CMSC498F/838C & ENEE759N: Advances in XR

TuThu 3:30-4:45 cs.umd.edu/class/spring2022/cmsc498F/

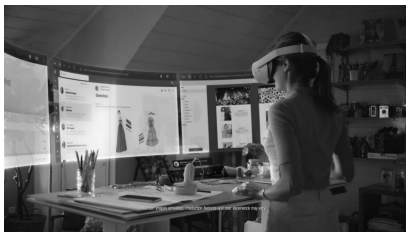


Instructional Team


Prof. Ming Lin


Nick Rewkowski

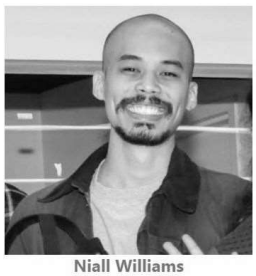
Niall Williams

Ming C. Lin

http://www.cs.umd.edu/~lin/

- BS, MS, PhD in Electrical Engineering & Computer Science
  University of California at Berkeley
- B. Mersky & CapitalOne E-Nnovate Endowed Professor 2021-
- UMD Distinguished University Professor, 2019-
- Former Elizabeth Iribe Chair of Computer Science @ UMD, 2018-2020
- J.R. & L.S. Parker Distinguished Prof. Emeritus @ UNC Chapel Hill
- ACM, IEEE & Eurographics Fellow; ACM SIGGRAPH Academy
- Areas of Research: **Virtual Reality, Robotics, AI/ML/Vision & Autonomy** with focuses on *physically-based modeling/simulation*, *multimodal interaction* (haptics & audio technology), *animation*, and *human-computer interaction*; applications in autonomous driving, virtual try-on, healthcare, digital design, rapid prototyping, and personalized fabrication/cybermanufacturing

Nick Rewkowski (nickvr.me)

- PhD student working on multimodal XR applications
- Undergrad & MS at UNC Chapel Hill
- Taught HCI-focused VR class at UNC
- Research journey: Surgical reconstruction → 3D audio → Locomotion → Haptics → HRI → AR surgical training → simulating humans → metaverse/spatial documents
  Programming games for ~13 years (~8 years Unity & UE4)
- Adobe collab for the past year on our informational metaverse vision (spatial documents)

My guinea pigs >

Niall Williams (niallw.github.io)

- PhD student working on VR locomotion.
- Undergrad from Davidson College.
- Research interests in XR, human perception, HCI, and robotics.

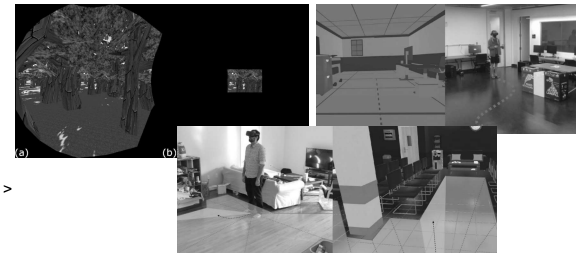

Some of my research >

What is Extended Reality (XR)?

- Software & hardware that replaces or mixes real world stimuli with synthetic
- Different types of XR mostly differentiated by **display** and **tracking** methods

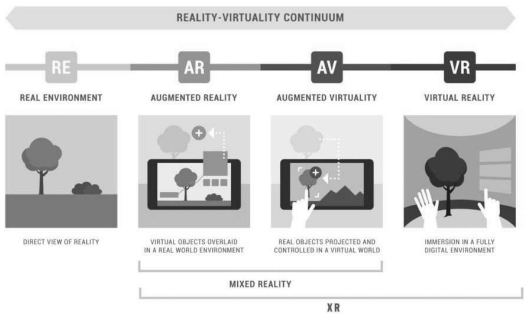

Sutherland 1965: The Ultimate Display

It (the Ultimate Display, referring to VR and AR) is a looking glass into a mathematical wonderland……. If the task of the display is to serve as a looking-glass into the mathematical wonderland constructed in computer memory, it should serve as many senses as possible. - Ivan Sutherland, 1965

