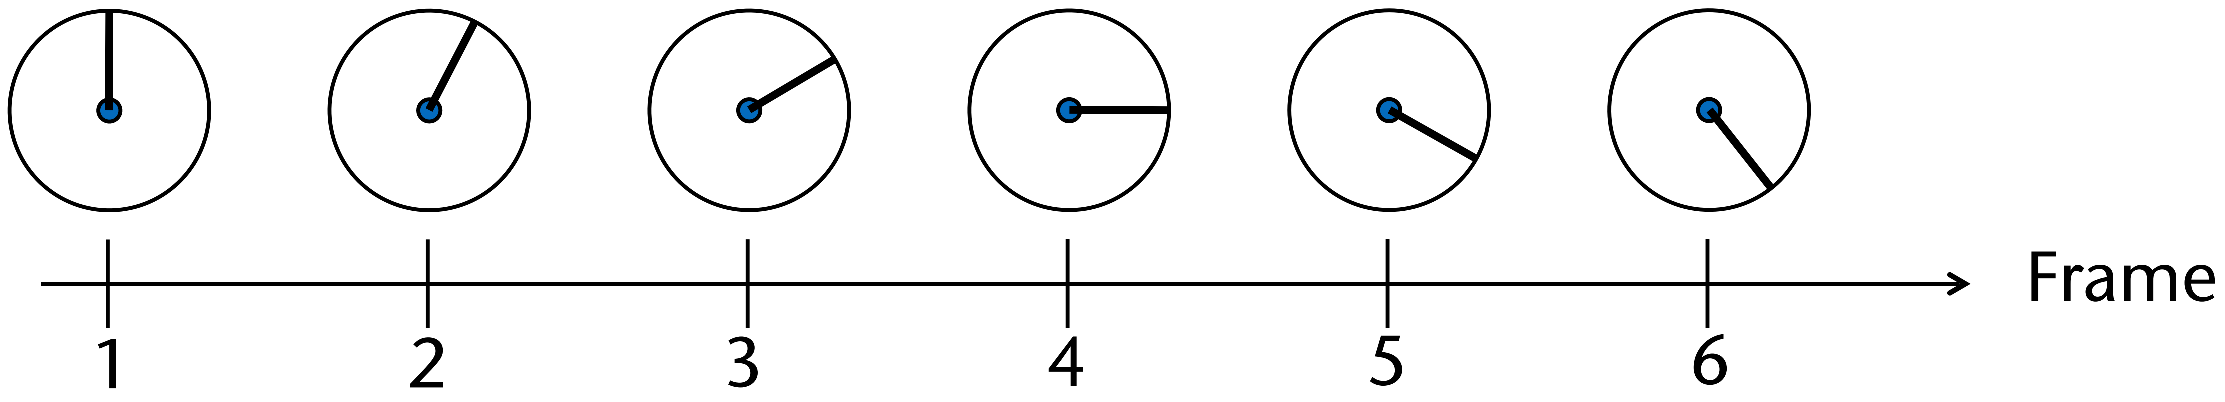

- Wagon wheel that turns *once per frame*:

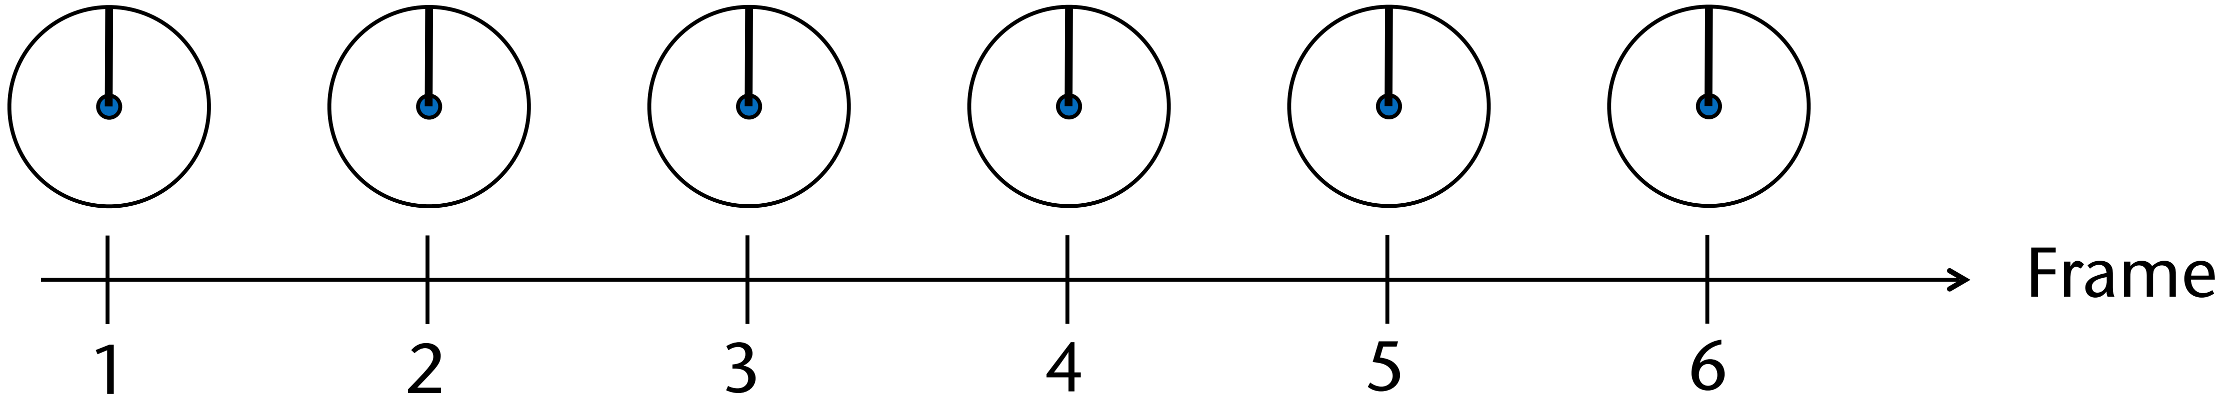

- Wagon wheel that turns *slightly less than once* per frame (e.g., 350°/frame):

