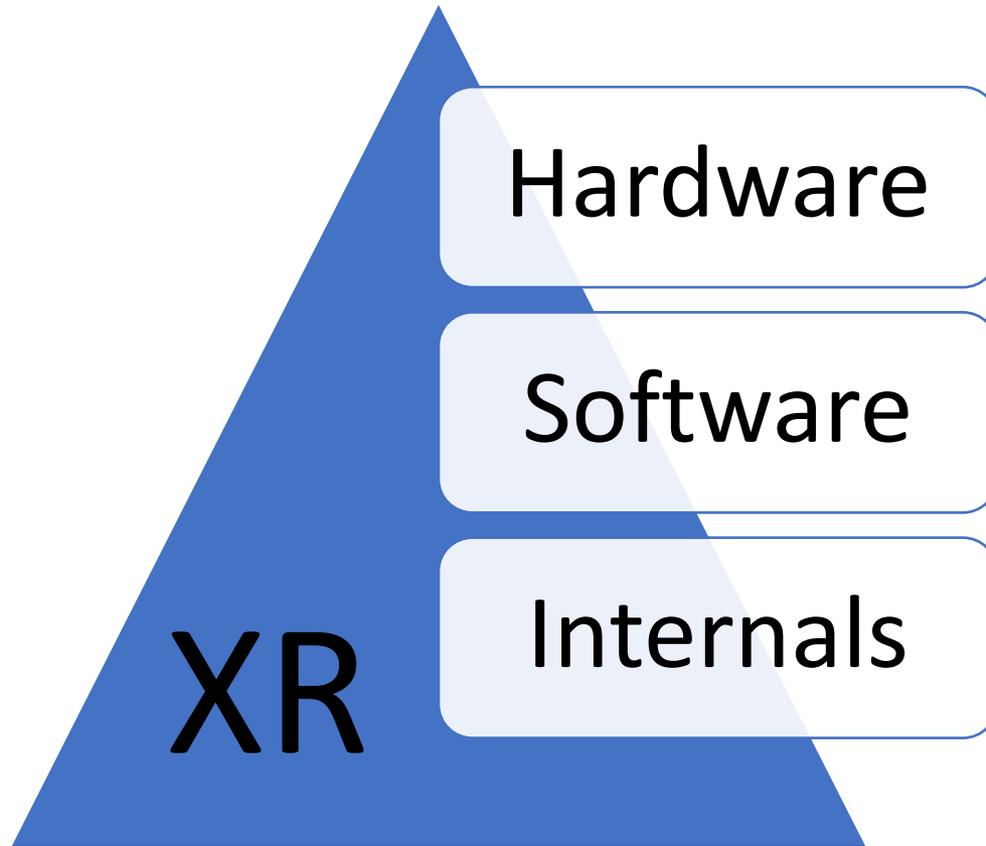


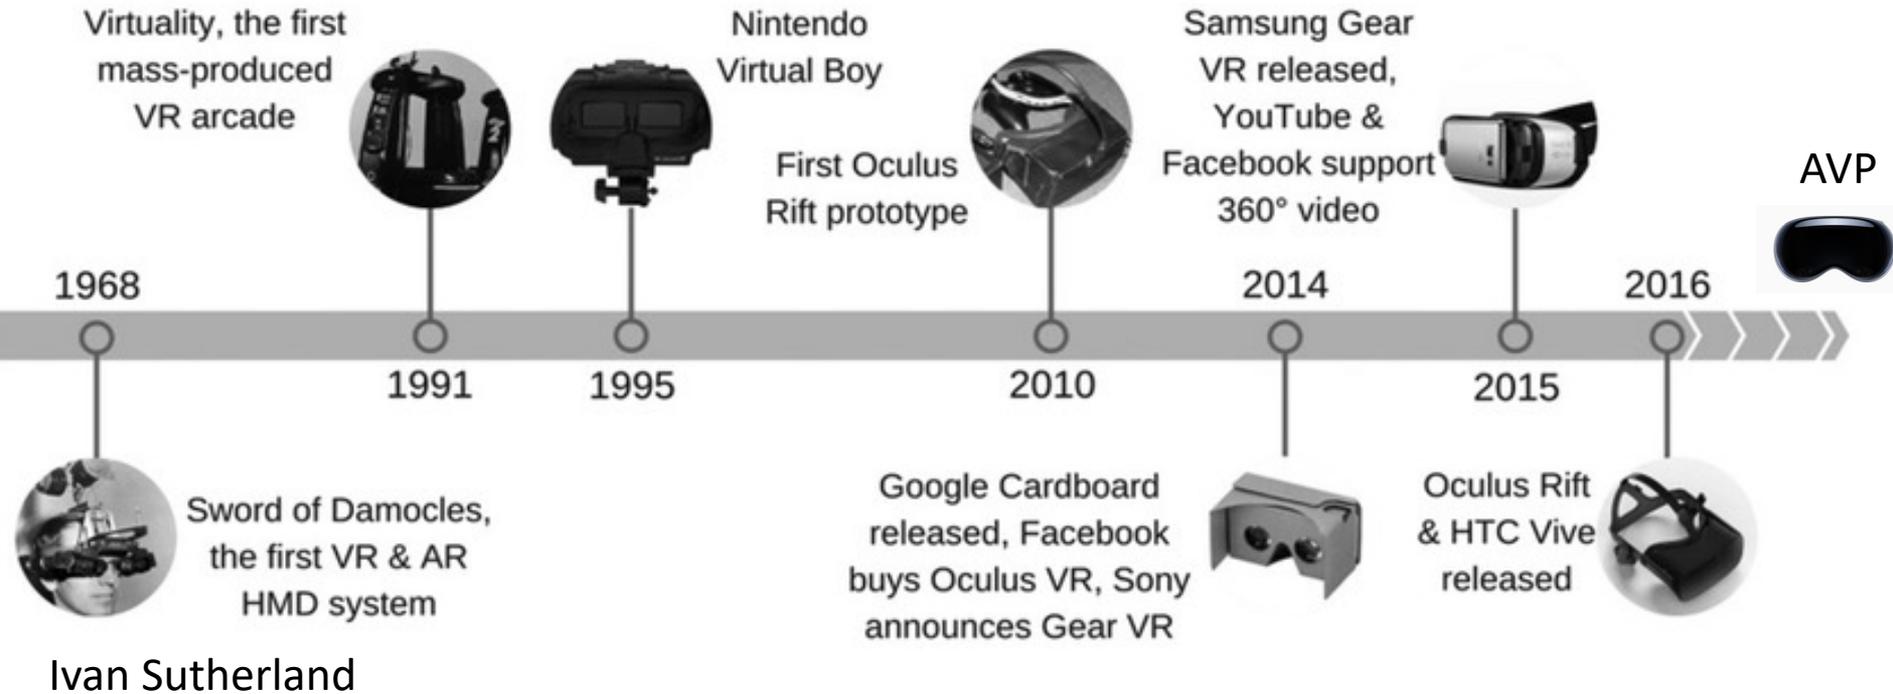
EECE5698

Networked XR Systems

Lecture Outline for Today



Timeline of XR Headsets



Popular XR Headsets Today



Oculus Headset Series - Rift

Display Type & Size	Dual low-persistence AMOLED (PenTile subpixel matrix)
Display Size	TBA
Resolution	1200 x 1080 (per eye)
Refresh Rate	90Hz
Field of View	~100-degrees
Lens Type	Hybrid Fresnel
Lens Adjustment	IPD (58-72 mm), lens-to-eye distance (adjustable with optional glasses spacer)
Sensors	Accelerometer, gyroscope, magnetometer
Tracking Technology	6 DOF Constellation camera optical 360-degree IR LED tracking
Integrated Camera	No
Audio	Microphone, integrated supra-aural 3D spatial audio headphones (removable)
HMD Ports	Proprietary headset connector (HDMI/USB 3.0)
HMD Cable Length	4 m
Materials Used	Plastic, IR-transparent fabric, glass, foam rubber
Dimensions	~171 (~216) × ~102 mm(W (width including headphones) × D)
Weight	470g (excluding cable)



Tethered to a PC

Oculus Headset Series - Rift

Integrated constellation tracking: IR LEDs
under fabric shell



Oculus Headset Series – Quest1

● SoC	Qualcomm Snapdragon 835
● Display	Dual 1440x1600 72Hz OLED panels
IPD Setting	Mechanical IPD adjustment (range undisclosed)
Storage	64GB or 128GB of internal flash storage
Audio	Integrated speakers and microphone, dual 3.5 mm audio jack (one on each side), in-ear headphone accessory available
RAM	4GB
Battery	Built-in Lithium Ion battery (mAh undisclosed)
Facial Interface and Strap Material	Knit Mesh, Nylon Micro Yarn, Spandex Materials
● Tracking Technology	Oculus Insight inside-out camera-based 6-DoF tracking with motion controllers
Input	2nd-generation Oculus Touch controllers
Play Space Requirements	Stationary or Room-scale. Room-scale requires a minimum of 2 x 2m or 6.5 x 6.5 feet of obstruction-free floor space
Dimensions	193 x 105 x 222mm
● Weight	571g



Standalone – No need of PC