- A virtual world, through a HMD, appeared realistic thru augmented 3D sound & tactile feedback
- Computer hardware to create the virtual word and maintain it in real time
- The ability users to interact with objects in the virtual world in a realistic way

Sutherland 1968: First XR headset

- "A head-mounted three dimensional display" Ivan Sutherland 1968
- Concept first introduced by a concept paper to DARPA in 1965

Virtual Reality (VR)

- Main "illusions" needed for immersive experience (Mel Slater 2009):
  - **Place illusion**: feeling like you're in the virtual world and not the real world
  - **Plausibility illusion**: feeling like what happens in the virtual world is really happening
- "Pure VR" only handles virtual world and assumes real stimuli completely replaced with synthetic stimuli (e.g. Oculus Rift, HTC Vive, Google Cardboard)
  - If the synthetic stimuli not handled correctly, people can get simulator sickness, lack of "presence" or immersion, or worse (e.g. trauma)
- Focus typically on games

Augmented Reality (AR)

- Overlays digital 'elements' onto real world, including graphics, images, video, sound, GIS data, text, animation, etc.
- Through head-mounted display, hologram, video-passthrough, etc.
- Focus on simple experiences or information easily available
- "Pure AR" only overlays information on real world but virtual world and real world do not interact/real world info is not needed (e.g. Google Glass, old Snap Spectacles, marker-based tracking–Nintendo 3DS, volumetric AR display, etc.)

Mixed Reality (MR)

- Merge of real world and virtual world, co-existing & interacting in real time
- Focus on practicality, productivity, integration with day-to-day life
- Tracking real world features & hands makes the **Oculus Quest** an MR headset
- **Hololens** handles real world & virtual 3D world simultaneously, so it's also MR
- The **class HMD** will be MR since exposed camera allows for inside-out tracking

XR Trends Over Time

XR Trends Over Time: Pre-1970s

- Fundamental tracking, rendering, & display technology (esp. For simulators)
- First VR headsets (stereo headsets in the 1800s!)

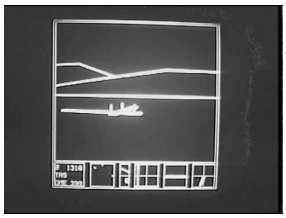

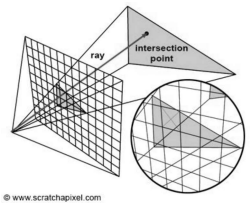

XR Trends Over Time: 1970s

- Fundamental 3D graphics technology (e.g. Phong shading)

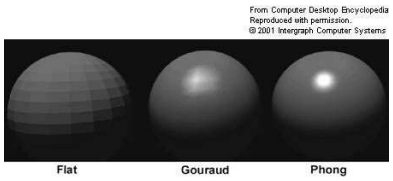

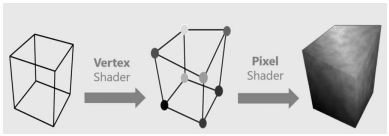

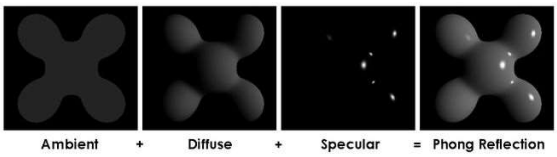

XR Trends Over Time: 1980s

- Physically-Based Rendering (PBR)
- VR Simulators (esp. flight)
- Interactive Games
- Display adapters (old GPUs)
- Multimodal XR apps

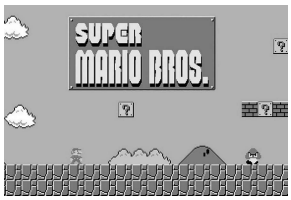

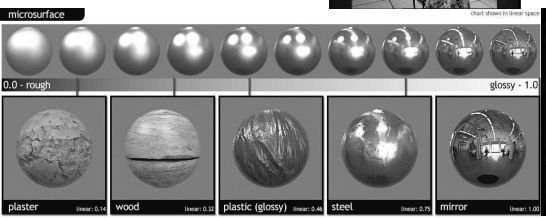

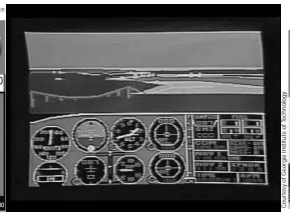

XR Trends Over Time: 90s & Early 2000s

- Clumsy but effective XR headsets
- Boom of XR psychological studies
- 3D Game Engines
- Interactive 3D graphics
- Collision detection & interactive physics
- Boom in areas like haptics & locomotion