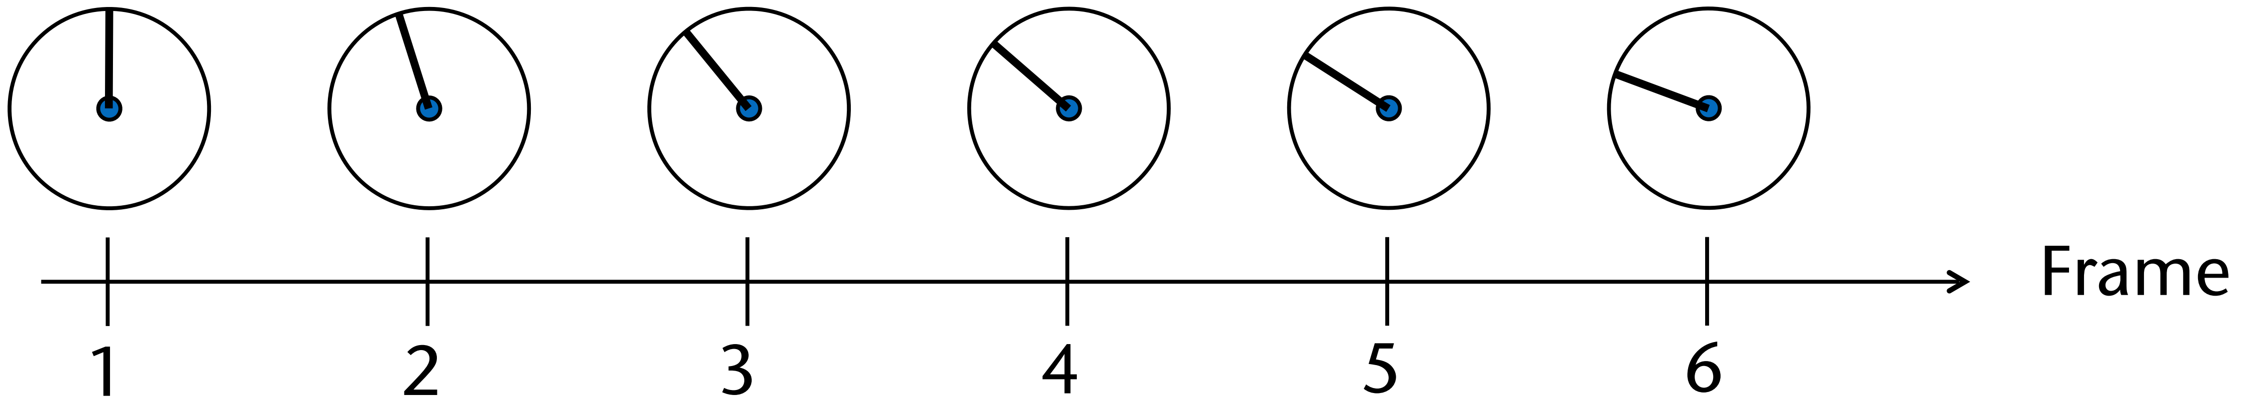

- Wagon wheel that turns *slightly faster than once* per frame (e.g., 370°/frame):

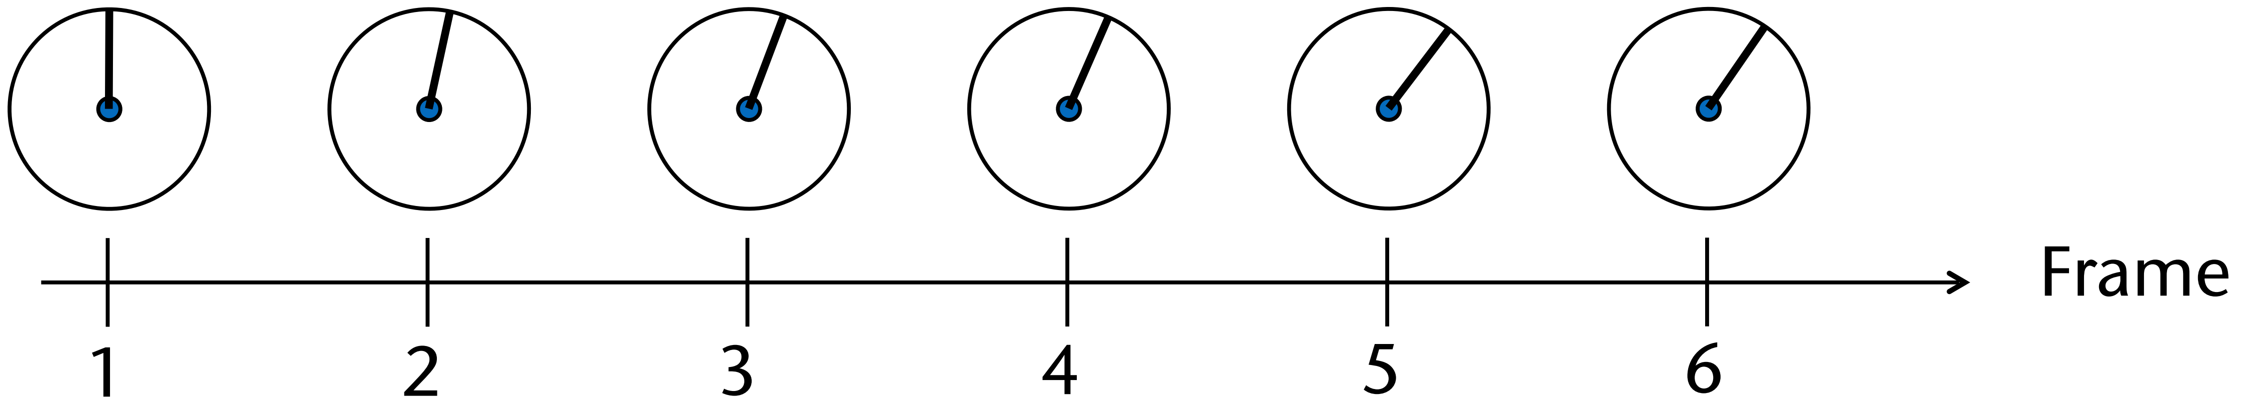

- Consequence: framerate (FPS) affects the display fidelity of motion being rendered on a screen!

- Note: this does not explain the shearing effect on the window

Digression: Utilization of the Temporal Aliasing

- Experiment setup:
  - Water droplets coming out of faucet at 60 Hz
  - Regular camera at 60 Hz
  - Strobe light at 60± Hz
- Effect: using just a regular camera, you can simulate a high-speed camera and produce a slow-motion video

Space-Time Diagrams

- Space-time diagram = curve showing positions of objects as a function of time, where *t* is one of the axes
- Example: 3D space-time diagram of a journey on a 2D map
  - I.e. curve = ($x(t)$, $y(t)$, t)
  - In general, they are 4-dimensional
- Simplification in the following: consider only the x-position of objects ⟶ 2D space-time diagrams
- Example: a point staying still on the *x*-axis

