22 – Augmented and Virtual Reality

3D UI’s, Virtual Environments, VR, AR, …

- 3D Interaction
  - Interactive systems performed in real or virtual 3D space
  - Doesn’t require 3D displays, 3D devices
- 3D UIs are UIs with 3D interaction

LaViola et al - 3D UI's

- Virtual Environments are 3D worlds w/ real-time control by user
  - AR, VR, MR are variations of VEs
  - XR is ill-defined (I use it as xR, meaning any-R)

Huge variety of non-immersive 3D tech

Image from https://commons.wikimedia.org/wiki/File:NTTICC_ImmersaDesk.jpg

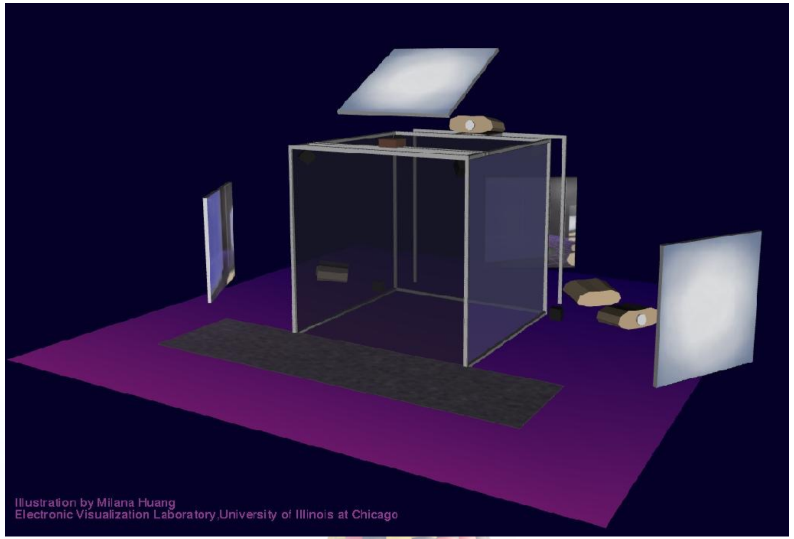

Image from https://zspace.com

What is Virtual Reality?

- Immerse in a VE with real-time tracking and sensing, immersive displays that try to replace a user's perception of the world

What is Augmented Reality?

- Enhance user’s senses with computer generated information
- Augment, not replace using see-through displays

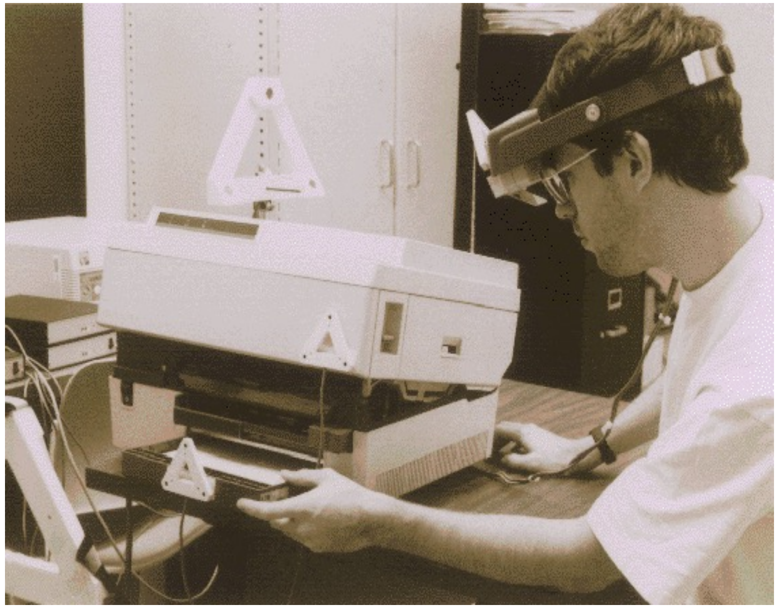

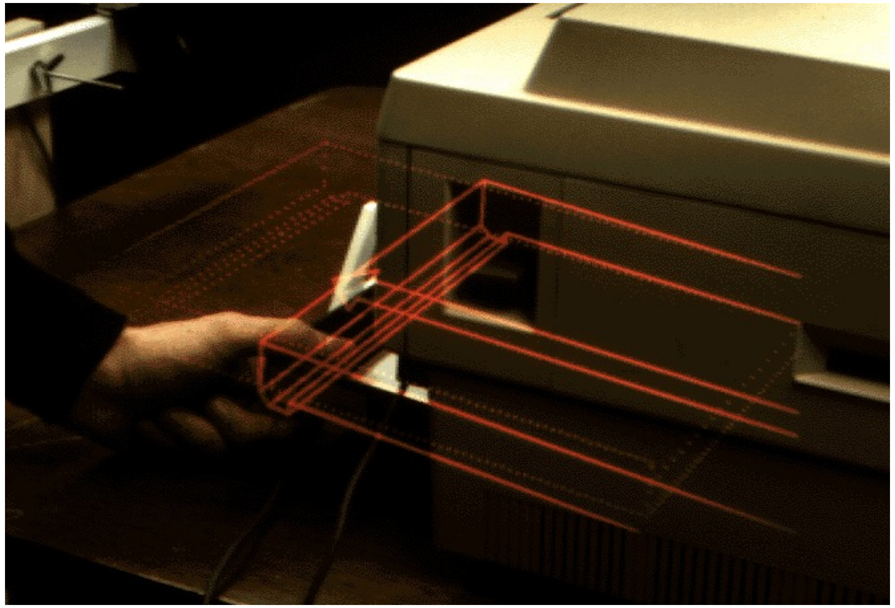

Images courtesy of Blair MacIntyre. Copyright Columbia University.

In 1994, Milgram Defined AR in Relation to VR and Reality

- AR sits between the real world and virtual reality, combining both. Some amount of real-world knowledge is needed to know where to put the virtual content.
- AR is closer to the real world than to the virtual world, since most of what the user sees is real.

Technical View → Display

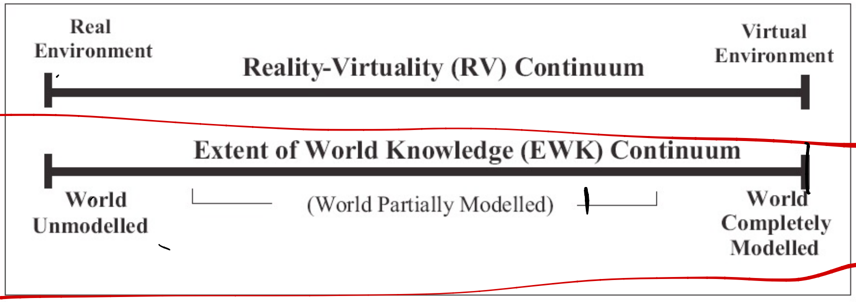

Figure 1.1 Reality-Virtuality (RV) Continuum, in parallel with Extent of World Knowledge (EWK) Continuum.

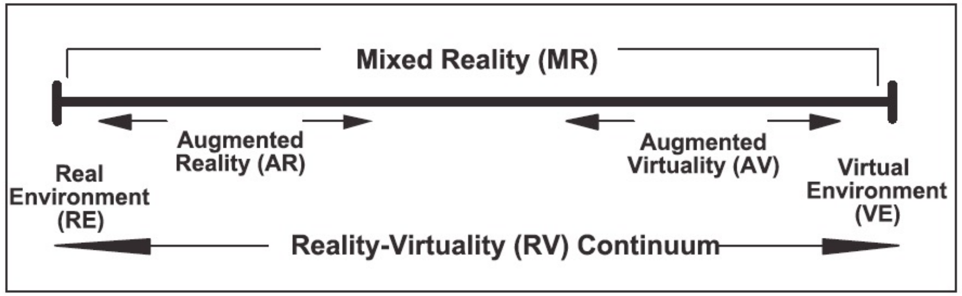

Figure 1.7 Definition of Mixed Reality, within the Context of the RV Continuum.

Immersive AR/VR

- Perceptual immersive (or partial immersion) in VE
  - Sense/track user's head pose relative to world
  - Content appears to be "around you"
- Technology Components:
  - Tracking and sensing head, hand, objects
  - Creating images
  - Displaying images to user

The First 3D (AR) Display

- Ivan Sutherland led creation of the first see-through head-mounted AR display at Harvard in 1968
- The display was tracked in 3D by the mechanical linkage holding it
  to the ceiling

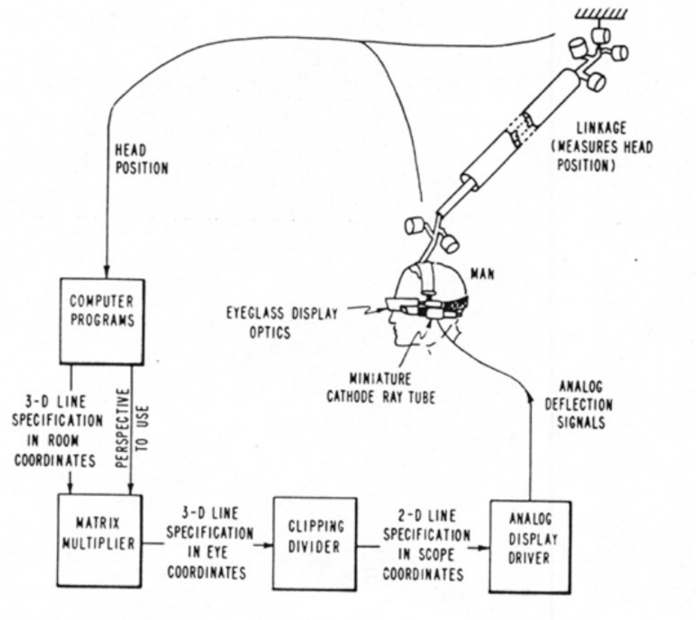

FIGURE 1—The parts of the three-dimensional display system

Ivan E. Sutherland, *A Head-Mounted Three-Dimensional Display*, 1968 Fall Joint Computer Conference

Computer History Museum: Sutherland and Sproul
https://www.youtube.com/watch?v=Y2AIDHjylMI

Nothing till the 80's and 90's

- Jaron Lanier coined the term "Virtual Reality", attempted to commercialize VR at VPL (created Eyephone, Dataglove, DataSuit)
- Too expensive, too low quality
  - computers
  - displays
  - Input
    - Used magnetic trackers
    - Gloves and suit fragile

Image from https://en.wikipedia.org/wiki/Wired_glove

Augmented Reality begins around the same time, late 1980s

- NASA Ames did AR experiments in the late 1980's
- In 1990, a group of Boeing researchers imagines aligning manufacturing instructions with a person's view of the world using see-through head-worn displays
- The project eventually results in a working prototype to assist with wire-bundle assembly in 1997

Image courtesy of David Mizell.