Oculus Headset Series – Quest2

● Display	Fast-switch LCD: 1832 x 1920 resolution per eye, 72 Hz or 90 Hz refresh rate
● IPD Setting	3 mechanical pre-sets (58mm, 633mm, 68mm)
Storage	64GB or 256GB of internal flash storage
Audio	Integrated speakers and microphone, single 3.5 mm audio jack, third-party accessories available
RAM	6GB
Battery	Built-in Lithium Ion battery (mAh undisclosed); 2-3 hours estimated runtime, 2.5 hour charge time
Facial Interface and Strap Material	Knit Mesh foam cushion, flexible fabric head strap
Tracking Technology	Oculus Insight inside-out camera-based 6-DoF tracking with motion controllers
Input	3rd-generation Oculus Touch controllers
Play Space Requirements	Stationary or room-scale; Room-scale requires a minimum of 6.5 x 6.5 feet (2m x 2m) of obstruction-free floor space
Dimensions	7.5 x 4 x 5.6 inches (191.5 x 102 x 142.5mm)
● Weight	1.1 pounds (503g)



Nothing changed significantly

Oculus Headset Series – Quest3

● Display	2064 x 2208 per eye
Display Type	LCD
Refresh Rate	2Hz, 80Hz, 90Hz, 120Hz (experimental)
● Processor	Qualcomm Snapdragon XR2 Gen 2
RAM	8GB
Storage	128GB or 512GB
● Field of View	110 degrees horizontal, 96 degrees vertical
Degrees of Freedom	6 DoF
Audio	Dual open-air speakers
● Wireless Connectivity	Wi-Fi 6E, Bluetooth 5.2
● Battery Life	Up to 2.9 hours
● Weight	1.13 pounds (515 grams)



Passthrough
Mixed Reality Headset

Meta Glasses

- Hardware: Snapdragon AR1 Gen 1 platform
- Live Streaming
- AI Voice Assistant
- Snap pics, record video, listen to music, etc



<https://www.youtube.com/watch?v=gvpybQpB46k&t=1s>

Magic Leap & HoloLens

Optical see-through, eye tracking



	Magic Leap 2	HoloLens 2
Resolution	1440 x 1760 per eye	1440 x 936 per eye
Optics	Waveguide	Waveguide
Refresh Rate	120 Hz	60 Hz
Field of View	70° diagonal	52° diagonal



No controllers, hand tracking

Apple Vision Pro

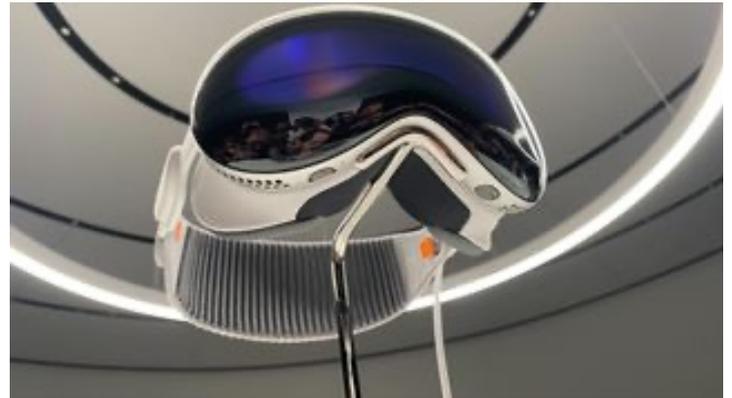
Two processors:

M1 – for general purpose

R1 – for sensing, IMU, cameras

4K micro-LED panel for each eye

Spatial videos

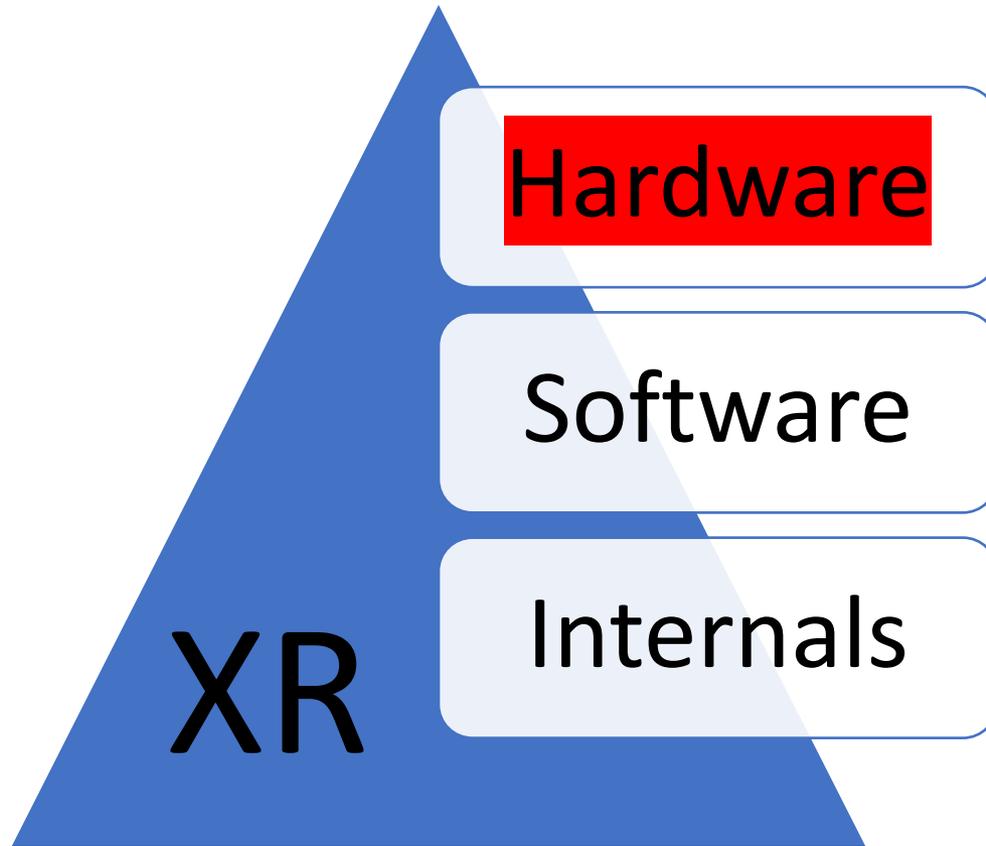


XReal

● Display	1920 x 1080 per eye
Display Type	OLED
Brightness	500 nits
● Field of View	46 degrees
● Degrees of Freedom	3 DoF
Audio	Dual open-air speakers
Connectivity	USB-C
● Weight	2.54 ounces (72 grams)



Lecture Outline for Today



XR Software Tools

- Web Programming
- Standards
- SDKs
- Native Renderers
- 3D Modeling
- Game Engines
- 3D Scanners

XR Web Tools

- WebGL

WebGL is a JavaScript API for rendering interactive 2D and 3D graphics within any compatible web browser without the use of plug-ins. WebGL is fully integrated with other web standards, allowing GPU-accelerated usage of physics, image processing, and effects in the HTML canvas.

<https://get.webgl.org/>

XR Web Tools

- Three.js

JavaScript 3D library

The aim of the project is to create an easy-to-use, lightweight, cross-browser, general-purpose 3D library. The current builds only include a WebGL renderer but WebGPU (experimental), SVG and CSS3D renderers are also available as addons.

```
import * as THREE from 'three';

const width = window.innerWidth, height = window.innerHeight;

// init

const camera = new THREE.PerspectiveCamera( 70, width / height, 0.01, 10 );
camera.position.z = 1;

const scene = new THREE.Scene();

const geometry = new THREE.BoxGeometry( 0.2, 0.2, 0.2 );
const material = new THREE.MeshNormalMaterial();
```

<https://github.com/mrdoob/three.js/>

XR Web Tools

- A-Frame

A-Frame is a web framework for building virtual reality (VR) experiences. A-Frame is based on top of HTML, making it simple to get started. But A-Frame is not just a 3D scene graph or a markup language; the core is a powerful entity-component framework that provides a declarative, extensible, and composable structure to [three.js](#).

<https://glitch.com/~aframe>

XR Standards

- **WebXR**

The WebXR Device API provides access to input (pose information from headset and controllers) and output (hardware display) capabilities commonly associated with Virtual Reality (VR) and Augmented Reality (AR) devices. It allows you develop and host VR and AR experiences on the web.

Target hardware

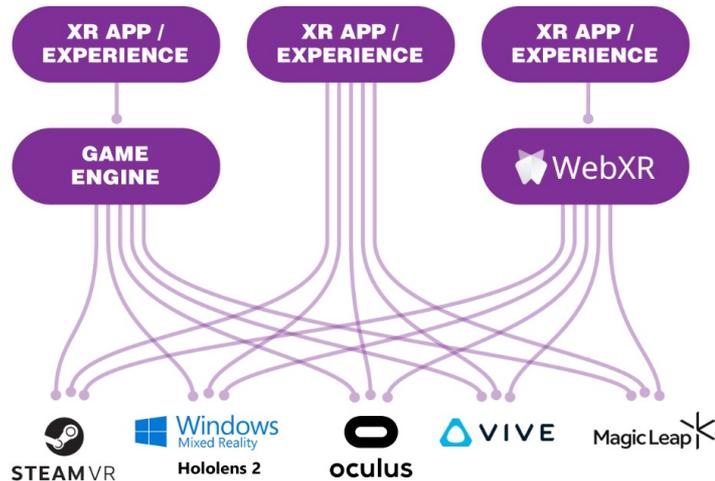
Examples of supported devices include (but are not limited to):

- [ARCore-compatible devices](#)
- [Google Daydream](#)
- [HTC Vive](#)
- [Magic Leap One](#)
- [Microsoft HoloLens](#)
- [Meta Quest 1, 2, and Pro](#)
- [Samsung Gear VR](#)
- [Windows Mixed Reality headsets](#)

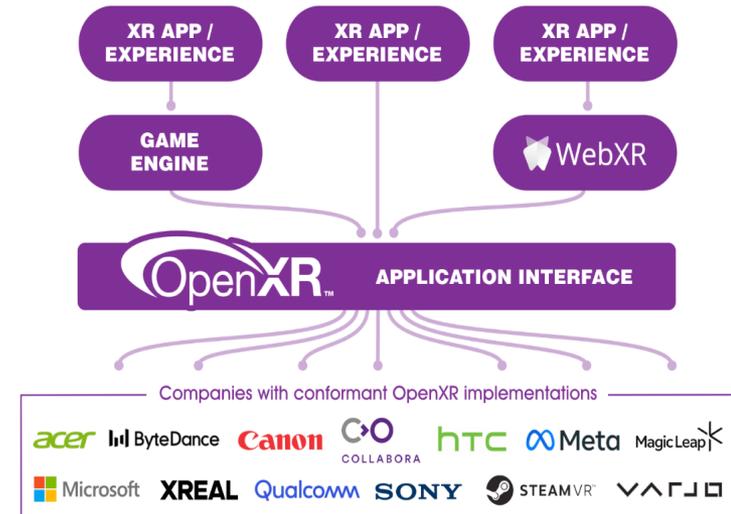
XR Standards

- OpenXR (<https://www.khronos.org/openxr/>)

OpenXR provides cross-platform, high-performance access directly into diverse XR device runtimes across multiple platforms. OpenXR enables applications and engines, including WebXR, to run on any system that exposes the OpenXR APIs.



Before OpenXR: Applications and engines needed separate proprietary code for each device on the market.