5 View of UNC nanoManipulator connected to an atomic force microscope.

3 NASA-Houston's "Charlotte" Virtual Weightless Mass lets astronauts practice handling weightless massy objects.

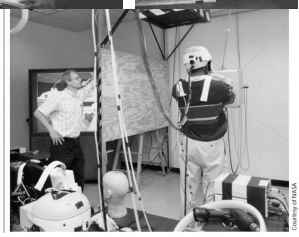

XR Trends Over Time: Late 2000s

- Research on different modalities, audio, locomotion, rendering→all evolved close to what they are today
- Unity, Unreal 2/3, idEngine, Source Engine, Autodesk Maya, etc. established many interactive 3D graphics conventions

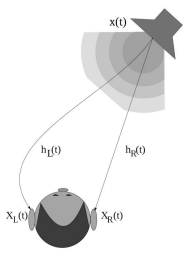

XR Trends Over Time: 2010s

- Good mobile technology
- Strong GPUs
- Decent commercial headsets
- Industry getting involved

XR Trends Over Time: Early 2020s

- Everyone is making XR headsets
- Metaverse hype
- Social XR

XR Trends Over Time: Predicting Late 2020s

- Focus on application development
- Procedurally-generated content
- Better, large-scale HRI
- Virtual assistants
- Attempts at neural interfaces

Challenges in XR

Challenges in XR: HMD Design

- Weight
- Physical dimensions & portability
- Field of view (FOV)
- Battery vs. computing power
- Pixel opacity
- Optics

Challenges in XR: Graphics

- Interactive PBR & realistic materials
- Foveated rendering & optimization
- Physically-based interactions

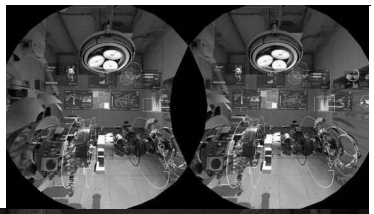

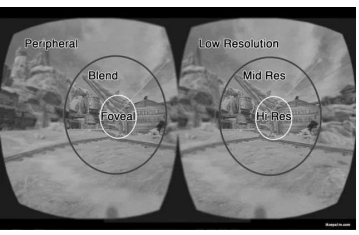

Challenges in XR: Audio

- Procedural generation vs. sound synthesis (esp. on mobile)
- Sound propagation & 3D acoustics (esp. on mobile)
- Personalized 3D audio display

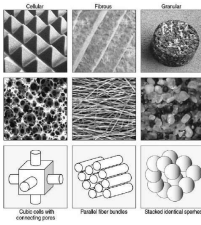

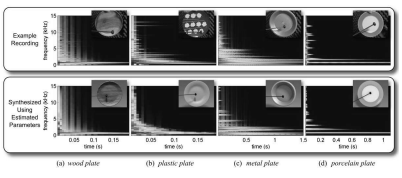

(a) *wood plate* (b) *plastic plate* (c) *metal plate* (d) *porcelain plate*

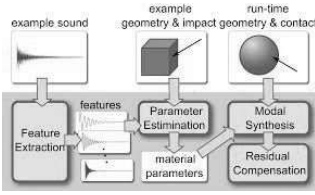

Challenges in XR: Virtual Humans

- Uncanny valley
- AI behavioral modelling
- Social cues
- Procedural animation & rigging
- GANs & ML-based creation
- UV mapping/texturing

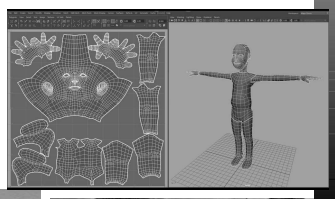


Challenges in XR: Content

- Procedural generation
  - Doesn't only apply to environments! Used to speed up getting good textures, models, animation, audio, etc.
- GANs & ML-based content
- User-generated vs. company-generated



Challenges in XR: Application Design & Utility

- How to keep people using it?
- How to reduce barrier of entry?
- How to integrate with people's lives?