After VPL, the VR Hype Bubble Burst

- VR didn't live up to the hype
- Problem: no obvious market
  - Cost / benefit, esp given quality
- Motion sickness a big problem
  - Latency a huge issue

- Work continued in labs through 1990's and 2000's on both AR and VR

AR research in the 1990's: KARMA

- KARMA (Knowledge-based Augmented Reality Maintenance Assistant) was one of the first AR demonstration systems, 1991 at Columbia University

Using a hand made display and four computer to handle the computation, KARMA had an integrated knowledge-based system that design maintenance instructions, dynamically based on where the user was looking, where the objects were located, and what needed to be communicated.

Images courtesy of Blair MacIntyre. Copyright Columbia University.

Application Research in AR and VR: e.g., Merging Ultrasound Imagery

- Henry Fuchs and others at UNC Chapel Hill presented a first demonstration of merging ultrasound data of a fetus onto a live patient in 1992.
- In 1996, doctors tested a follow-on system for needle biopsy guidance.

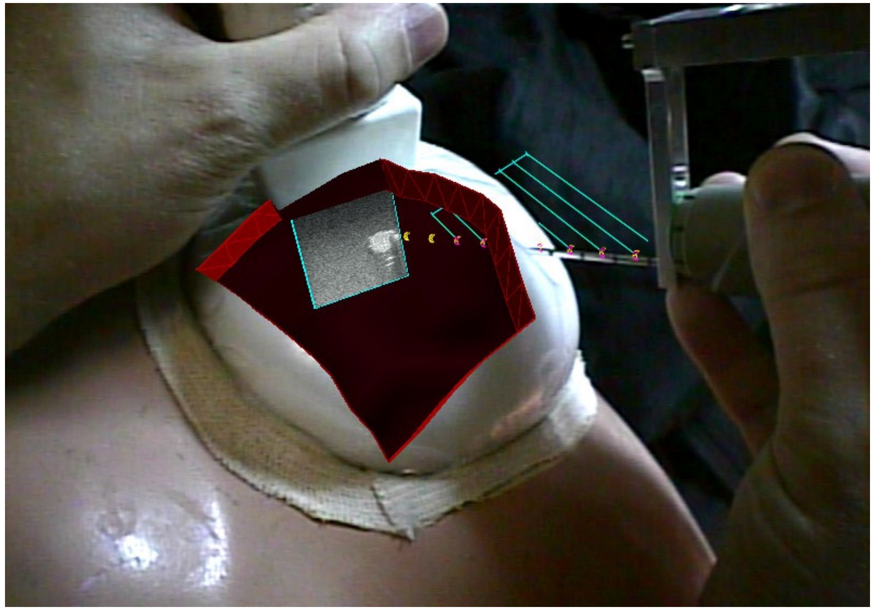

Images courtesy of Ron Azuma. Copyright UNC Chapel Hill.

Application Research in AR and VR: e.g., Studierstube

- In 1996, Dieter Schmalstieg created Studierstube, the first comprehensive architecture for collaborative AR and VR.
- Many applications and technology concepts were developed and demonstrated in Studierstube and other research and commercial systems over the years

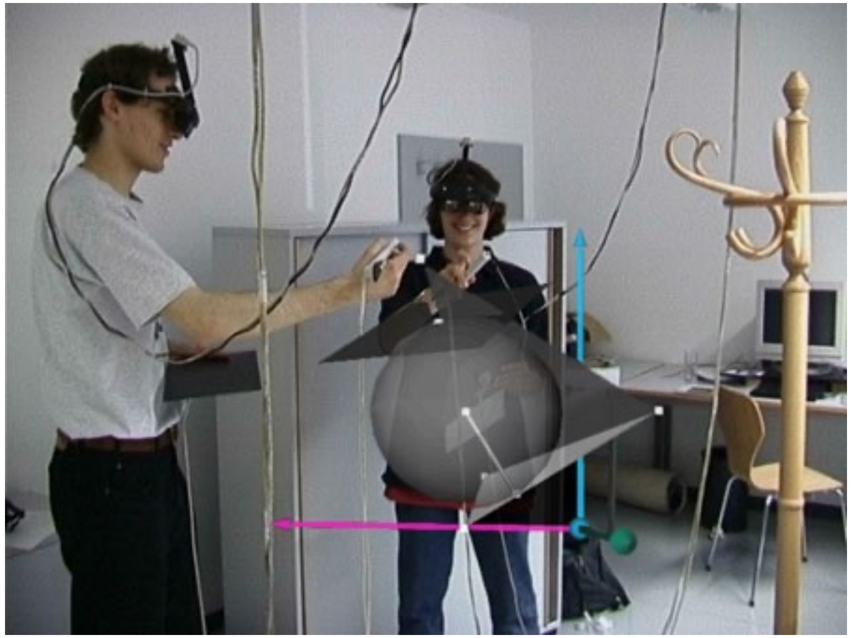

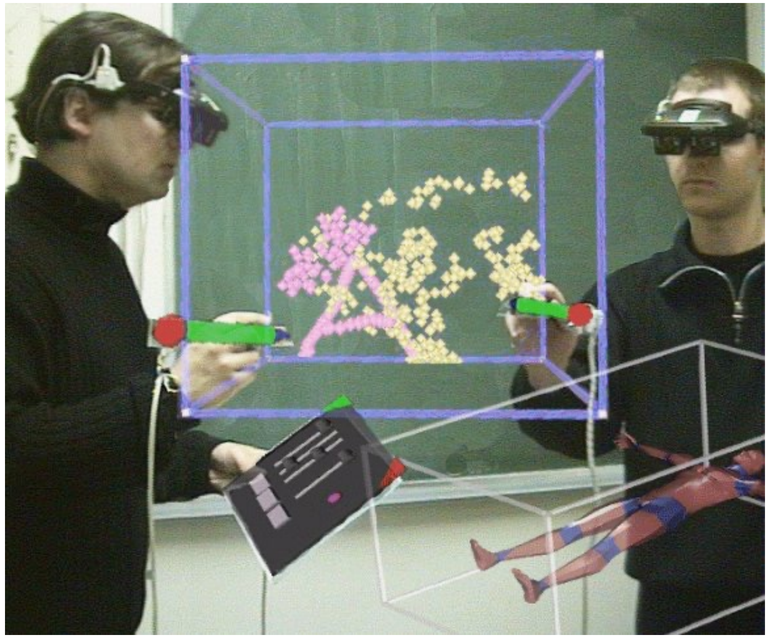

Images courtesy of Dieter Schmalstieg. Copyright TU Vienna.

Core Problem: Tracking and Registration

- Throughout the early 1990's, researchers at UNC Chapel Hill worked on a high-accuracy optical ceiling tracker
- Ronald Azuma's PhD work demonstrated near-perfect registration of a simple scene using this system combined with inertial sensors and carefully integrated with custom rendering hardware and a custom head-worn display

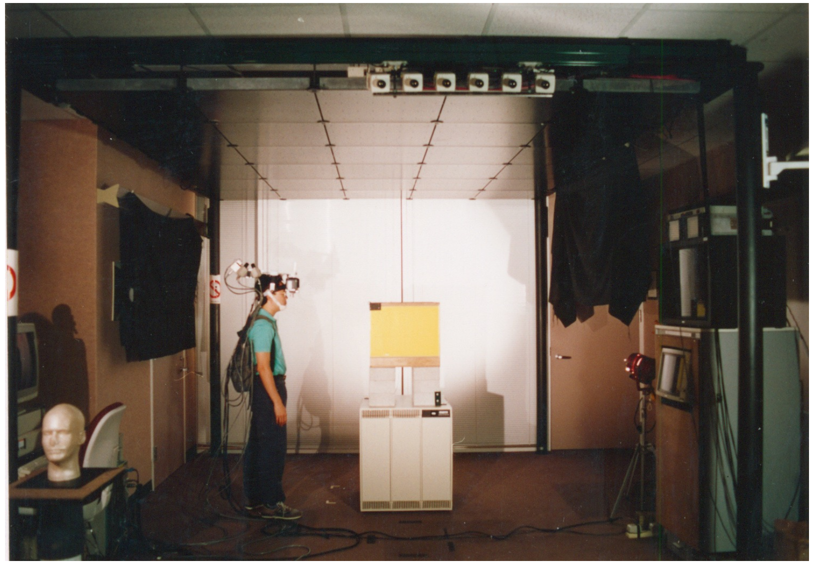

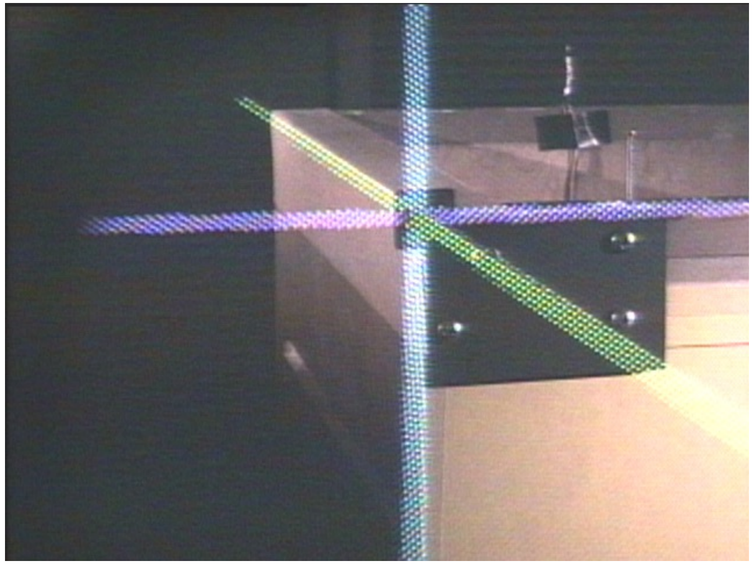

FoxTrax and Broadcast “AR”

- The NHL started US broadcasts in 1995, and a common complaint was difficulty following the puck. The FoxTrax puck was developed by Fox Sports, and used between January, 1996, and the end of the 1998 season.
- The technology was spun off in 1998 to form Sportvision, which created much of the broadcast TV sport enhancements seen today.
- These systems work because the fixed broadcast cameras are heavily instrumented, the environment is measured and accurately modeled, and the the broadcast signal can be delayed to give engineers time to correct problems. Often, multiple live engineers monitor a system.

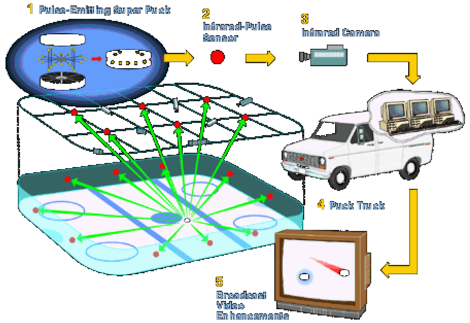

AR Moves Outdoors: The Touring Machine

- Built in 1996 at Columbia University, the "Touring Machine" was the first working demonstration of a mobile AR tour guide
- Backpack contains
  - GPS, wireless network, and computer
- Display overlays labels on the surrounding buildings
  - handheld tablet displays additional detailed information


Images courtesy of Blair MacIntyre. Copyright Columbia University.