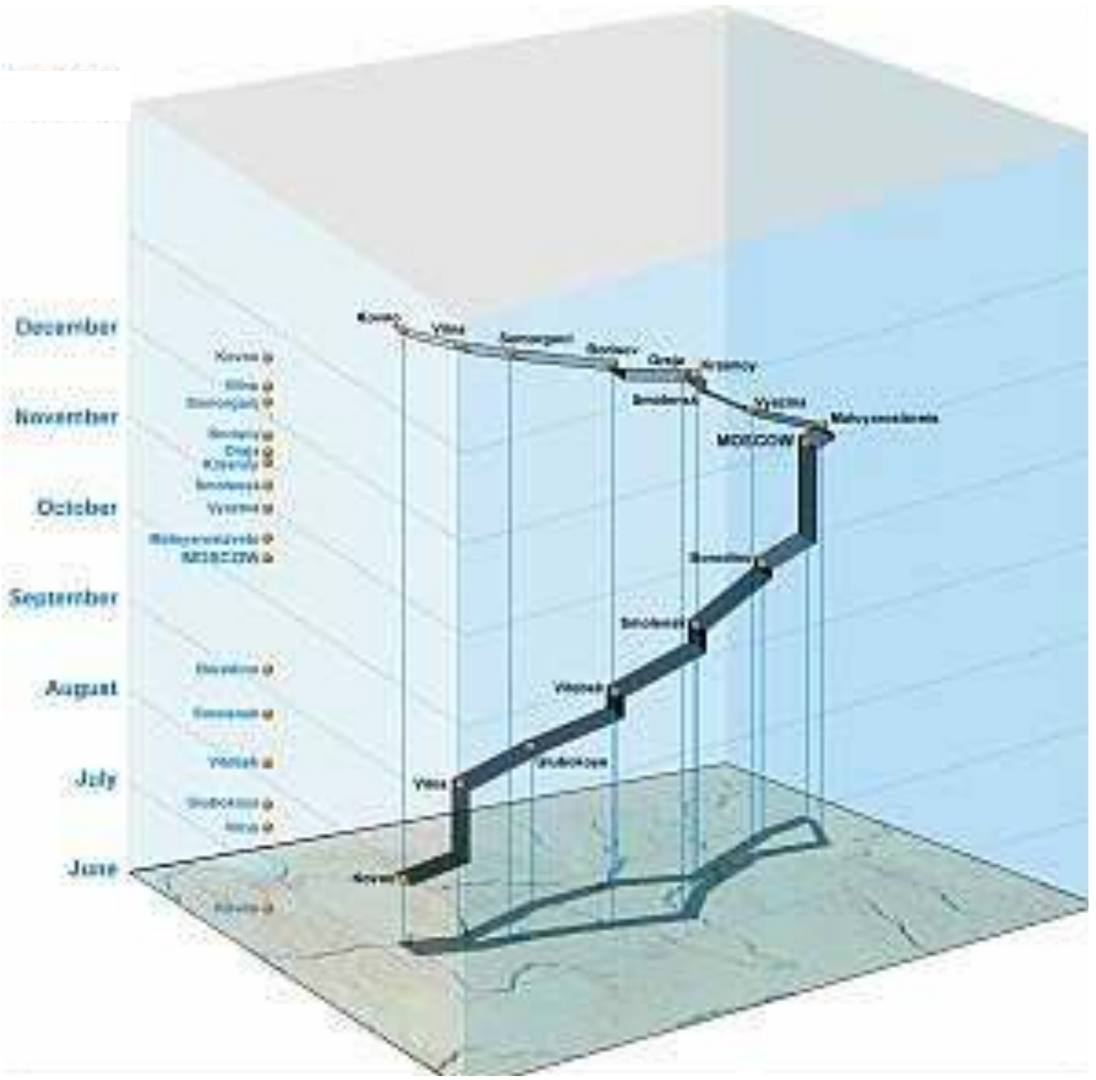

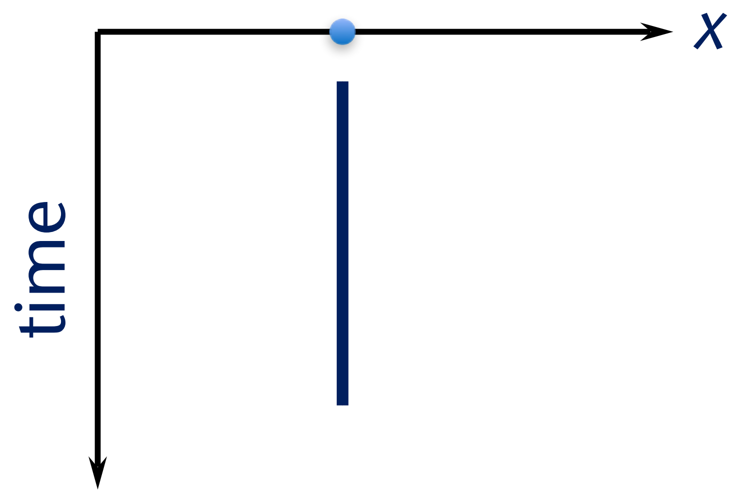

- Example: a point moving with constant speed along *x*

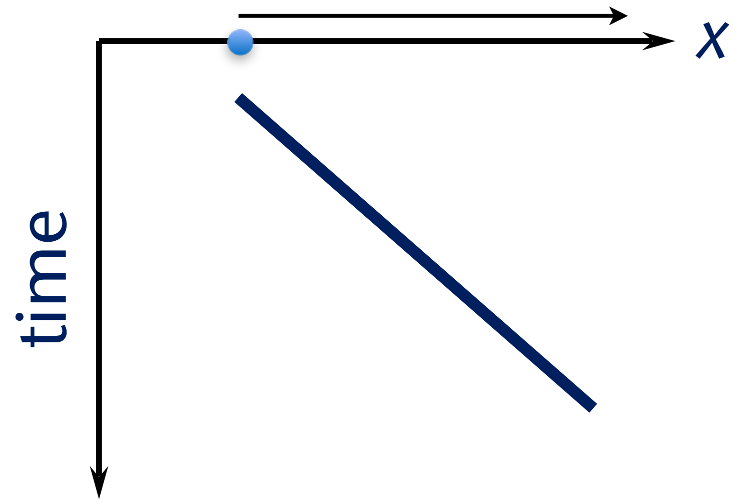

- A point in VE being moved steadily by a simulation along *x* with constant speed; space-time curve of its rendition on a monitor with full persistence
  - Remember: "sample-and-hold" display

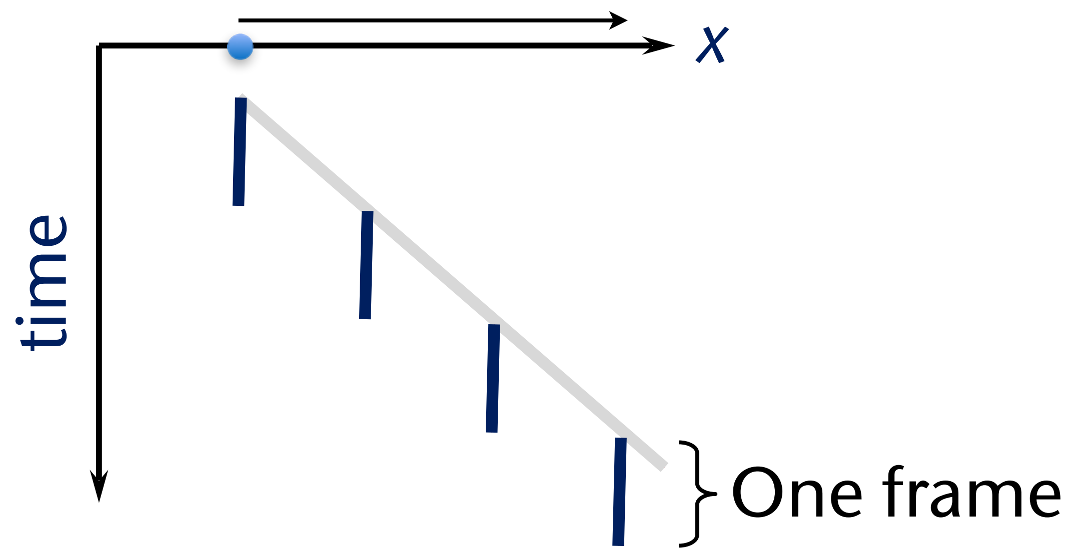

- Same again on *low persistence* display
  - E.g. CRT's, laser displays
  - LCD's and OLED's can be turned into low persistence displays (reduces brightness significantly)