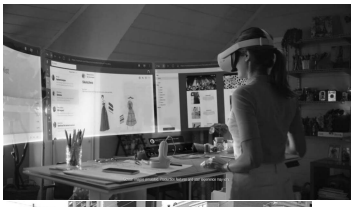


Challenges in XR: Natural Locomotion in VE

- Getting people to walk around naturally with limited space
- Handling sickness & perception (knowledge of the "illusion" changes its effect!)

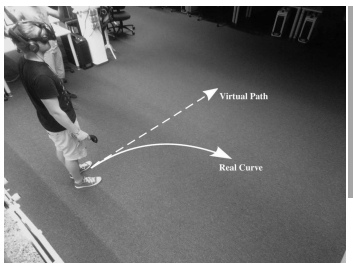

Challenges in XR: Tracking & Reconstruction

- **Local** tracking of headset more or less solved
- **Global** tracking still challenging
  - How to register virtual environments to real ones?
  - How to link the virtual environments together?
  - How to overlay real and virtual worlds?
- Reconstructing people, real environments, etc. to make virtual world as convincing as real world
  - Photogrammetry is the state-of-the-art for high-fidelity asset creation but it's incredibly slow & tedious


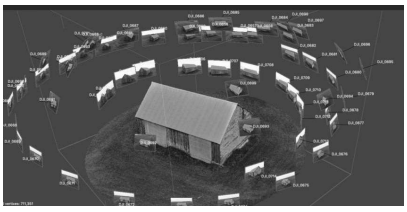


Challenges in XR: SW/HW Compatibility

- Huge issue with Metaverse, NFTs, etc.
- How can assets trivially move between XR applications as if they are part of a unified metaverse?
- **Walled garden** vision of Facebook vs. mostly **user-generated** approach of VRChat/Second Life

Current Research Directions

Challenges in XR: Natural Interactions with Full-Body Tracking

- Eyes, hands, etc.
- Predict user intent
- HMD's FOV of hands
- Gesture recognition
- Accessibility

Research Directions: Natural Virtual Locomotion

- Use distortions to trick people into walking differently than they realize.
- Maximize real walkable space.
- Must support a range of users and environment shapes.

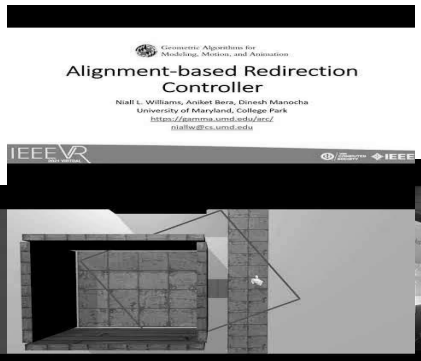

Research Directions: Intelligent Virtual Environments

- Future of work
- Context-aware interaction
- Integration with daily life

Research Directions: Emotive Virtual Humans

- Realistic virtual assistants w/ procedural animation
- Passing the uncanny valley

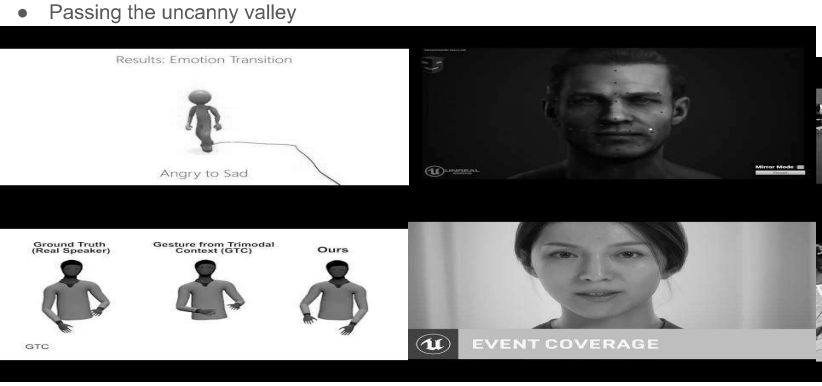

Research Directions: Human-Robot Interaction in XR

- Robot dogs to train the visually-impaired spatial mapping skills
- Robots in warehouses for collaborative XR