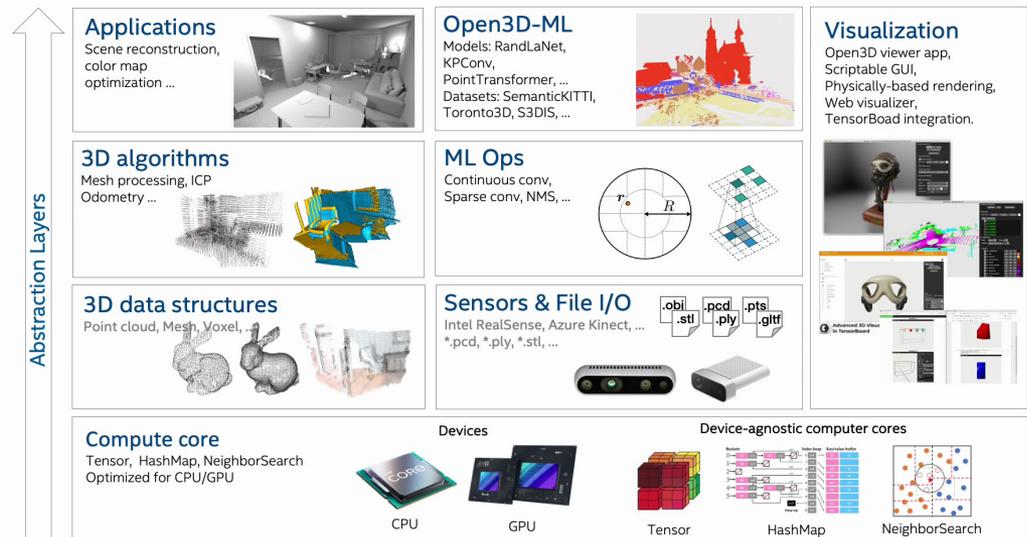


OpenXR provides a single cross-platform, high-performance API between applications and all conformant devices.

XR SDKs

- Open3D

Open3D is an open-source library that supports rapid development of software that deals with 3D data. The Open3D frontend exposes a set of carefully selected data structures and algorithms in both C++ and Python. The backend is highly optimized and is set up for parallelization.



<https://www.open3d.org/docs/release/tutorial/geometry/mesh.html>

XR SDKs

- ARKit – Mobile (Apple)
 - ARKit combines device motion tracking, world tracking, scene understanding, and display conveniences to simplify building an AR experience.

```
let session = ARKitSession()  
let worldInfo = WorldTrackingProvider()
```



XR SDKs

- ARCore – Mobile (Android)

ARCore is Google's augmented reality SDK offering cross-platform APIs to build new immersive experiences on Android, iOS, Unity, and Web.



Built-in sensors

GPS for position and compass for orientation.

Cloud Anchors API

Create a map of an area for other users to localize against.

Geospatial API

Leverage Google's global-scale 3D map as your canvas.

XR SDKs

- Oculus & MRTK



Input System



Hand Tracking
(HoloLens 2)



Eye Tracking
(HoloLens 2)



Profiles



Hand Tracking
(Ultraleap)



UI Controls



Solvers



Multi-Scene
Manager



Spatial
Awareness



Diagnostic
Tool



MRTK Standard Shader
Example View



Speech
& Dictation



Boundary
System



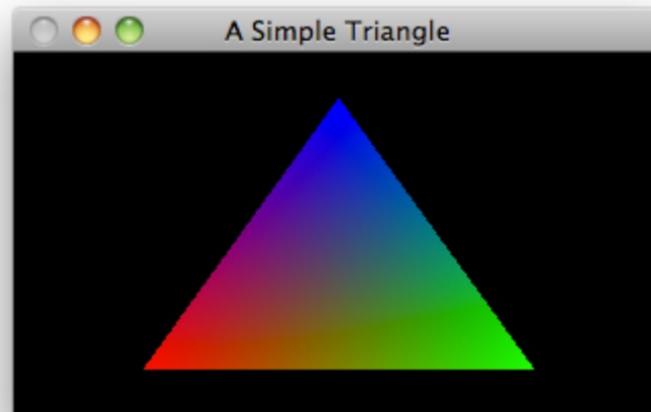
In-Editor
Simulation



Experimental
Features

XR Native Renderers

- OpenGL
 - OpenGL is a cross-language, cross-platform application programming interface for rendering 2D and 3D vector graphics. The API is typically used to interact with a GPU, to achieve hardware-accelerated rendering.



XR Native Renderers

- DirectX
 - Microsoft's graphics API

DirectX is composed of multiple APIs:

- [Direct3D](#) (D3D): Real-time [3D rendering](#) API
- [DXGI](#): Enumerates adapters and monitors and manages swap chains for Direct3D 10 and later.
- [Direct2D](#): 2D graphics API
- [DirectWrite](#): Text rendering API
- [DirectCompute](#): API for [general-purpose computing on graphics processing units](#)
- DirectX Diagnostics (DxDiag): A tool for diagnosing and generating reports on components related to DirectX, such as audio, video, and input drivers
- [XACT3](#): High-level audio API
- [XAudio2](#): Low-level audio API
- [DirectX Raytracing](#) (DXR): Real-time [raytracing](#) API
- [DirectStorage](#): GPU-oriented file I/O API
- [DirectML](#): GPU-accelerated machine learning and artificial intelligence API

<https://learn.microsoft.com/en-us/windows/win32/directx>

XR Native Renderers

- Vulkan
 - Vulkan is a low-level low-overhead, cross-platform API and open standard for 3D graphics and computing. It was originally developed as Mantle by AMD, but was later given to Khronos Group. It was intended to address the shortcomings of OpenGL, and allow developers more control over the GPU.

Vulkan is preferred over OpenGL nowadays

XR 3D Modeling

- Blender
 - Blender is a free and open-source 3D computer graphics software tool set used for creating animated films, visual effects, art, 3D-printed models, motion graphics, interactive 3D applications, virtual reality, and, formerly, video games.

<https://youtu.be/f-mx-Jfx9lA?t=236>

XR 3D Modeling

- Maya
 - Better modeling features compared to Blender
 - Mainly for enterprises – not free, not open

XR Gaming Engines

- Unity & Unreal
 - Unity was originally an Apple game engine but slowly spread to many platforms
 - Both provide game developers with a 2D and 3D platform to create video games.