Return of VR: the Oculus Kickstarter in 2012

- Palmer Lucky was a student of Mark Bolas at USC
  - Inspired by Bolas' novel displays, such as the Wide5
- Wildly successful Kickstarter
  - DK1, DK2 delivered. Had issues, but potential, and cheap
  - John Carmack (Doom, Quake, ID Software) early enthusiast
- Bought by Facebook in 2014 for $2-3 billion
  - Kicked off current VR boom
  - Had resources to attack core problems such as latency

VR is now a consumer technology

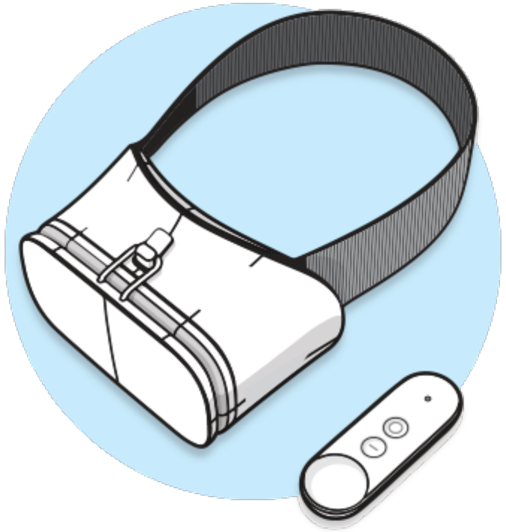

All Senses have been explored, attempted to be commercialized

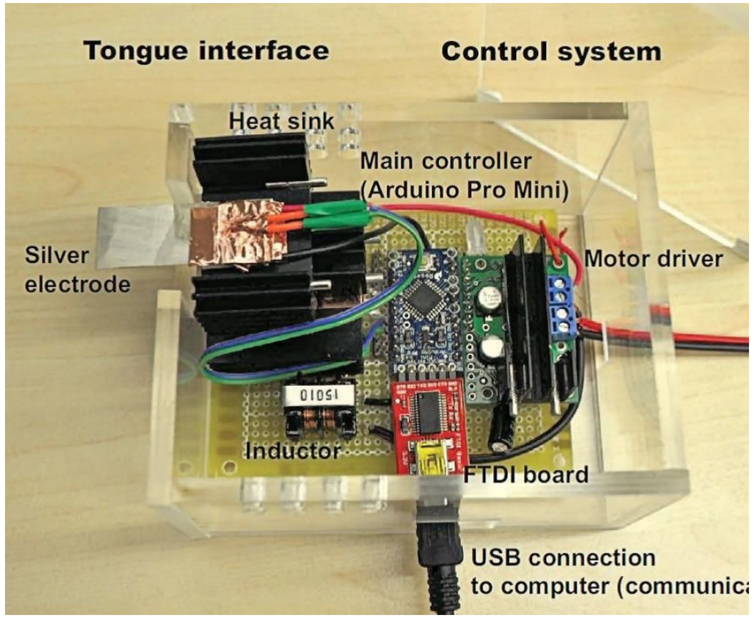

AR finally reaching enterprise / verticals

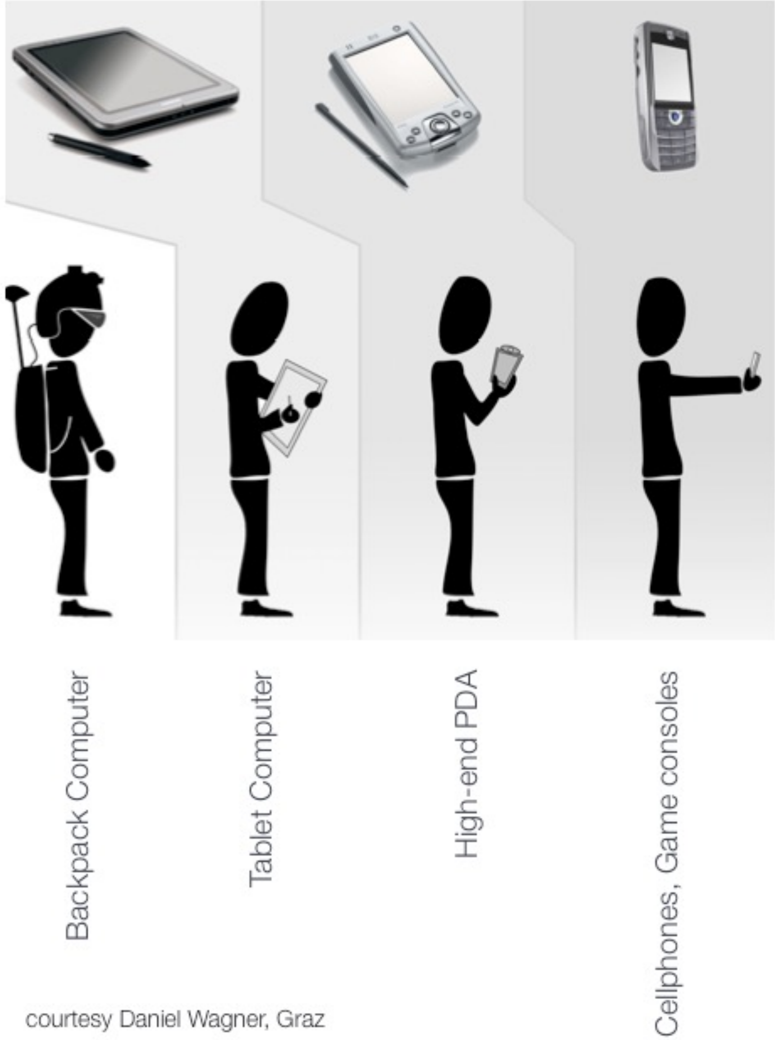

courtesy Daniel Wagner, Graz

Rendering and Display

- Leverage modern GPUs
  - Render millions of polygons at 60Hz or higher (currently 140Hz or higher)
  - Need to render both eyes with low latency
- Rendering is “solved”(?) and getting better
  - PC-attached VR leverages modern GPUs
  - Stand-alone VR limited to less-powerful mobile GPUS, and battery constraints
- Displays and Tracking/sensing are still far from “solved”

Typical Modern AR/VR Displays

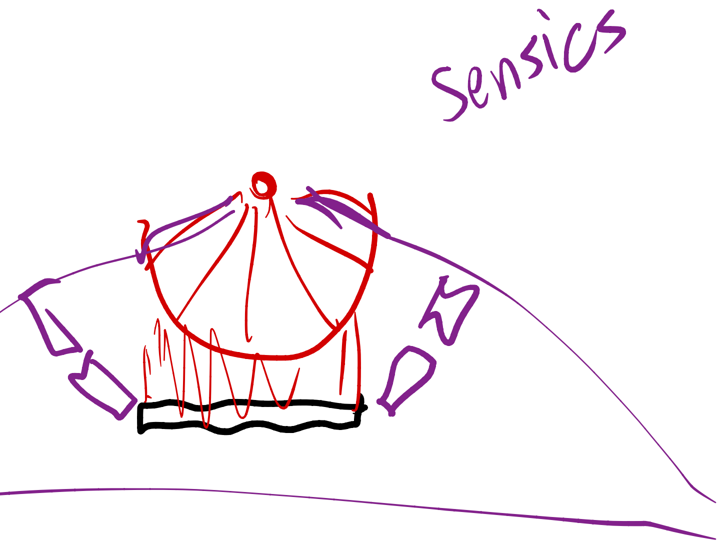

- Deliver separate images to each eye
  - Some delay between user motion and image being seen
  - Small lag causes motion sickness
- Use small displays (or lasers), either via mirrors and folded optics (AR) or directly viewed with distorted optics (VR)
  - AR field of view << VR field of view
- Fixed focal distances
  - Major problem that needs to be solved. Lightfield displays are promising idea (many focal depths)
  - Magic Leap tried to create Lightfield displays but has not been able to deliver