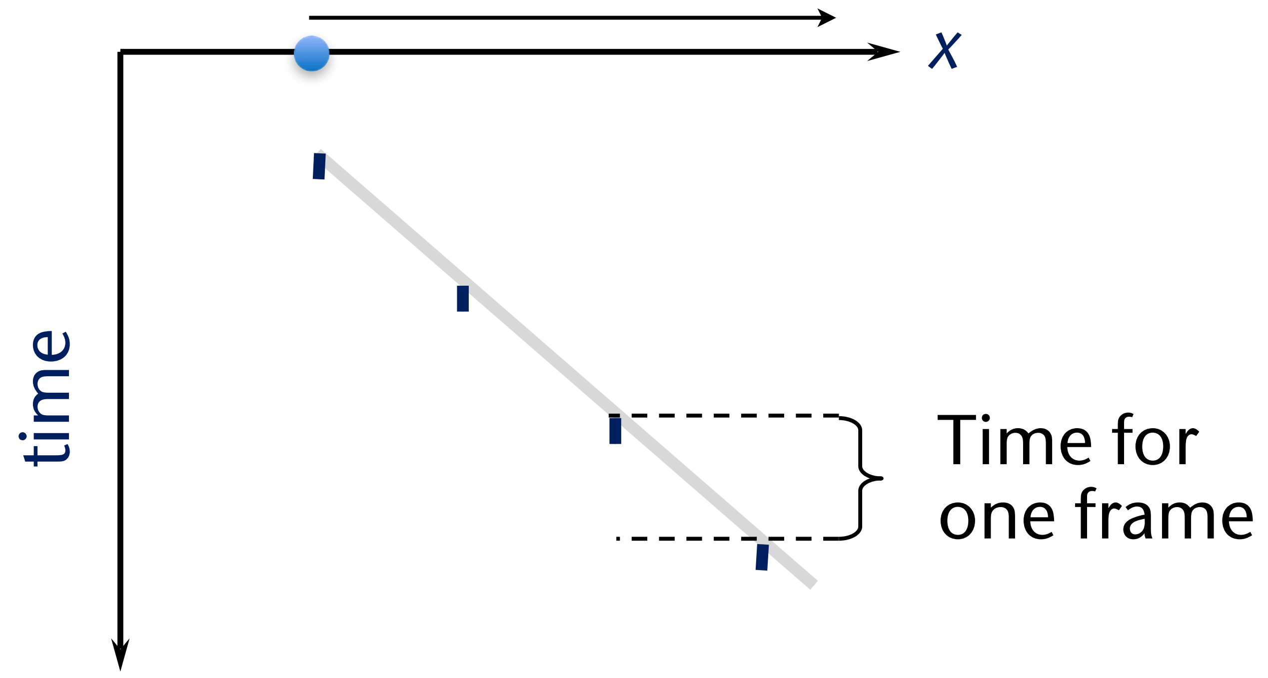

- Consider a slight change:
  - Point is moving in the VE along *x* at constant speed
  - Full-persistence monitor renders it at *n* FPS
  - Eye is *tracking* the virtual point (i.e. following its position)
  - What is the space-time diagram of the image of the virtual point *on the retina*?

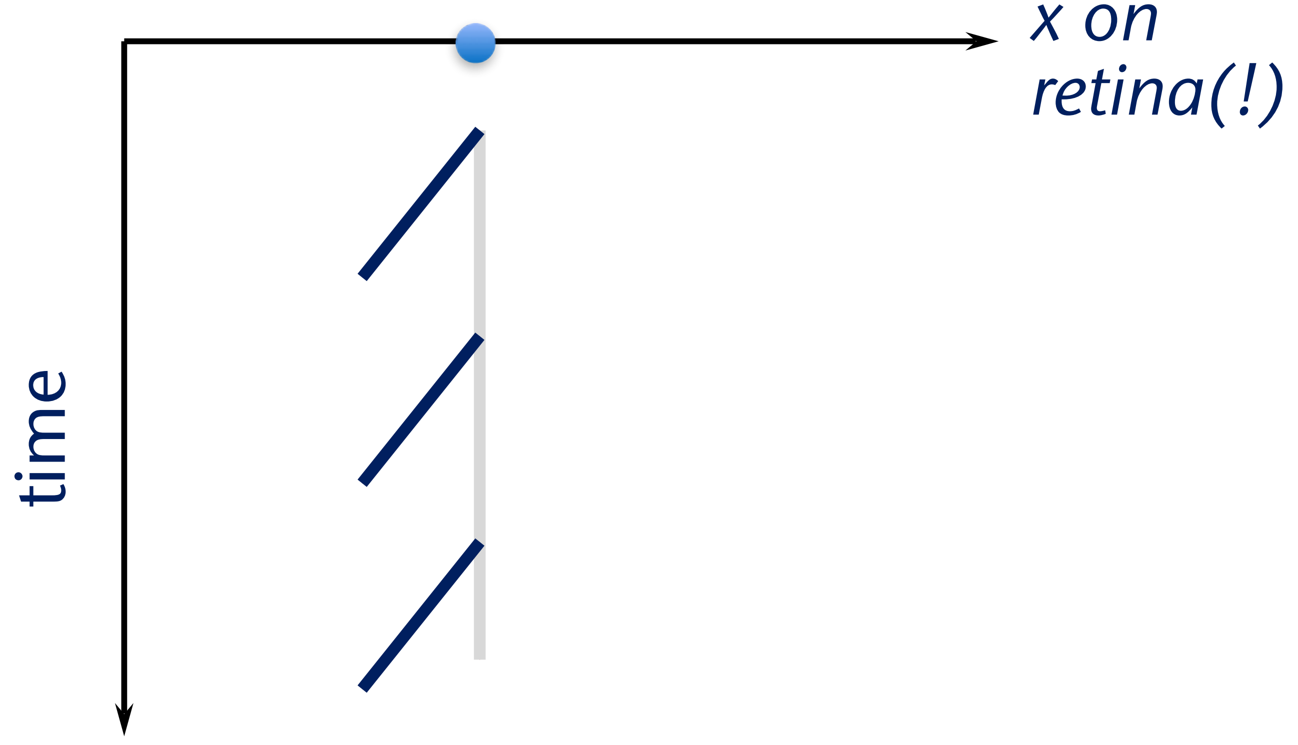

- This is the cause of judder and/or smear
- Effect of doubling the frame rate:
  - Still judder, but less "smeared out"
- Consider this case:
  - User is wearing an HMD
  - Point moves constantly in the VE
  - Assume *no* latency from HMD tracking to image
  - User tracks point such that eye fixates always the same pixel, i.e., HMD and eye do not move relative to each other; instead, user turns head
  - Space-time diagram of image of point on retina?