	Unity	Unreal
Developer	Unity Technologies	Epic Games
Written in	C# (Unity Scripting API)C++ (runtime)	C++
Supported platforms	Mobile, desktop, web, console, VR/XR	Mobile, desktop, console, VR/XR (less than Unity offers)
Primary audience	Mobile, indie, and beginner developers	AAA devs and indie teams striving for realism
Ease of use	Beginner-friendly interface	Steep learning curve
Open source	No	Yes
Price	Free to use (until the product has earned more than \$100k in the last 12 months)	Free to use (a 5% royalty if the product earns more than \$1 million)
2D/3D support	Yes	Yes (limited for 2D)

XR 3D Scanners

- Matterport (<https://matterport.com>)



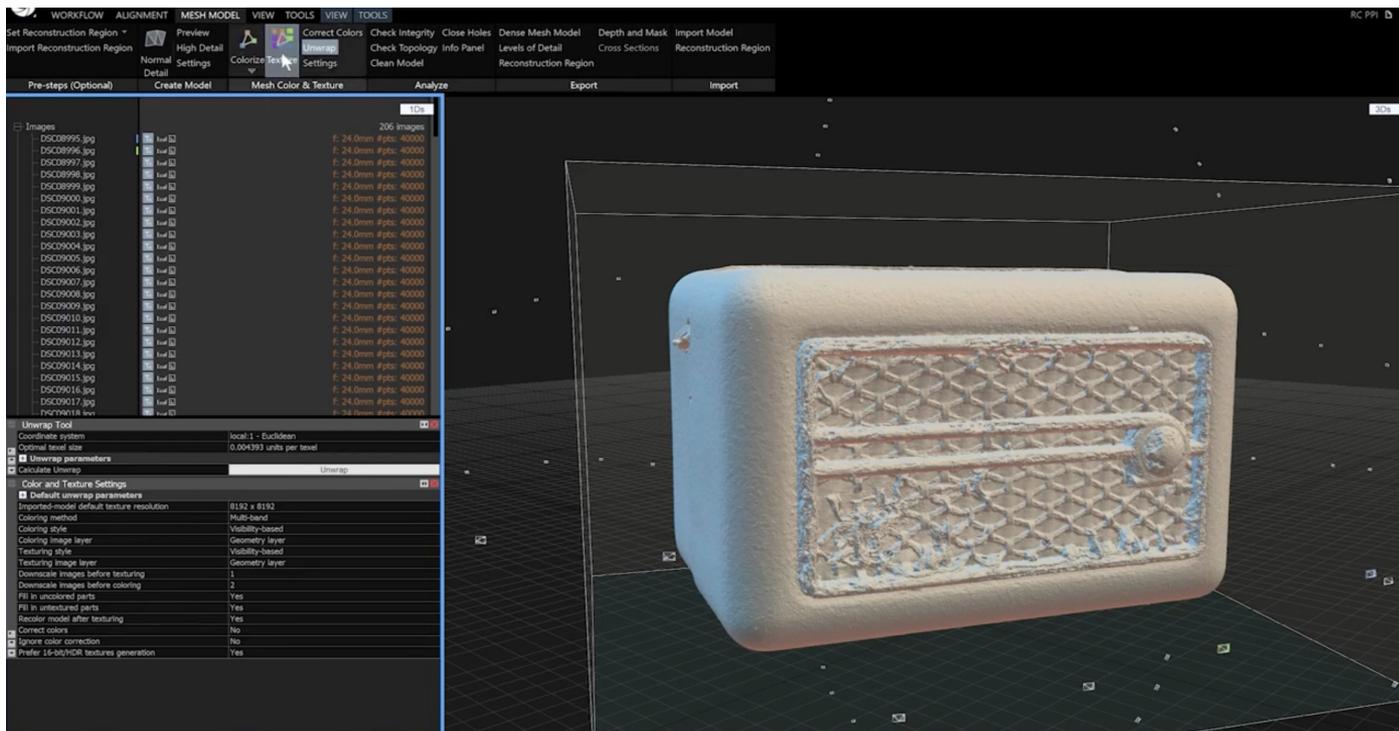
<https://jamesandharrisoncourt.com/virtual-tours>

XR 3D Scanners

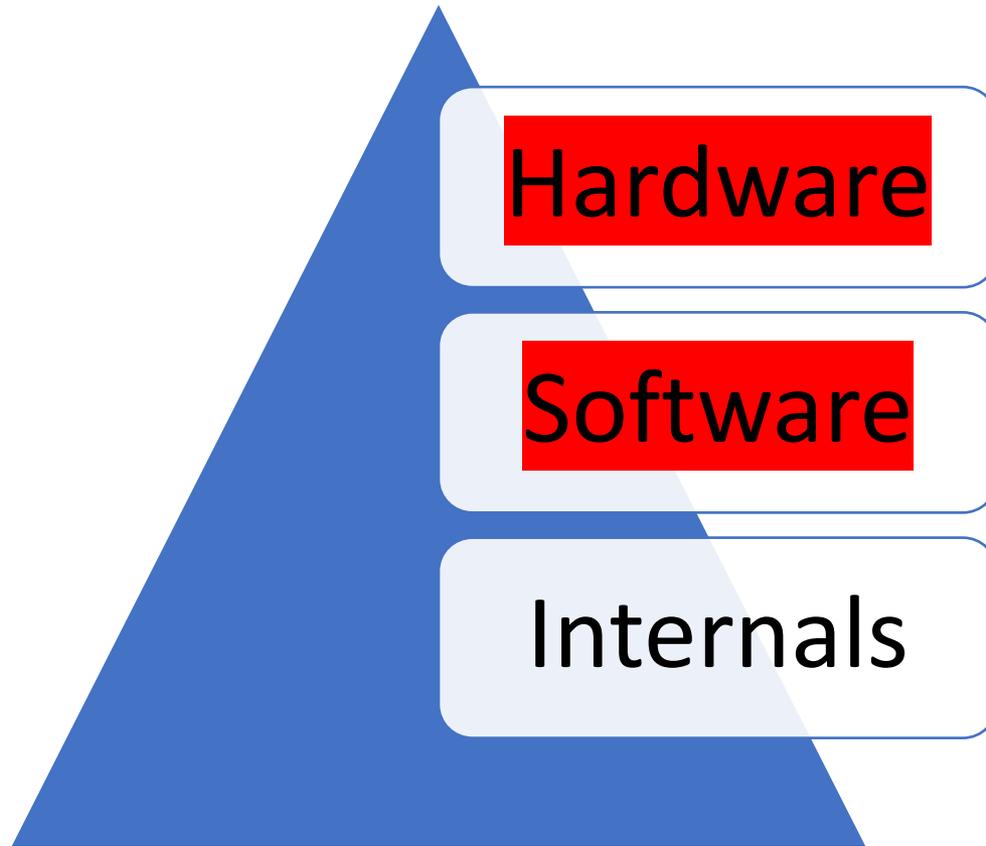
- Scaniverse (<https://scaniverse.com>) - Mobile
- a 3D scanning app that supports all recent iPhones and iPads, including those without LiDAR. Scaniverse uses photogrammetry to accurately reconstruct objects, rooms, and even whole buildings and outdoor environments.