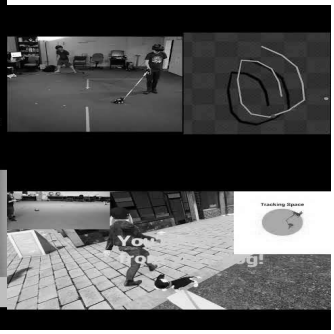

Research Directions: Metaverse

- (Seamlessly-Integrated Real & Virtual Environments)
- **Human-centric mapping between spaces**: pocket dimensions/ portals/ symbolic links, human-centric motion planning, context, behavior prediction
- **“Internet of XR”/ “IoX”**: sensor/device fusion, connected HMDs, maximizing knowledge of real world

Research Directions: Intelligent, Responsive XR

- Immersive, XR Interfaces for autonomous systems
  - Driverless vehicles
  - Autonomous drones
  - Intelligent sensor networks
- XR Systems for rapid design and personalized fabrication/manufacturing
- XR Systems for remote collaboration, tourisms, social events
- XR Systems for personalized healthcare and wellness
- XR Interface for virtual try-on
- XR Systems for social good
  - Immersive scenario replay for Police Training, Bias Training, AI for Fairness, etc.
- Audio-Visual Reconstruction

Course Information

Course Goals

- Understand **multimodal XR design**
  - Basic technology is not enough!
  - Basic principles of audio, haptic, and visual rendering
  - How multiple modalities work together and interact with each other
  - Challenges of Metaverse & future of XR
- Understand “realism”, “illusion”, and “presence”
  - Roles of locomotion
  - Sense of ‘being there’
  - Avatar & object animation/simulation
  - Social interaction via virtual humans
- Application design & development
- Use the state-of-the-art APIs & XR tools

Course Components

- Lectures & Participation
- Homework Assignments [30%]
- Final Project [35%]
- Midterm & Final exams [35%]

Lectures

- Will be recorded, but with emerging tech, the discussion is very beneficial.
- Rough sequence:
  - Basic 3D graphics
  - Basic game development
  - Immersion & presence
  - Virtual Locomotion
  - 3D Audio
  - 3D Animation
  - Virtual agents & AI
  - Tracking & Reconstruction
  - Displays, Optics, Lenses, etc.
  - Advanced Topics (light fields, haptics, olfactory, GANs, holography, etc.)
- Invited speakers from industry & academia

Why Are We Learning Game Development?

- **Game engines**: powerful, realtime, interactive, multimodal 3D applications
- De facto standard for XR development and most interactive 3D consumer programs
  - Unity typically better for mobile XR or indie developers
  - UE4 typically better for desktop VR or future jobs in game dev
- Both engines have libraries covering most XR topics!
- You can choose engine based on preference & future goals. Should not affect difficulty much. Unity is typically easier for new game devs.

UNREAL ENGINE

unity

Headset Logistics

- Possibilities:
  - Class-provided headset
    - Requires building for your own phone
    - ARCore (Android) and ARKit (iOS) mostly the same
  - Oculus Quests (IRB 0110)
  - Personal devices
- Assignments will use the same HMDs except for variation of implementation modules & API difference
- Some assignments require class HMD
  - (AR features provided by phone camera)
- Anyone using personal devices?