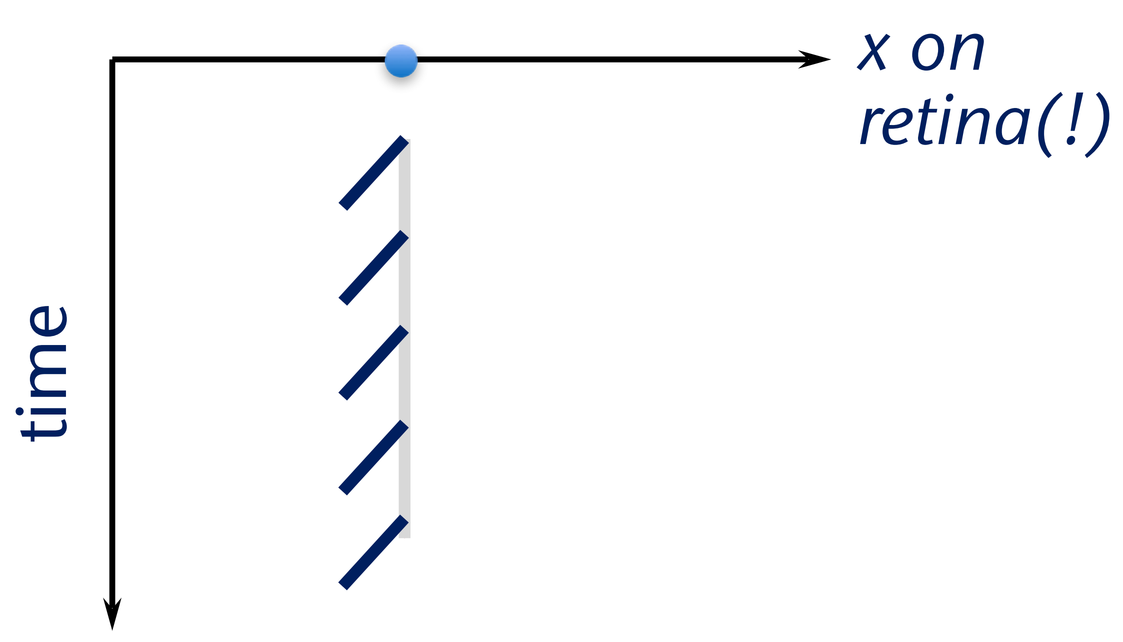

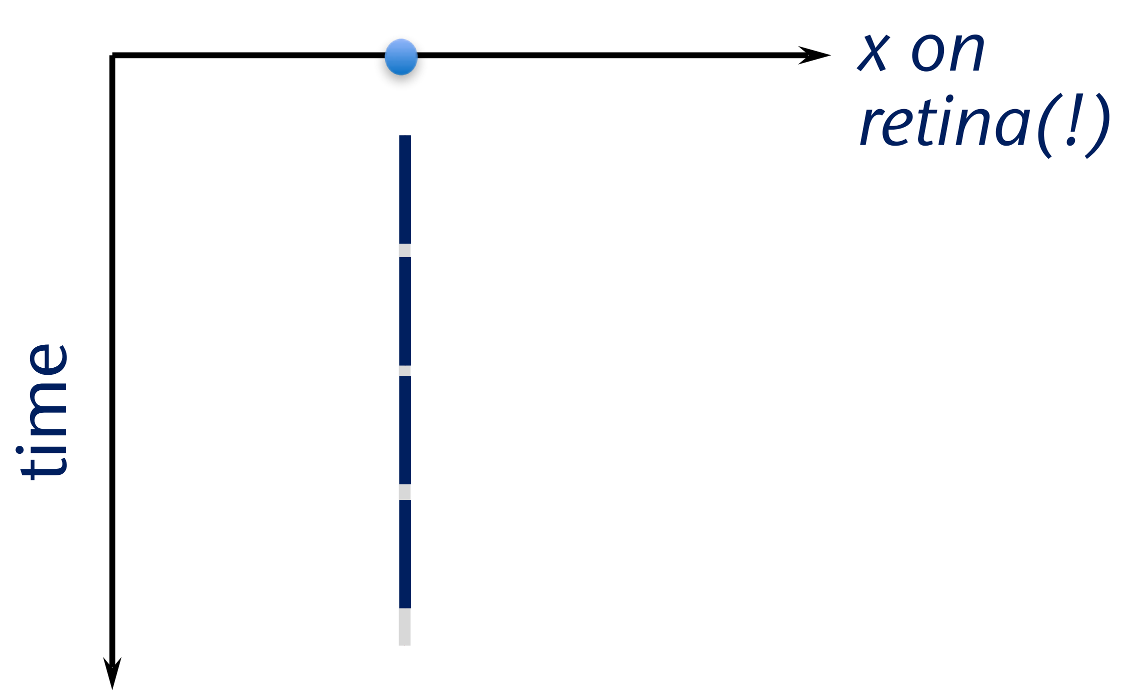

Animation of the Cause for Judder

Ideal HMD

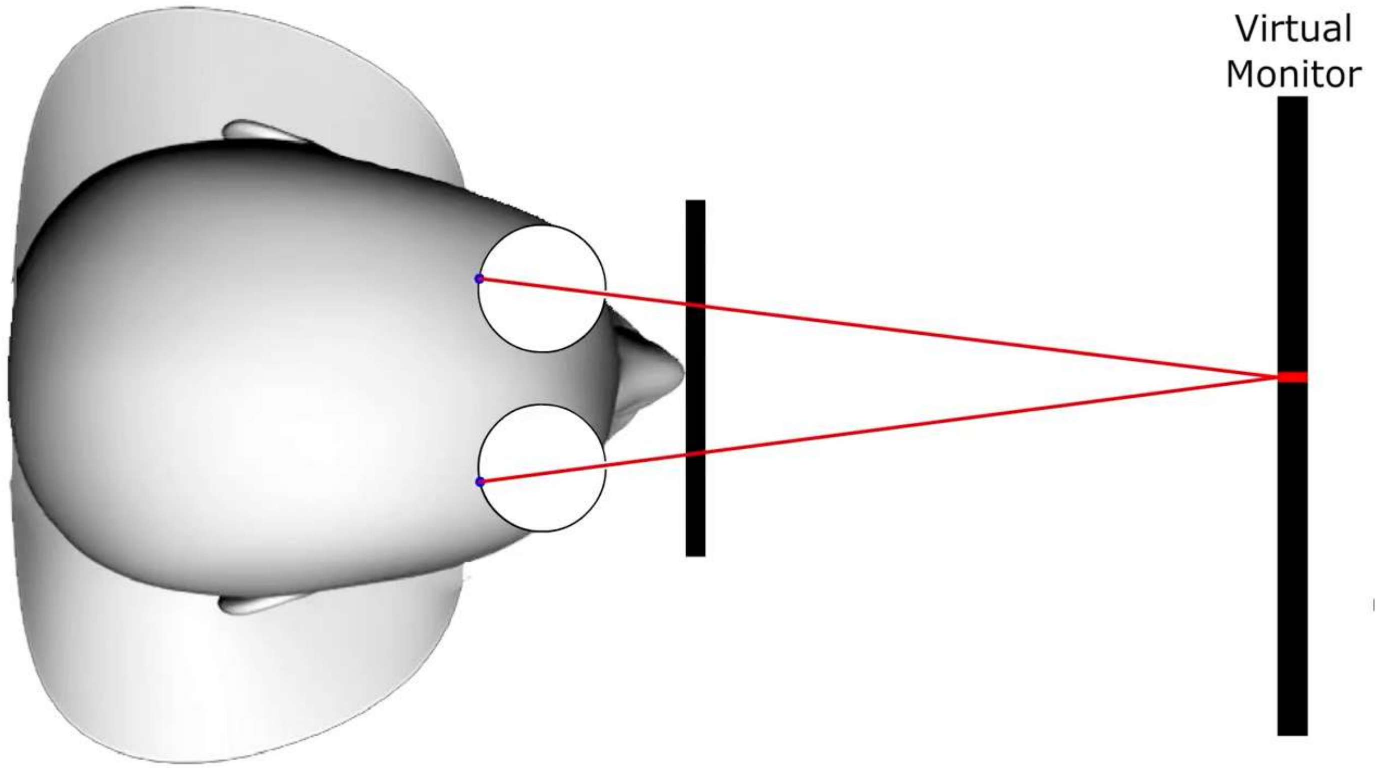

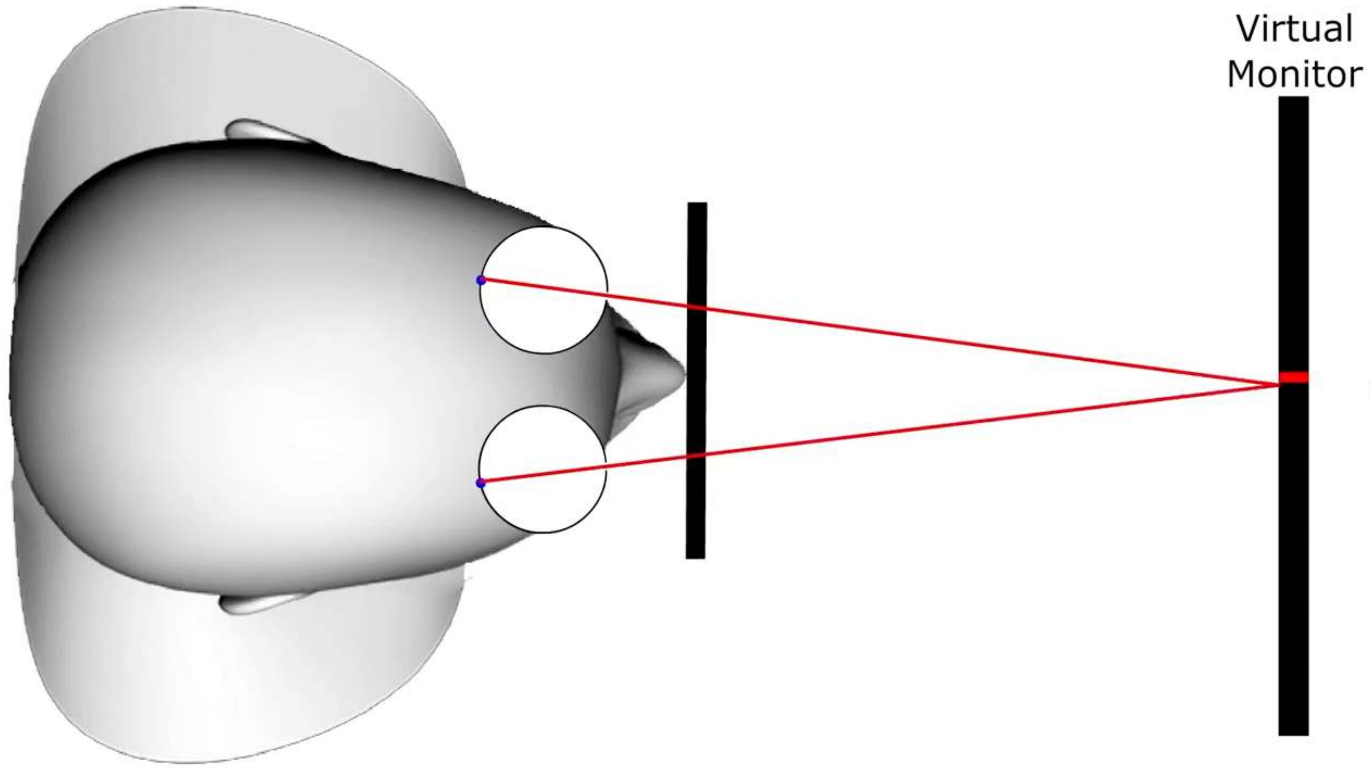