Fresnel Lenses in VR

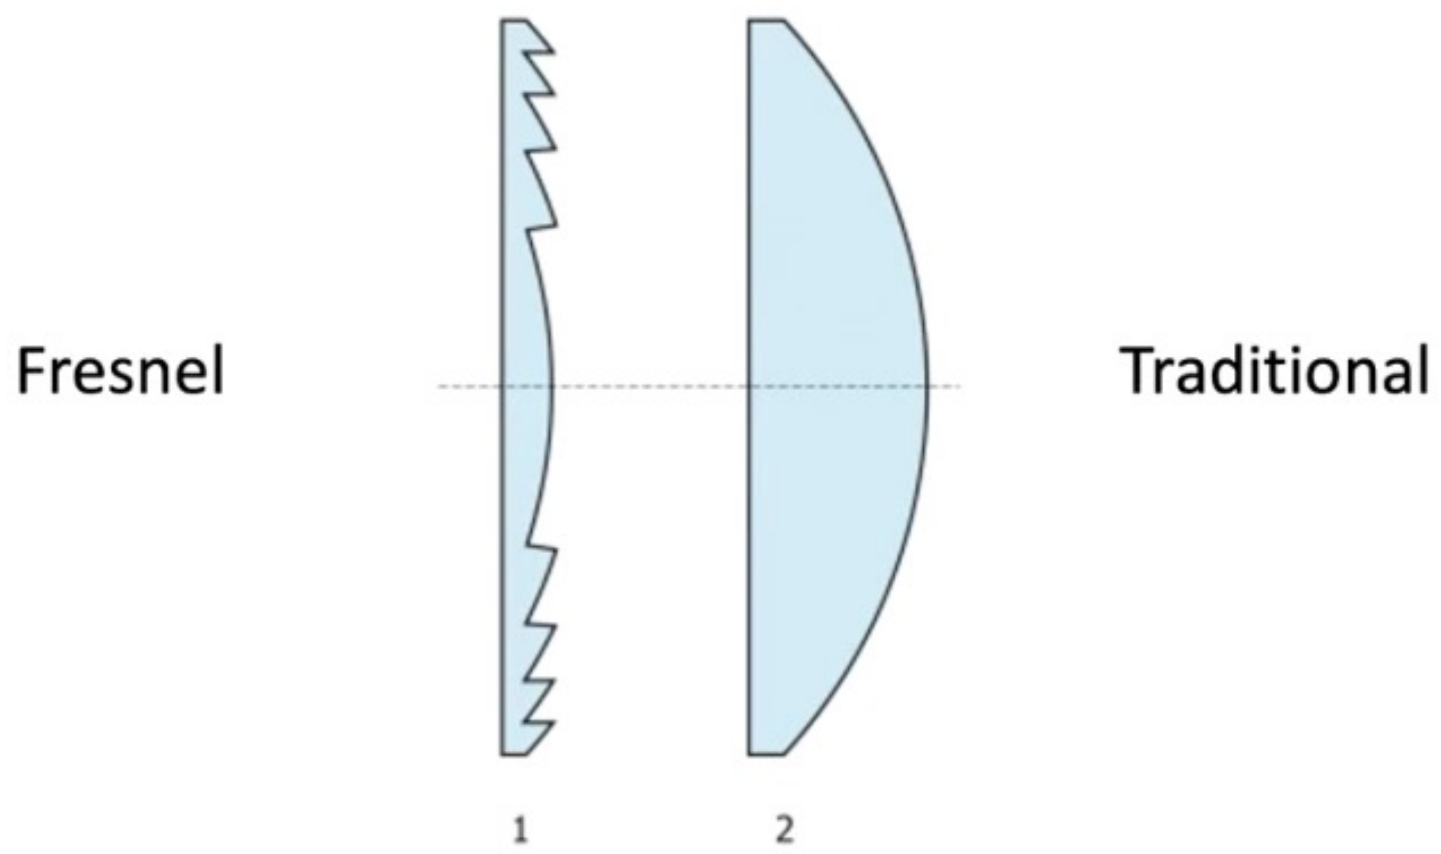

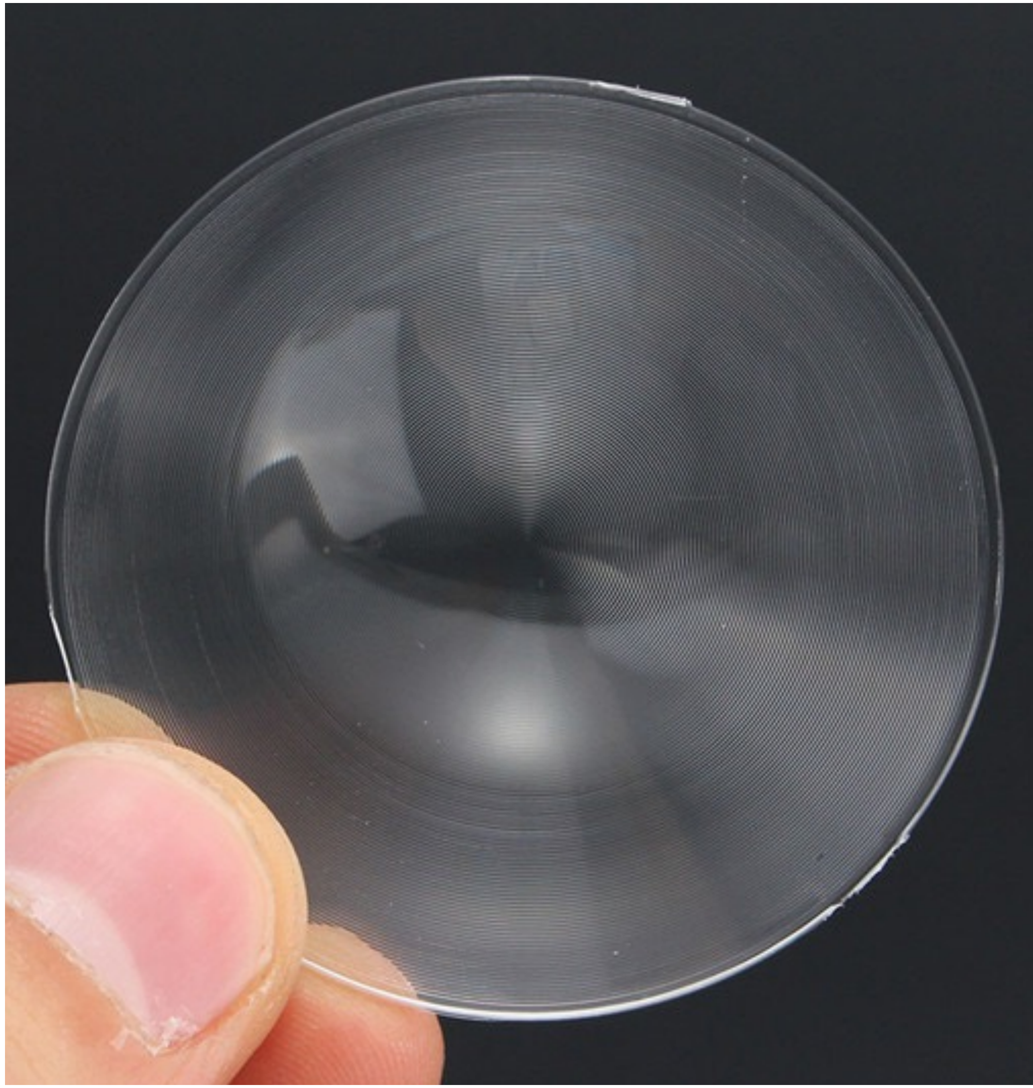

Waveguide Displays

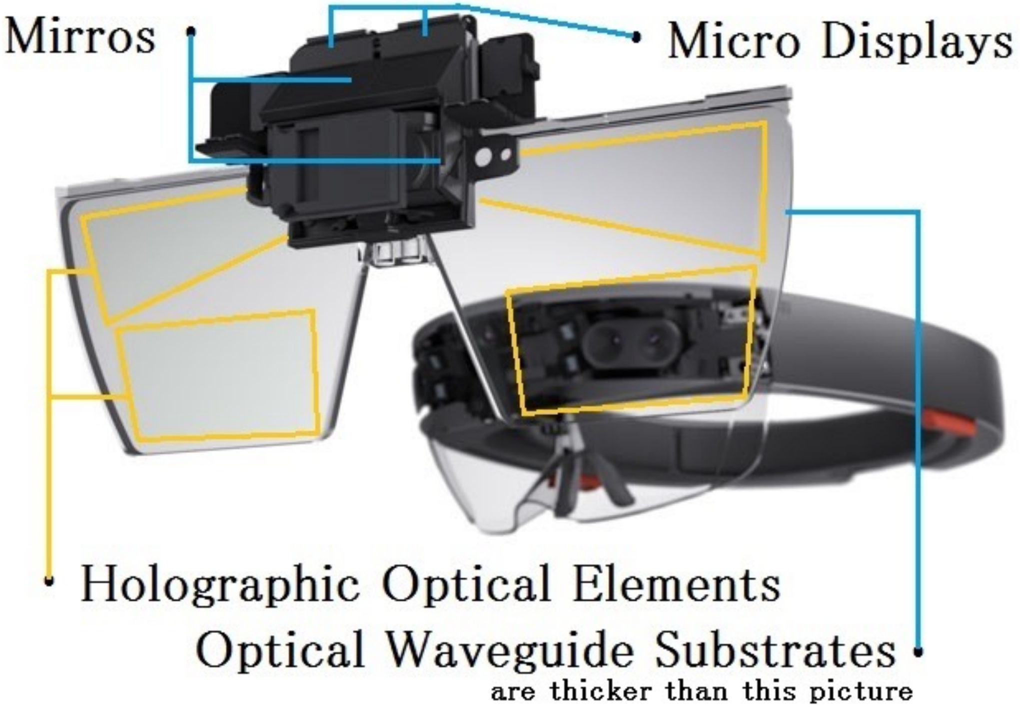

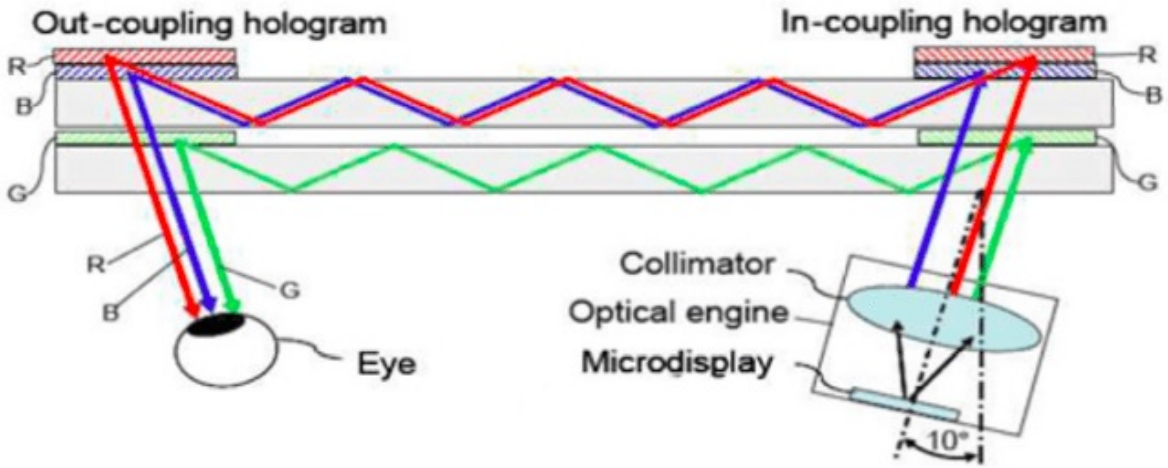

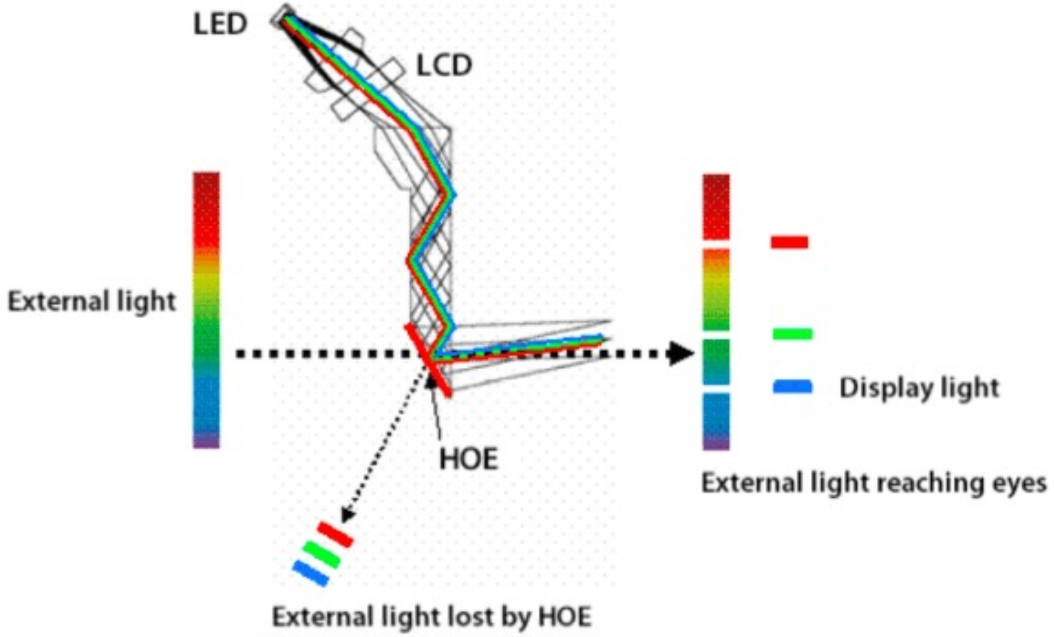

Tracking and Sensing

- Determine "pose" (3D position + 3D orientation) of head, hands, …
  - Ideally as fast as the display with low end-to-end latency
- Position virtual "cameras" at the eyes of viewer

- Kinds
  - Closed loop sensing. E.g., magnetic, ultrasonic, etc
  - Outside-in. E.g., Fixed cameras observe passive targets in mocap systems.
  - Inside-out. E.g., cameras on user/display sense lights/objects in world