Assignments

- Address the major modalities & concepts of XR design
  - Goal is **breadth** of experience; final project is opportunity to dive deep
- Build on each other to result in complete multimodal XR application
- Sequence
  - **A1:** Setting up the XR development environment
  - **A2:** 3D modelling and virtual environment creation
  - **A3:** Game development for XR; gaze-interacting with virtual objects
  - **A4:** Hand-tracking, hand interaction with your virtual environment, & AR markers
  - **A5:** Natural virtual locomotion
  - **A6:** Adding 3D sound to a virtual environment
  - **A7:** Animated virtual agents (adding autonomous virtual humans to the environment)
  - **A8(G):** Basic inverse kinematics (allowing you to commandeer a virtual avatar)
- May include non-technical readings
- Grad section will have slightly extra requirements & higher expectations
- Submission includes code/Git repo & video showing it works

Introducing Assignment 1

- Development Environment Setup
- Very different – depending on combination of platforms, but more or less the same moving forward
  - Try not to do it at the last minute! Almost Guaranteed to run into some kind of errors
- 2 fundamental readings:
  - **"What's Real About VR?"** Fred Brooks 1999
  - **"Recent advances in augmented reality"** Ron Azuma 2001
  - Will answer some short-answer questions about them

Final Project

- Will dive deeper into:
  - one of the topics from the suggested list
  - something not covered in depth (e.g. later class topics like HRI, haptic displays, body/object/scene reconstruction)
- Previous final projects
  - *Examples:*
    - Inter-human behavioral tree for realistic interactions with virtual humans in VR
    - Realistic bouncing audio "probes" to help visually-impaired people go through a maze
    - Objectively measuring frustration with design & mechanics in a virtual world
    - Creating 3D minimap for navigation in VR
    - Simulate depth-of-field in gaze-based interactions with 3D objects
    - Study how acoustic materials affect performance in a treasure hunt
    - Sound propagation for realistic instrument playing in VR

Exams

- Midterm
- Final
- Relatively high-level design questions
  - E.g. How you would advise a client who wants specific XR app built? What are the important considerations?
  - E.g. Someone wants to make an XR app a certain way; what problems are they not addressing correctly?
  - E.g. Why do we need to do [a particular thing] for XR?
  - E.g. How did [this paper] likely accomplish [this feature]?
- No programming, but will ask high-level questions about technical problems
  - E.g. why are VR lenses designed the way they are
  - E.g. what do waveguides and lightboxes do in near-eye AR displays?

Collaboration & General Policies

- Alone or pair for regular assignments
- 1-5 people in a group for final project
  - Final report must describe what each team member does
  - Each will be graded based on individual contributions specified
  - Group efforts should show higher levels of complexity and components
- Recommendation: try to stay with the same OS & platform combination
- No collaboration on exams
  - Planning to be in-class
- Lateness
  - A total of 24 hours grace period for the entire semester
  - 10% off per fraction of week late
  - Late assignments due to illness or unexpected events can be excused with doctor's notes or other forms of written indication

Office Hours (subject to change)

- Open Lab Hours: Monday 1pm-4pm for the first 3-5 weeks

  @ AR/VR Lab (IRB 0110) on the ground level

- Office hours
  - Ming: Tues/Thur after class or by appointment (email: lin@umd.edu)
  - Nick: Friday 2-5pm (Zoom ID "nrewkowski2") or by request (email nick1@umd.edu)
  - Niall: Wednesday 1pm - 3pm (Zoom ID "niallw") or by request (niallw@umd.edu)

Communication Method: Piazza, Slack, or Discord?

- Piazza:
  - Pros: made for classes, everyone knows how to use it
  - Cons: hard to have ongoing conversation, not used outside of classes, no videoconferencing, slow response
- Slack:
  - Pros: good for professional stuff, most people know how to use it
  - Cons: resource intensive, buggy, not great videoconferencing, somewhat overkill in features
- Discord
  - Pros: most gamers already use it regularly, great videoconferencing & group calls for OH, lightweight, basically same as Slack otherwise, good for hybrid classes
  - Cons: non-gamers probably don't use it, not really meant for classes
- VOTE! Survey at tinyurl.com/xrclasspoll

Background form

tinyurl.com/umdxrsurvey

Mask Policy

- Wear KN95 masks in classrooms
- Must have booster and follow UMD testing protocols
- Keep yourself and everyone else healthy and safe!

Questions?

Course Website:

https://www.cs.umd.edu/class/spring2022/cmsc498F/index.html