Full-persistence HMD

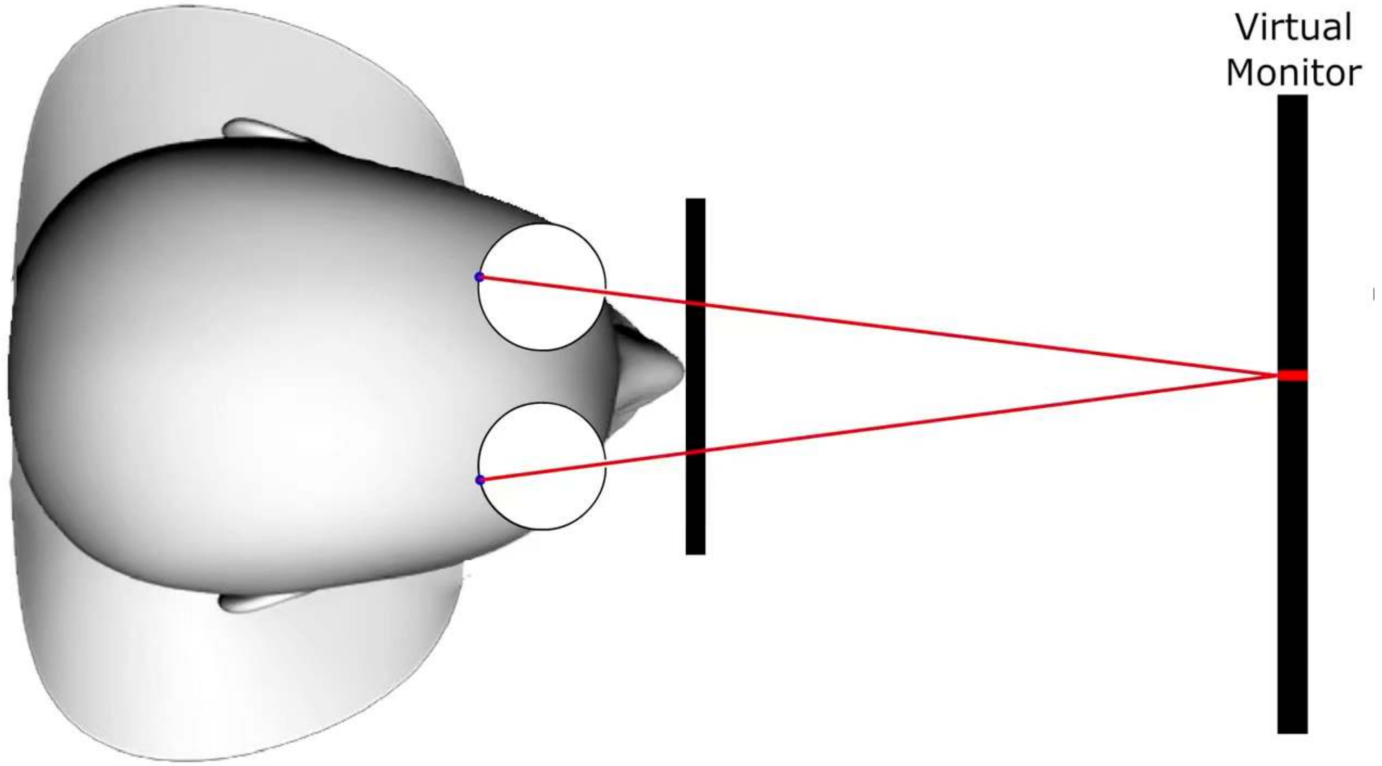

Low-persistence HMD

Assuming perfect HMD tracking, and 0 latency from the HMD tracker up to the frame buffer

Animations by eVRydayVR

How Blurry is Your Display?