Modern Tracking (Outside In): Oculus Rift, Playstation VR

Modern Tracking (Inside Out): HTC Vive Lighthouse

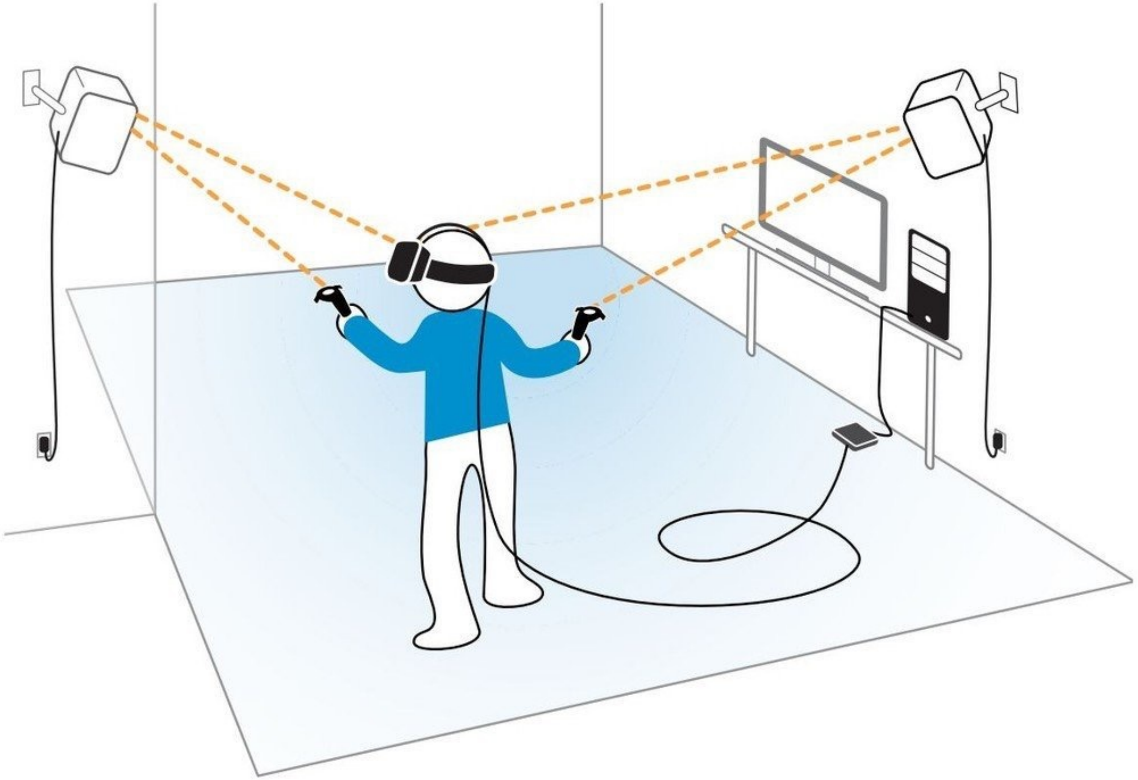

Modern Tracking (Inside Out): Camera-based SLAM and VIO

- SLAM: Simultaneous Localization and Mapping
- VIO: Visual Inertial Odometry
- Both used on Phones and Mobile AR/VR systems
  - ARKit on iOS, ARCore on Android
  - Hololens1 and 2, Magic Leap 1
  - Quest 1 and 2, Windows MR, all other stand alone displays
- Nothing installed in room: finds features and builds a map (Mapping) while tracking device relative to map (Localization)

https://youtu.be/2jY3B_F3GZk?t=33

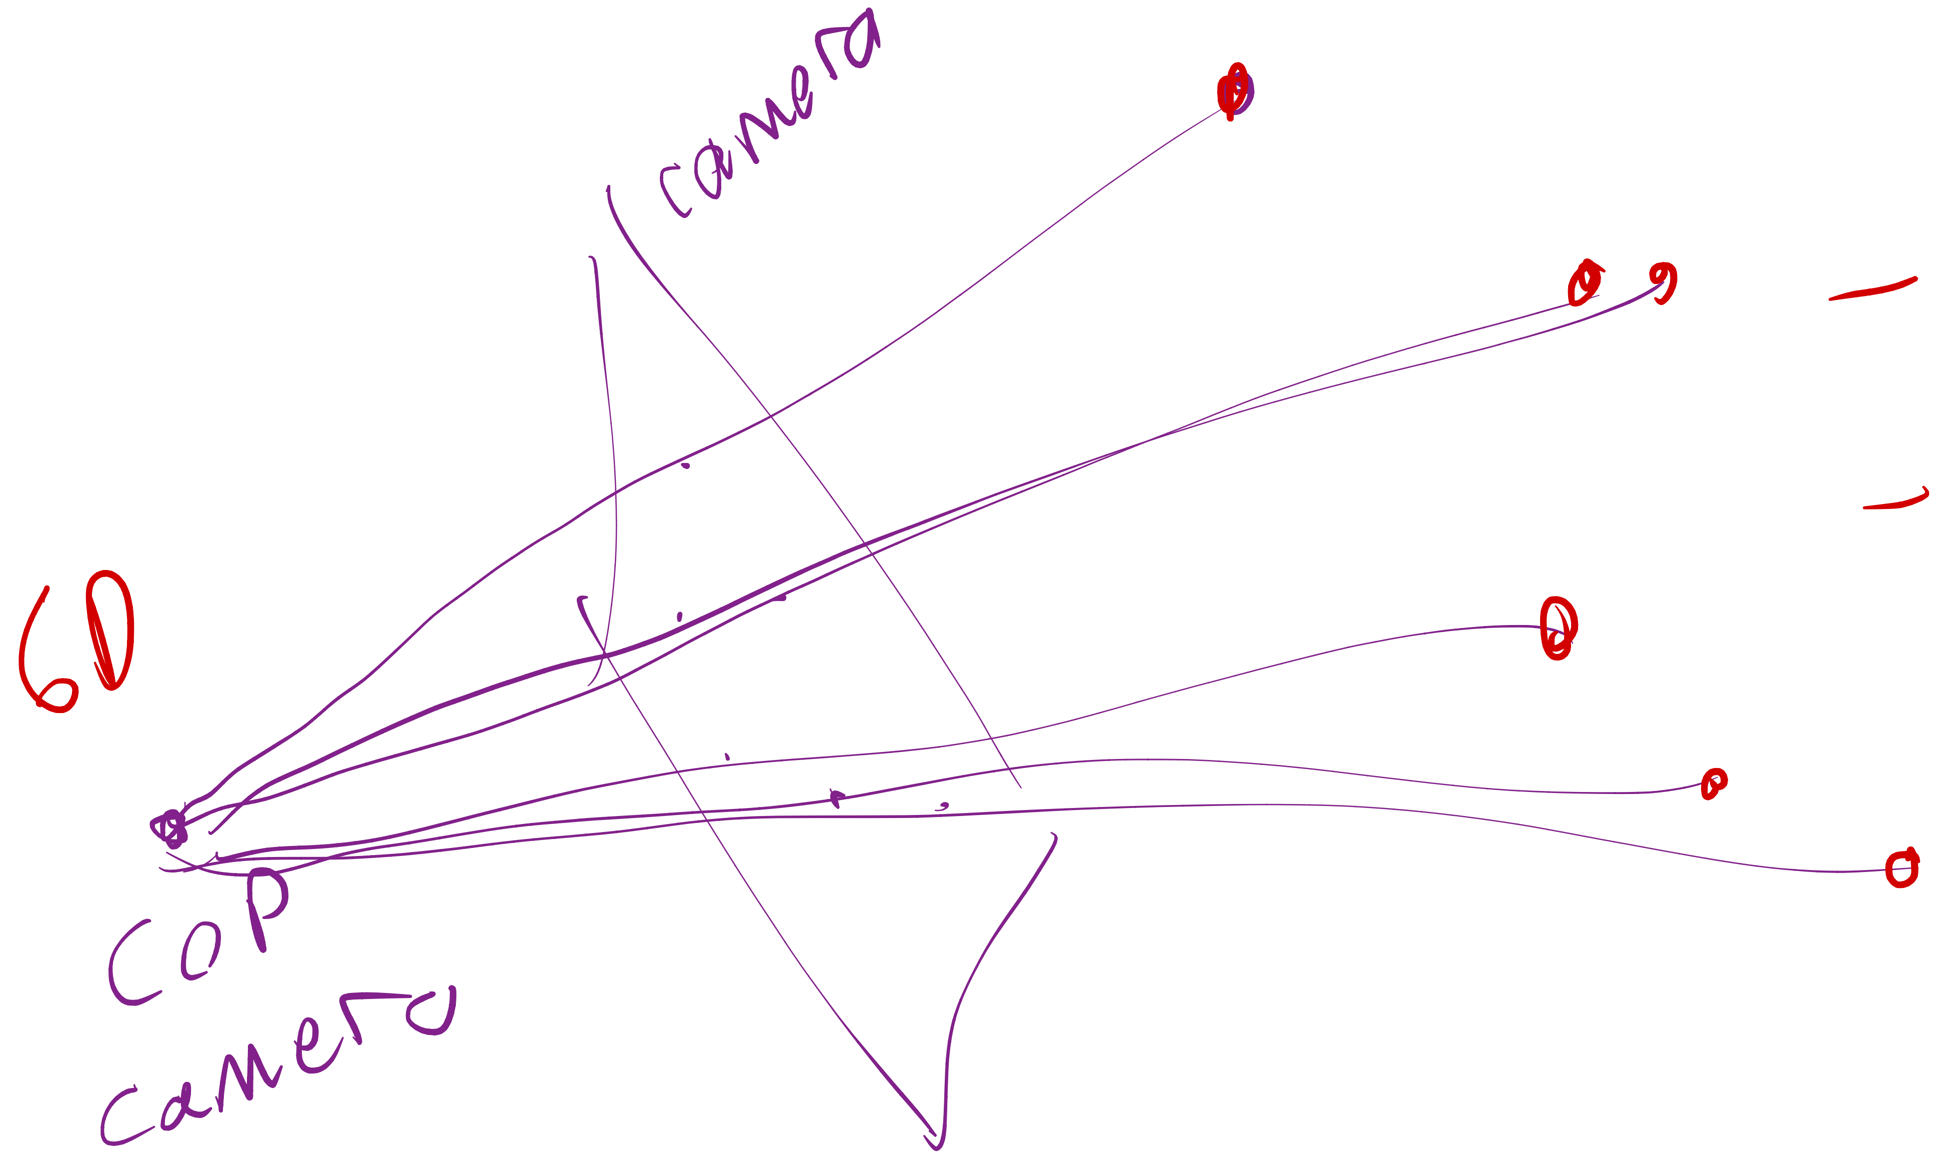

- more points
- wide spread

Core System Loop for VR/AR

- Once: Setup system to request features
- Each frame:
  - App gets state, renders frame, submits to system
  - Compositor renders to device
- State synchronized to a single time
  - Display intrinsics, controller poses, controller buttons and other peripheral state
- Render to 3D pre-distorted canvas using inputs
- Submit to VR runtime, compositor takes care of final rendering