XR 3D Scanners

- RealityCapture – photogrammetry + manual editing



Lecture Outline for Today

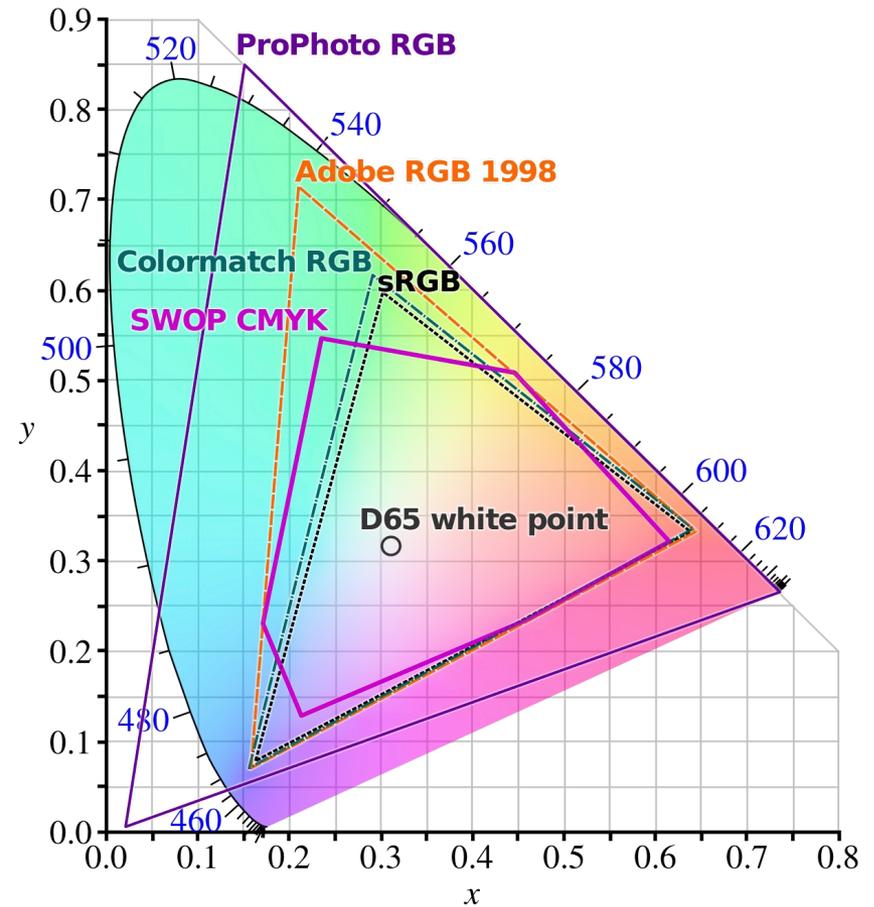


XR Internals

- Perception
- Motion to Photon Latency
- Positioning and Tracking
- 3D Reconstruction
- Real-time Rendering

XR Perception

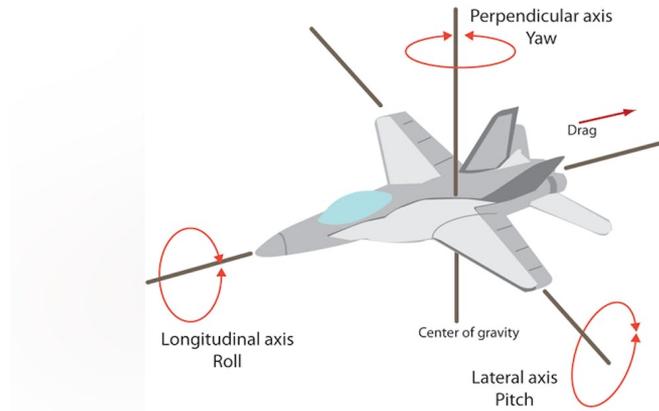
- Visual
 - Color
 - Quality/spatial resolution
 - Depth resolution
 - Temporal resolution
 - Field of view
- Non-visual
 - Sense of touch
 - Audio
 - Balance
 - Smell