1. Fixate the upper UFO with your eyes: you should see stationary black & white vertical stripes, with some black squares moving by
2. Track the lower UFO with your eyes – what do you see now?

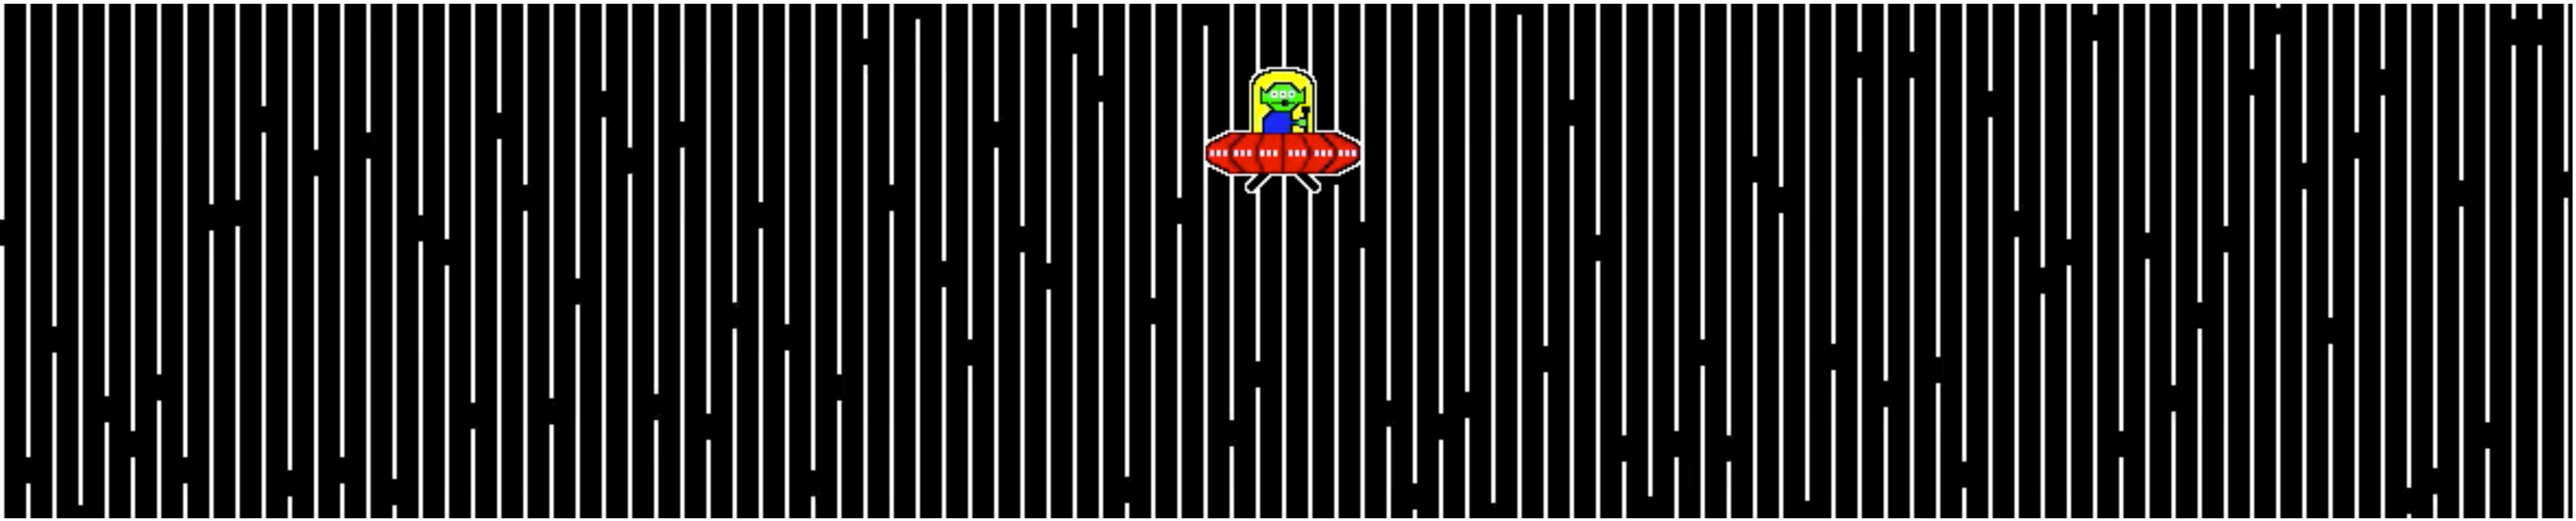

http://www.testufo.com/#test=eyetracking&pattern=stars

It Can Get Worse - With Field-Sequential Displays

- Field-sequential color (FSC) displays: first, only the red channel of all pixels of the frame buffer is transferred (and displayed), then the green channel, then blue channel
  - Reduces cost, size, wires, …
  - E.g., Google Glass used field-sequential color [2014] ; some (cheaper) projectors, too
- Space-time diagram of a stationary point on an FSC monitor

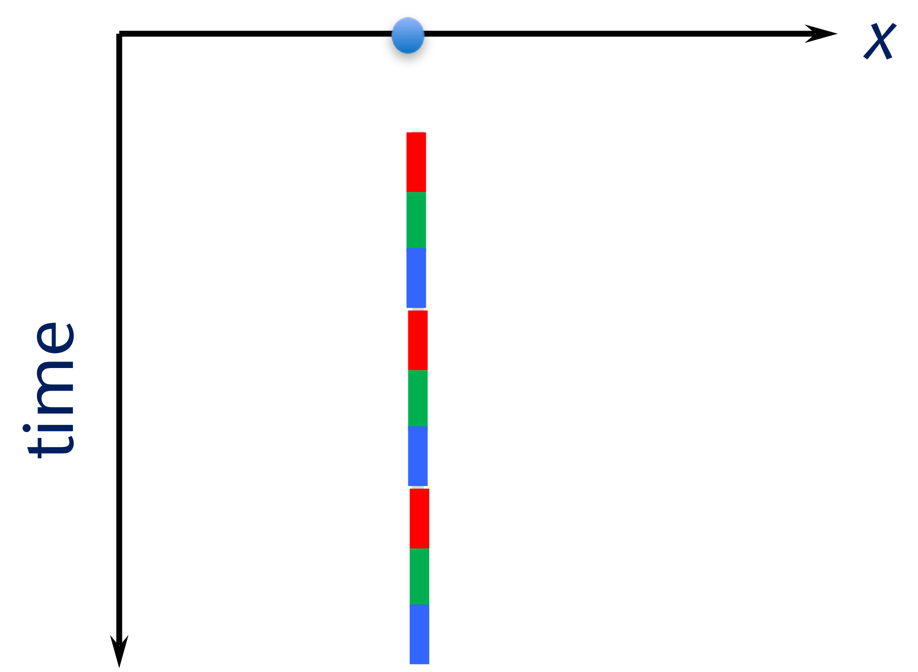

- Space-time diagram of a moving virtual point *on the retina*, with the eyes tracking its image on an FSC monitor

- Space-time diagram of a moving *disc* on the retina rendered on an FSC monitor, tracked by the eyes
  - Result: smear and *color fringes*!