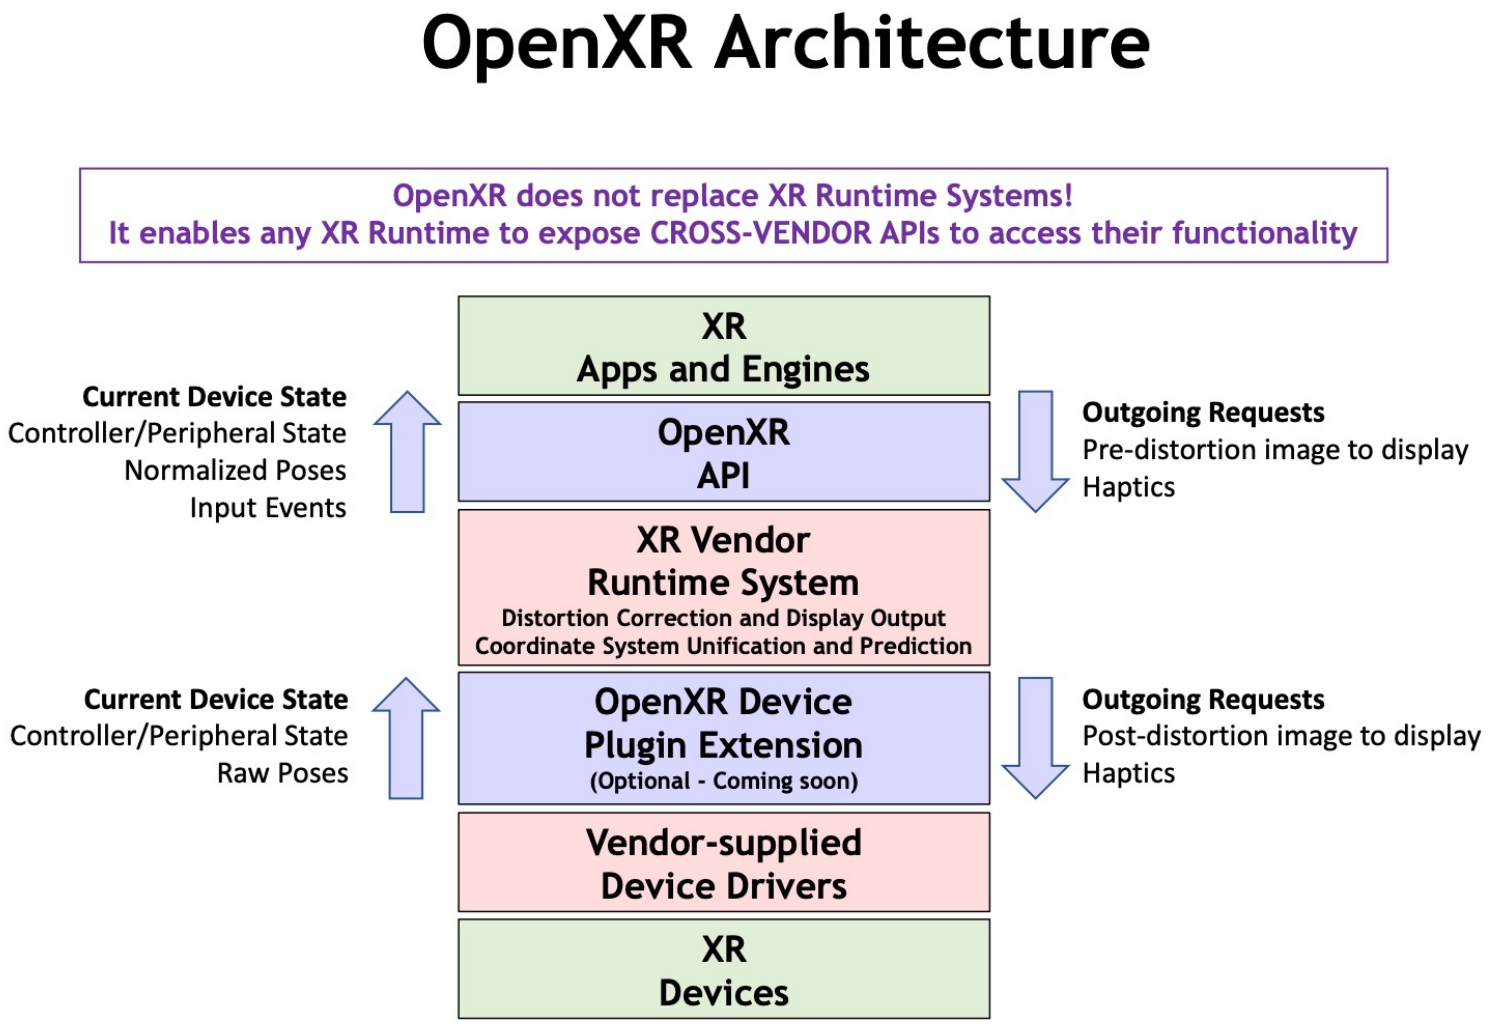

VR Runtime Compositor

- Two jobs: distort image to display geometry, handle latency
- Second job used to be ignored, or dealt with via prediction
  - Post rendering warping is key to modern VR/AR device usability
  - Timewarp, Spacewarp, Reprojection, Motion Smoothing, Asynchronous SpaceWarp
    - https://uploadvr.com/reprojection-explained/
    - https://uploadvr.com/asw-2-rift-released/

- Ideas have been around for decades, but Carmack / Oculus finally made it work in practice in 2013/2014
  - (https://web.archive.org/web/20140719085135/http://altdev.co/2013/02/22/latency-mitigation-strategies/)

Timewarp

- Reproject rendered frame to account for change in head rotation from when tracking information given to application

- Only uses orientation

Asynchronous Timewarp (ATM)

- Use Timewarp to deal with lost frames
- Keep frame around, asynchronously Timewarp old frame if a frame dropped

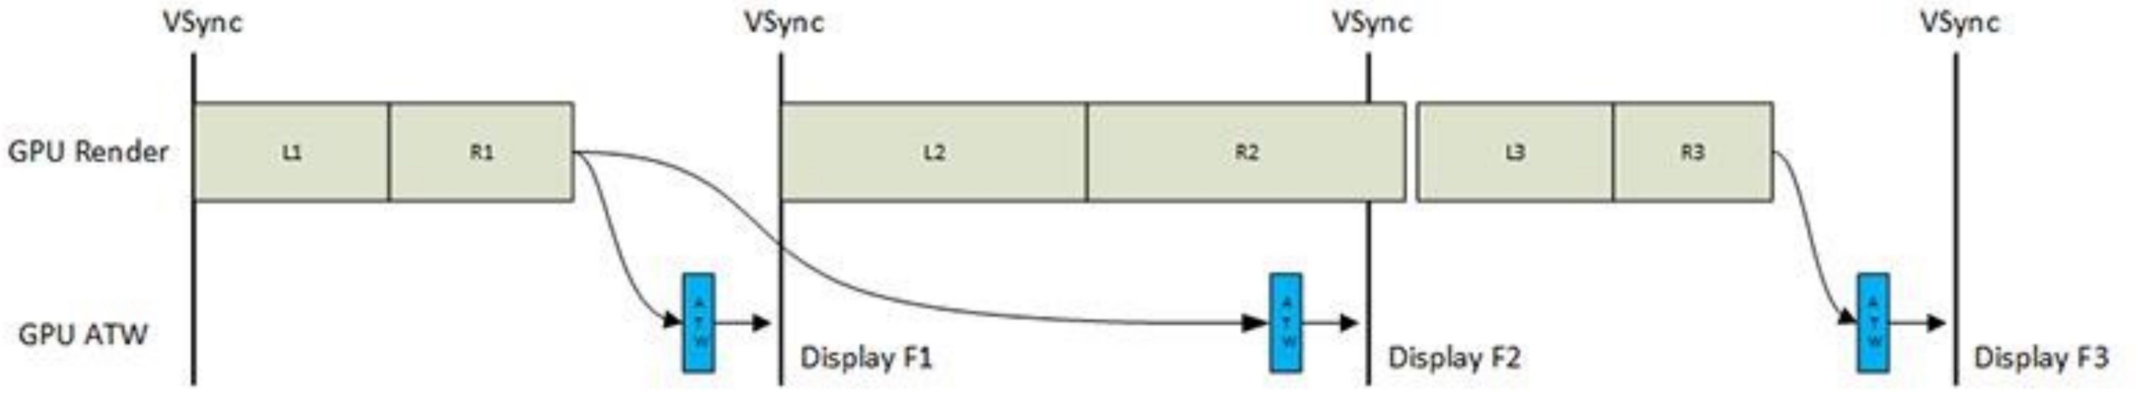

- Called Asynchronous Reprojection in the Steam SDK

Asynchronous Spacewarp

- Tries to account for motion as well as rotation by extrapolating image changes based on motion in previous frames

- ASW 2.0 incorporates depth information for better results

What is AR/VR Good For?

- No Killer App yet

AR/VR Areas: Education

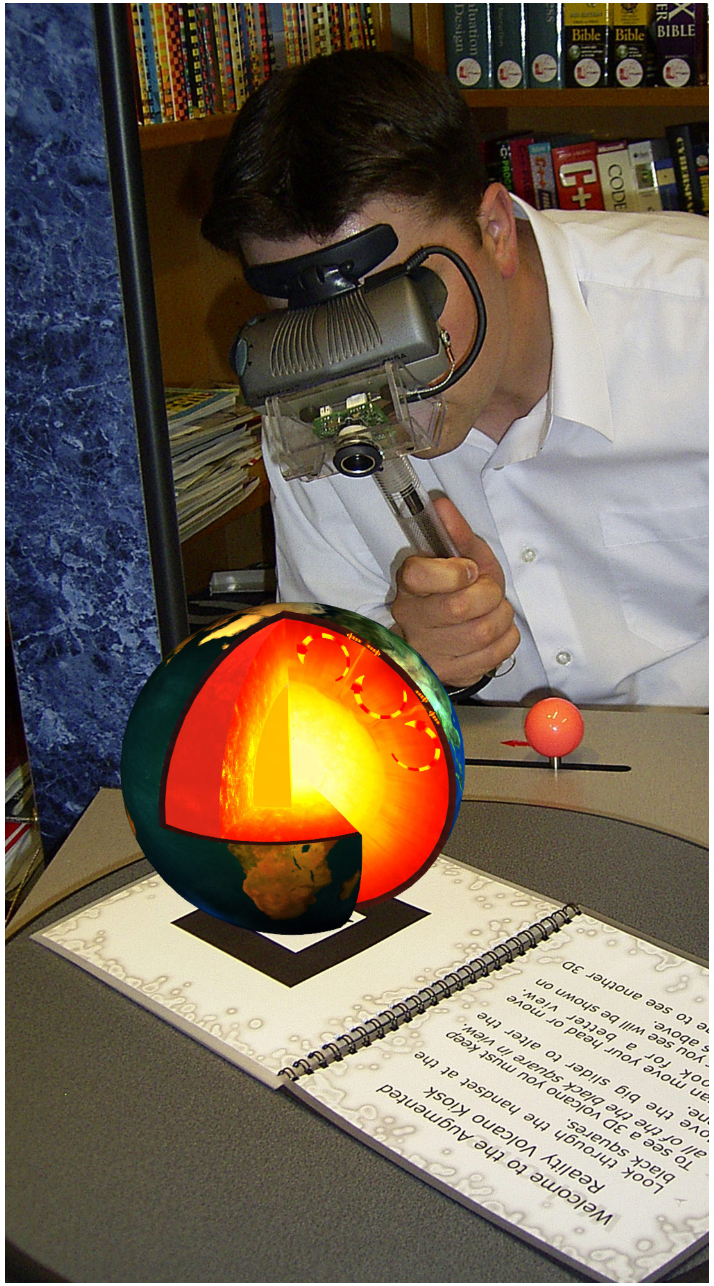

Image courtesy of Mark Billinghurst. Copyright University of Canterbury.