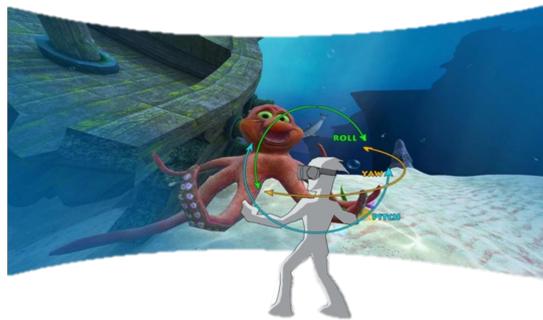
Positioning and Tracking

- You need to know where you are in the world

- GPS?
- Visual
- Inertial
- Lidar
- RF



- 3-DoF
- 6-DoF



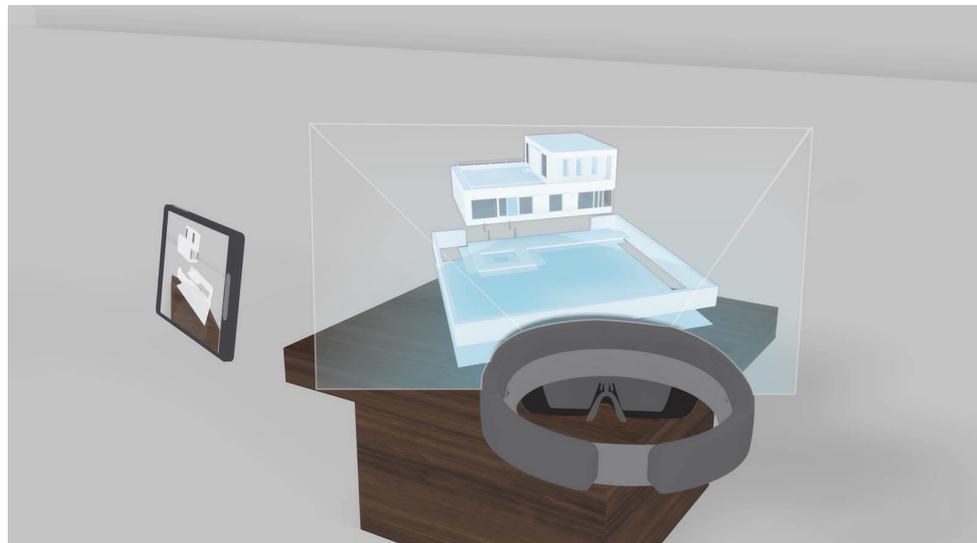
X, Y, Z & Yaw, Pitch, Roll

Positioning and Tracking

- Anchors

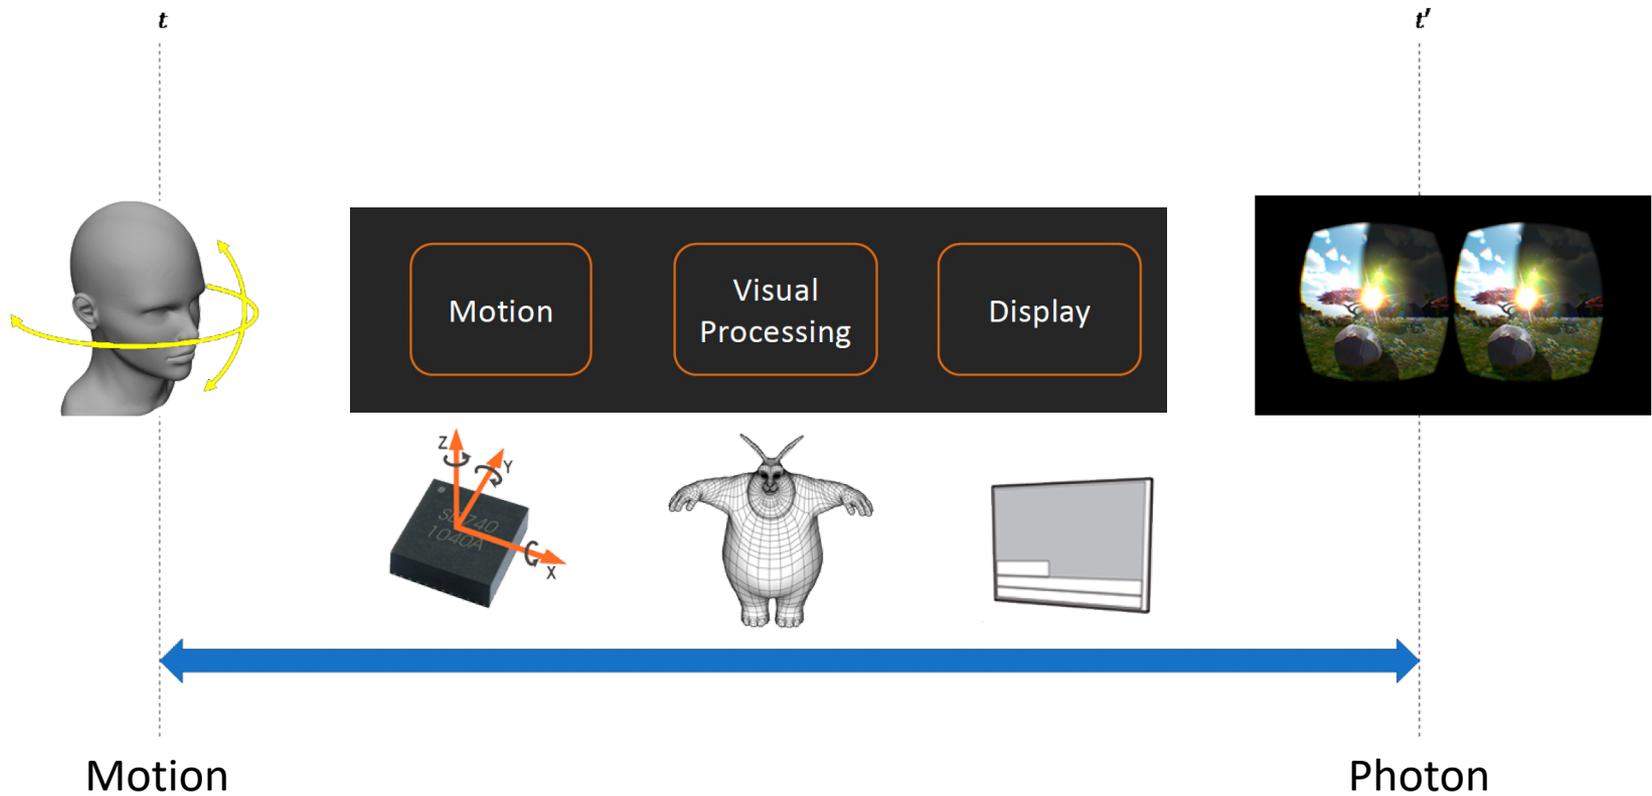
- Anchors ensure that objects appear to stay at the same position and orientation in space, helping you maintain the illusion of virtual objects placed in the real world.

- Plane
- Wall
- Floor
- Face...