- Similar stuff happens in HMD!

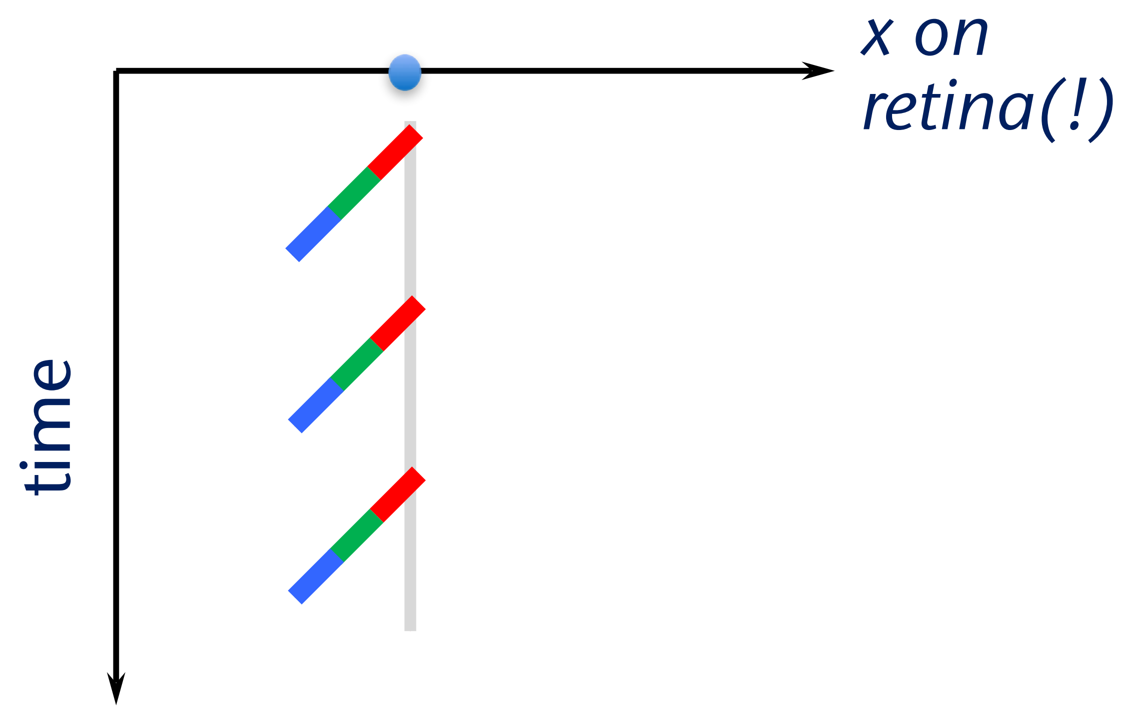

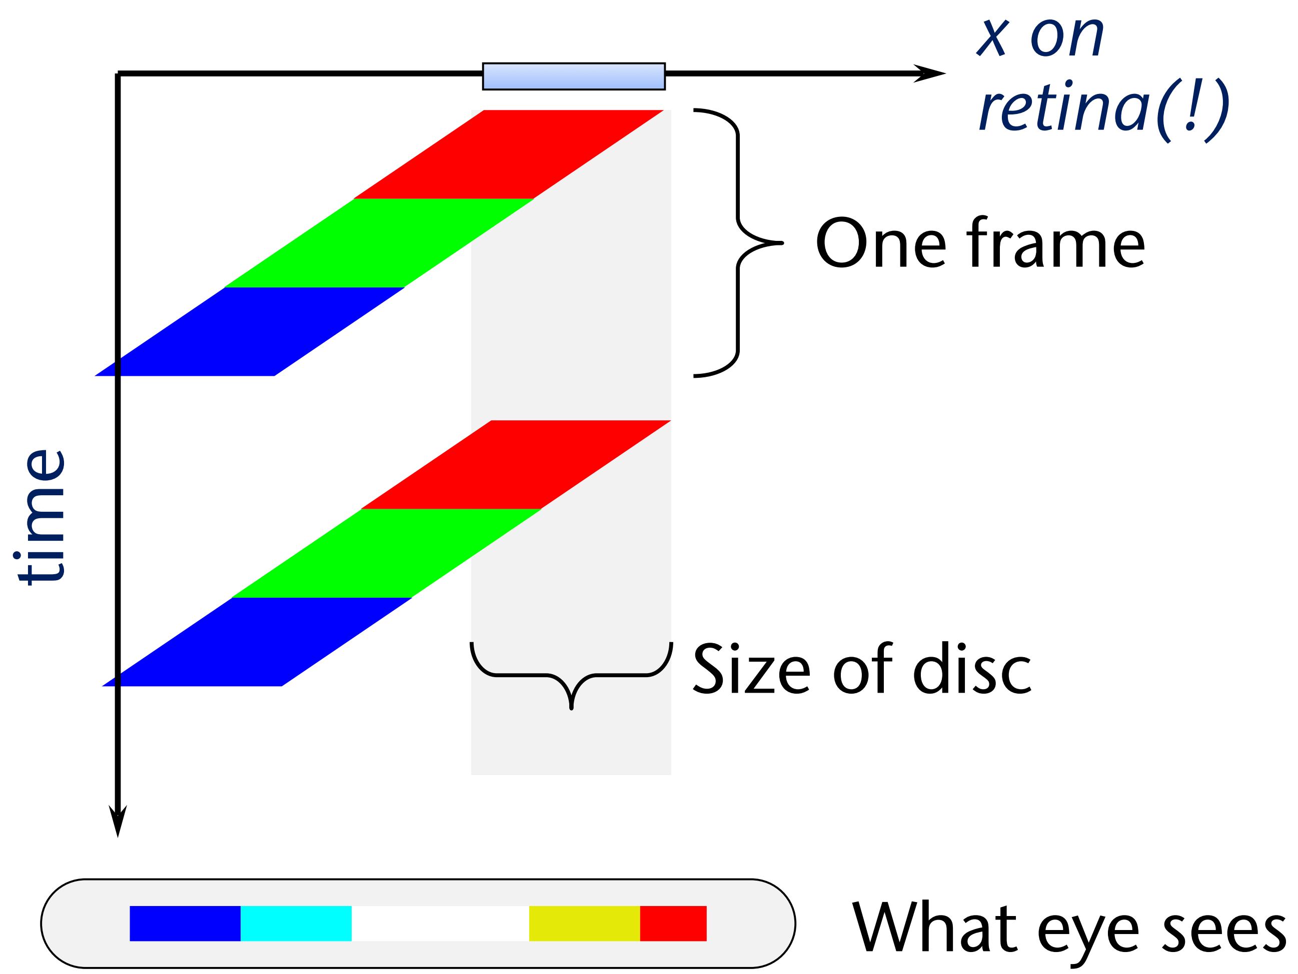

Possible Side-Effects of Low Persistence

- Low-persistence might introduce other problems
- Strobing: perception of multiple copies of the same object
  - Smear can hide strobing artifacts
- The short light bursts of a low-persistence display could interact/disturb saccadic masking
  - Saccadic masking = eye is effectively blind (to some degree) during a saccade
  - Consequence of the interaction: brain might lose frame of reference ⟶ visual instability

- Lots of perceptual research needed, and a good engineering idea!
  - The 1000 Hz display & rendering pipeline?