AR/VR Areas: Training

- Both in VR and AR, and hybrid
  - Experience more of real environment
  - ie. UAV Takeoff, Landing
- Integrate with simulation (ie. MODSAF)

AR/VR Areas: Tours of Historic Sites

- ARCHEOGUIDE (2002): use AR to improve on audio or text-based tour guides
- VR 360 images/video, immersive recreations

Mock-up of a mobile AR tour. Courtest IEEE Spectrum.

ARCHEOGUIDE.

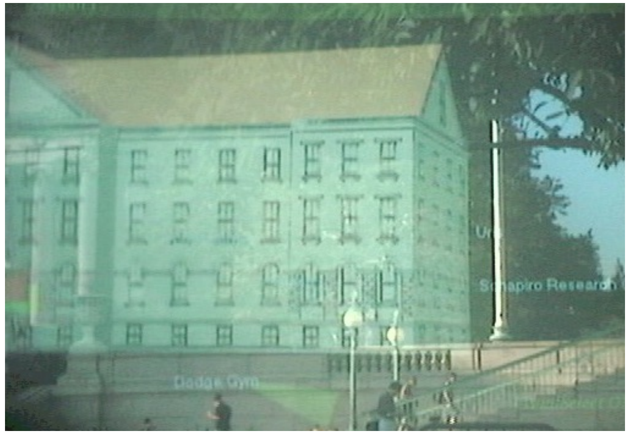

Recreation of a building at Columbia University.

AR/VR Areas: Games and Entertainment

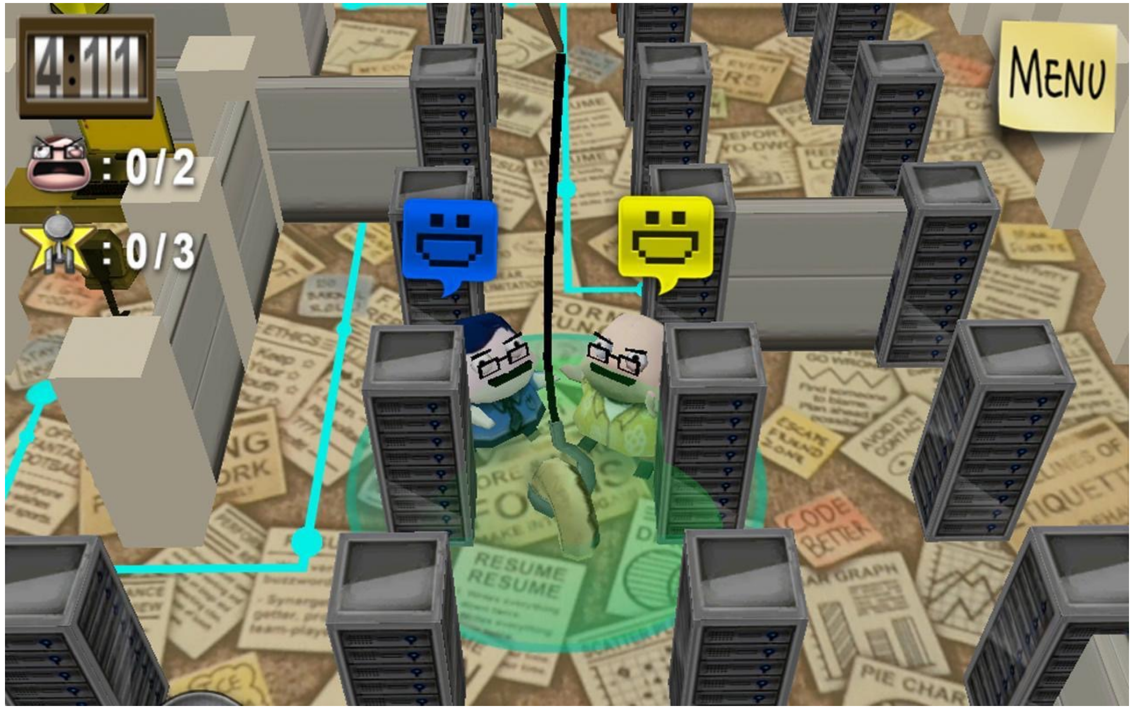

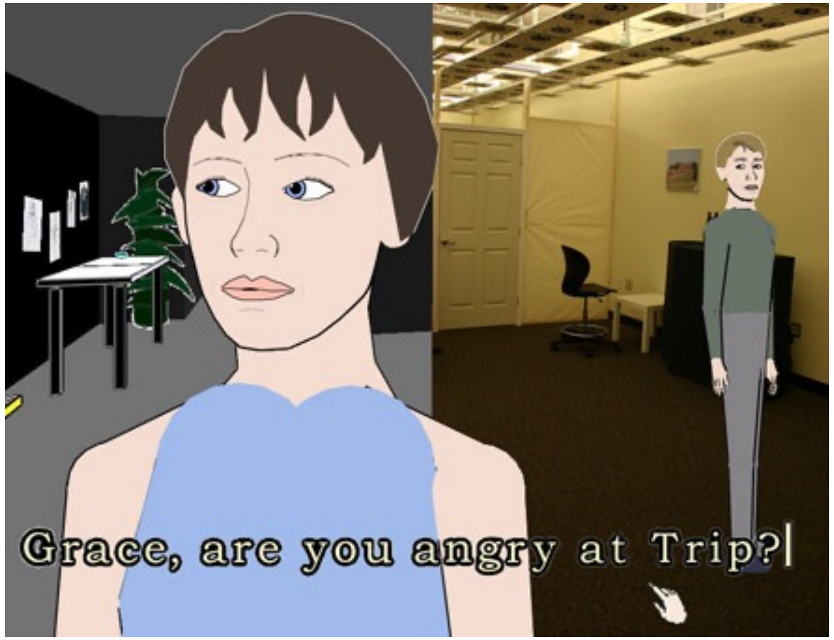

Courtesy MR Systems Lab

AR Quake. Copyright University of South Australia.