- Anything that you can identify well

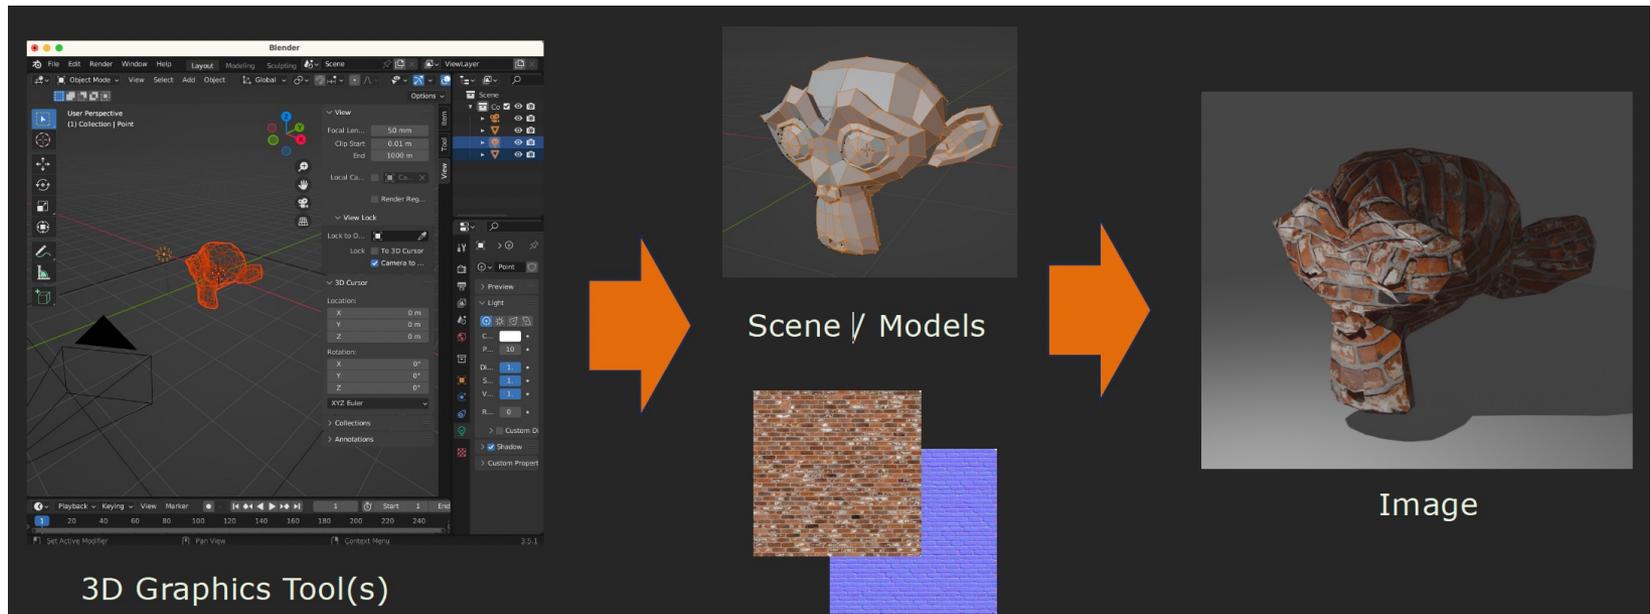
Motion to Photon Latency



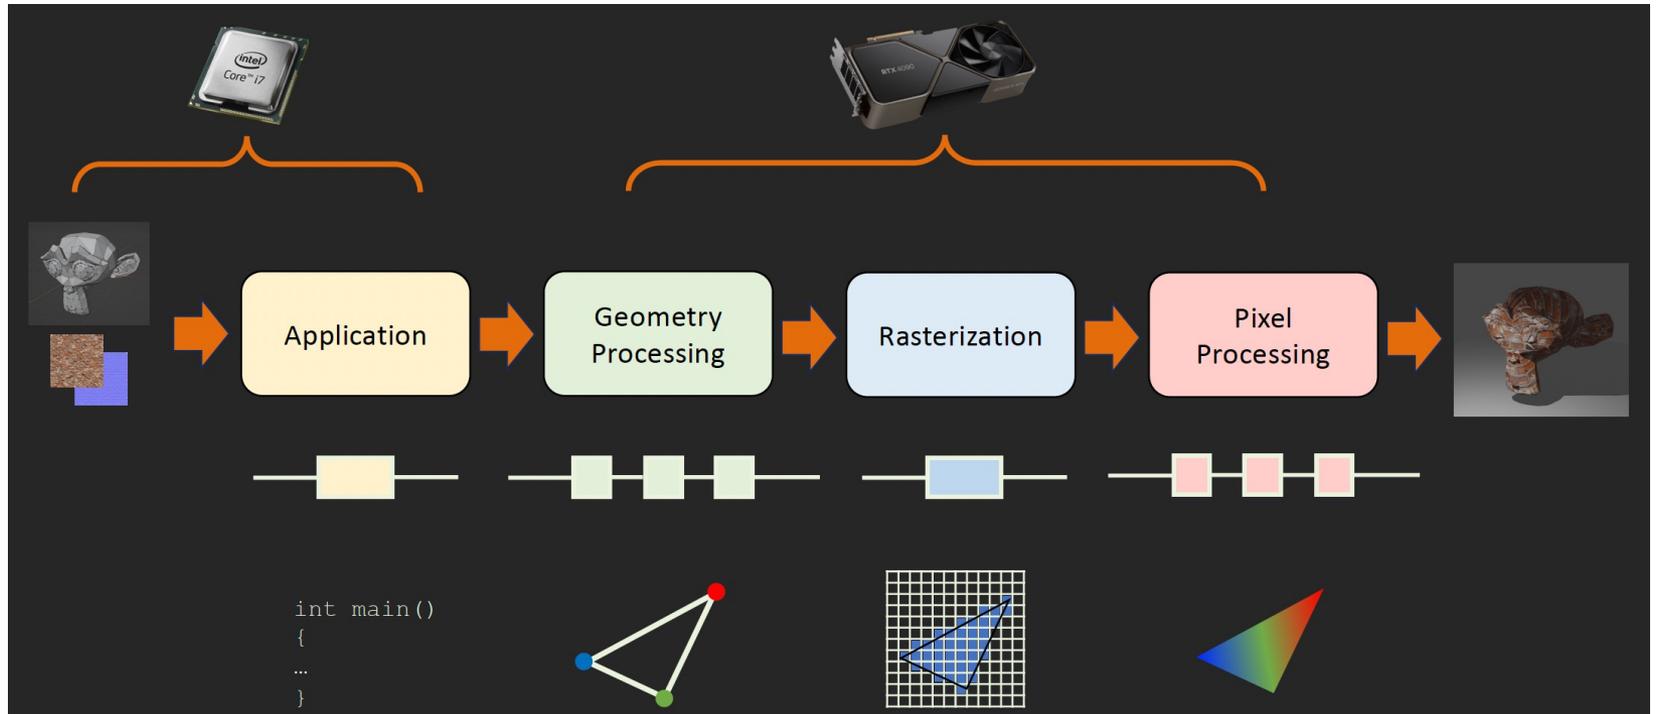
3D Reconstruction



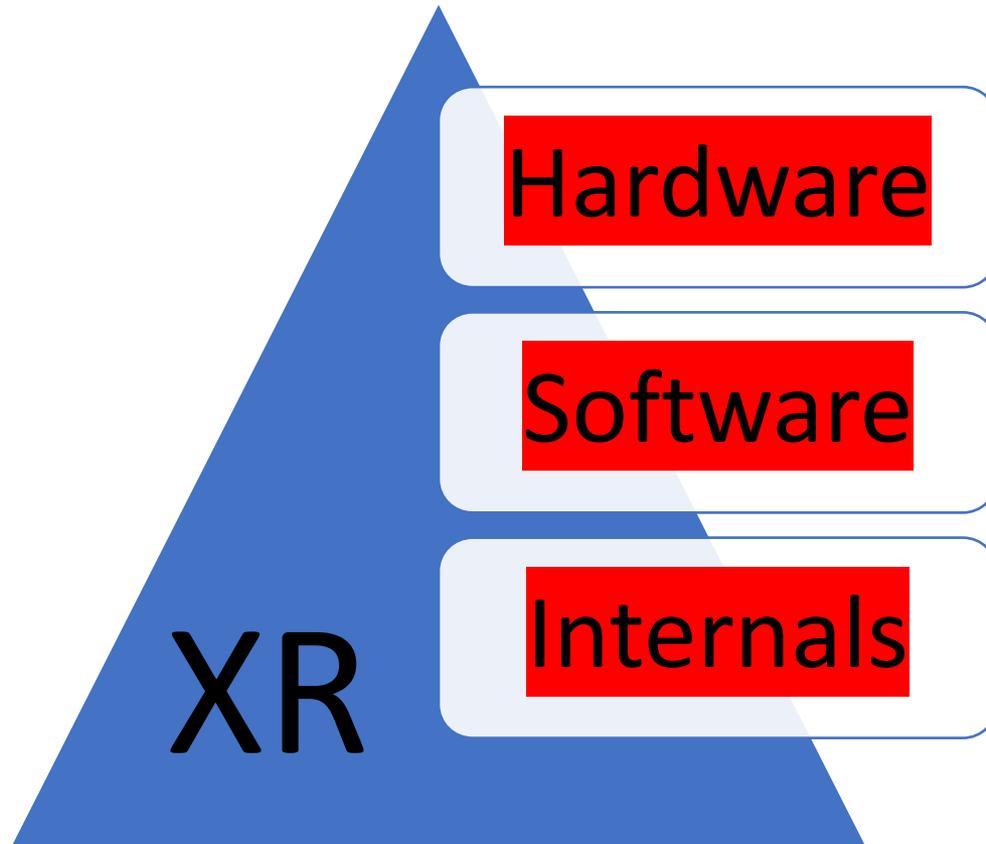
Real-time Rendering



Real-time Rendering



Lecture Summary



Next up: XR Sensors and Sensing Algorithms