AR/VR Areas: Fitness

AR/VR Areas: Remote Work and Collaboration

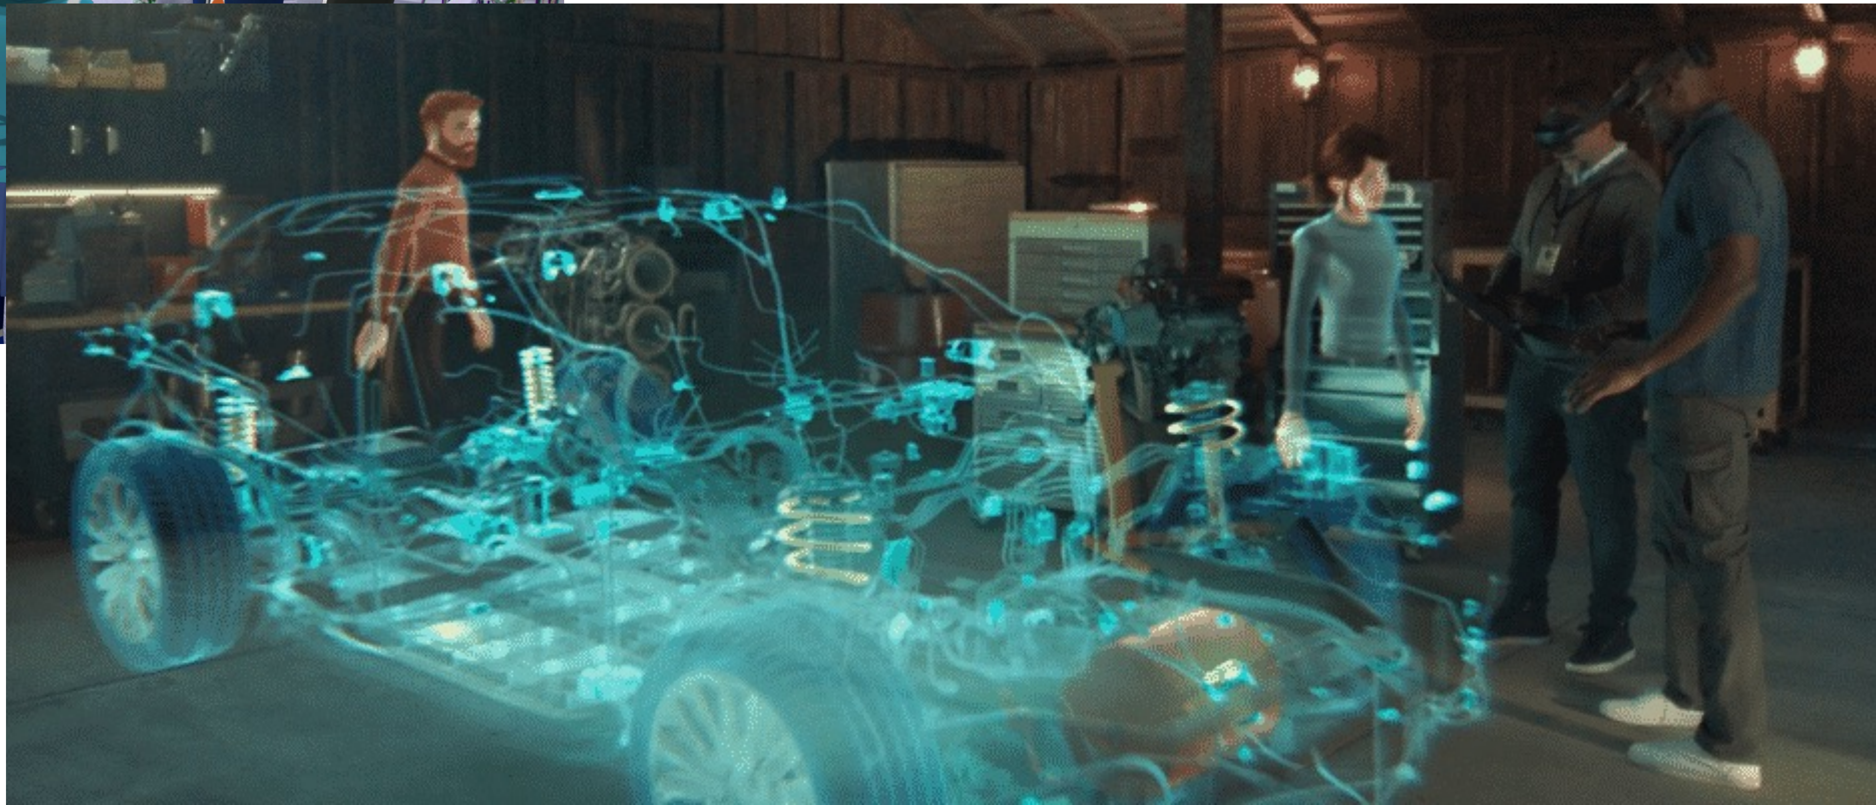

AR/VR Areas: X-Ray Views

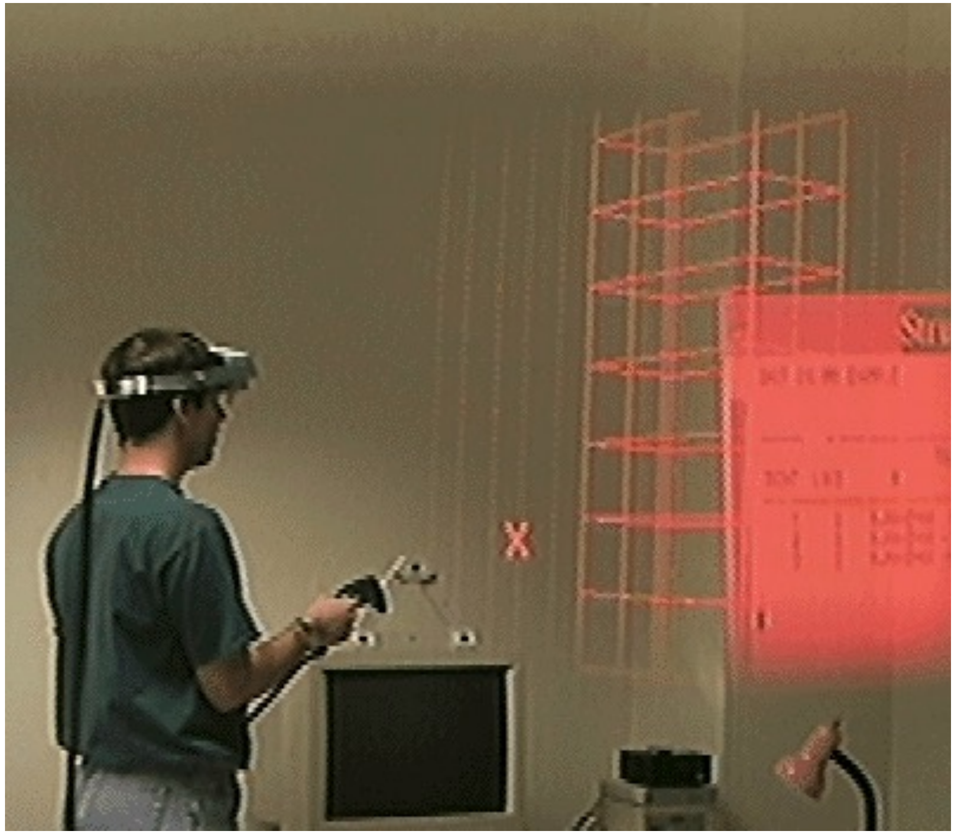

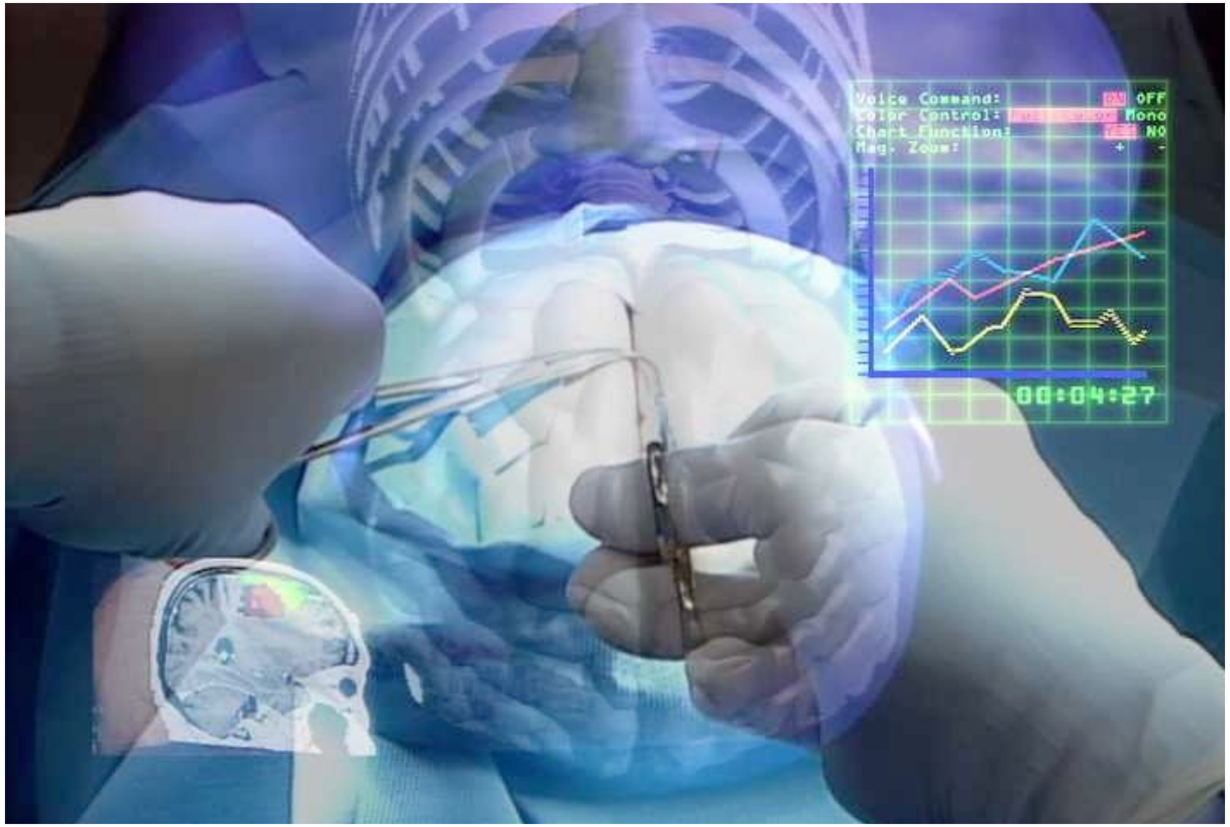

- Others: views through hills, buildings

AR/VR Areas: Enhanced Views in Enterprise Apps

- Merging virtual information and sensor data onto user view
  - E.g., aircraft HUD
- For both visible and invisible entities

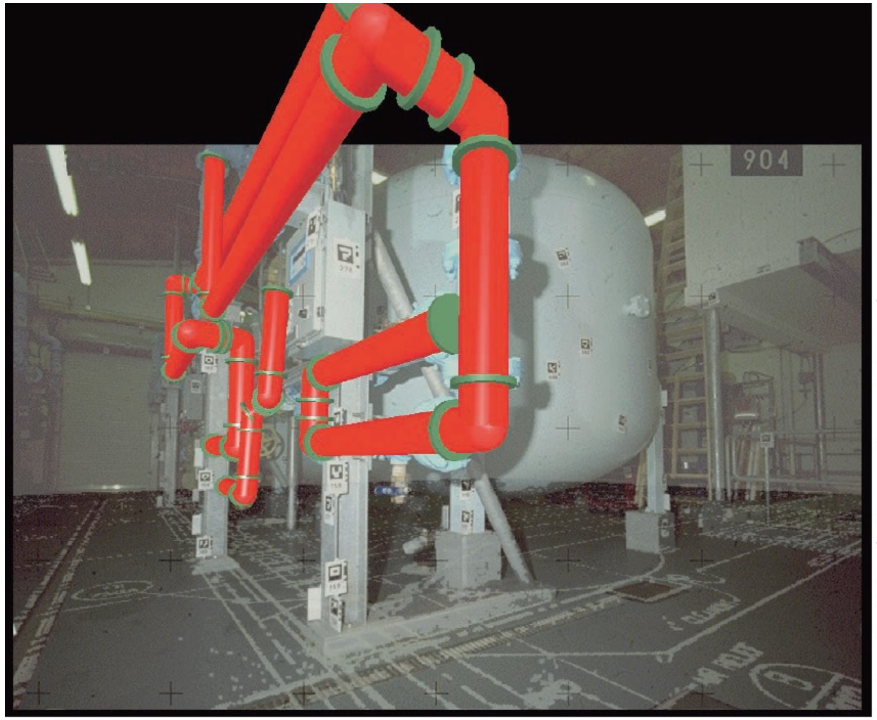

Courtesy Nassir Navab, Siemens Corporate Research

AR/VR Areas: Enhanced Views in Enterprise Apps

- Merging virtual information and sensor data onto user view
  - E.g., aircraft HUD
- For both visible and invisible entities

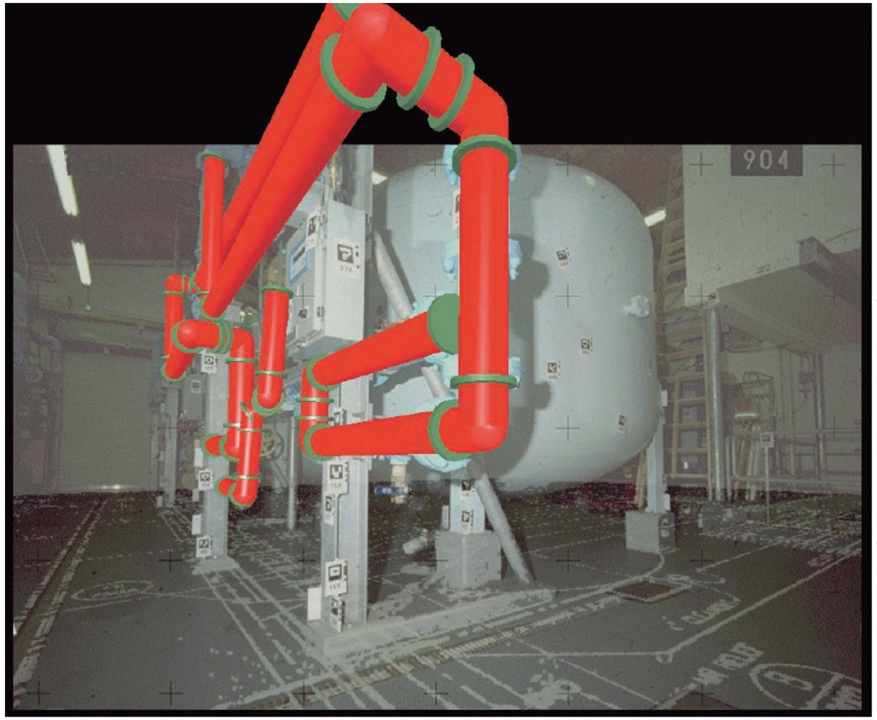

Courtesy Nassir Navab, Siemens Corporate Research