Lecture 20:

Course Summary + Graphics at Stanford Today

Interactive Computer Graphics
Stanford CS248, Winter 2022

As accomplished CS248 students you've now learned the basics of drawing shapes, representing surfaces/light/ materials/motion, and manipulating images, etc…

(and have been introduced to core graphics ideas like sampling, anti-aliasing, acceleration data structures, etc.)

What's Next?

More graphics classes at Stanford

SPRING

CS348B: "Image Synthesis Techniques", theory and practice of realistic, physically-based rendering (Fatahalian, James, Pharr) - TTh 3:15pm

CS348K: "Visual Computing Systems", creating efficient systems for computational photography, 3D graphics, and AI (Fatahalian) - TTh 1:30pm

CS348E: "Character Animation: Modeling, Simulation, and Control of Human Motion" (Liu) - MW 3:15pm

EE267: "Virtual Reality", focuses on display and tracking hardware for VR (Wetzstein) - MW 9:45am

FALL

CS 348I: "Computer Graphics in the Era of AI", deep learning methods for computer graphics (Liu, Wu)

CS 448B: "Data Visualization" (Agrawala)

WINTER

CS348C: "Animation and Simulation", deep dive into animation and simulation techniques (James)

EE367/CS448i: "Computational Imaging and Display", advanced course on display design (Wetzstein)

CS205L: "Continuous Mathematical Methods with an Emphasis on Machine Learning" (Fedkiw)

CS348B (Fatahalian, James, Pharr)

Rendering realistic images by modeling the physical process of light interacting with materials

With ray tracing as the mechanism to simulate these phenomenon

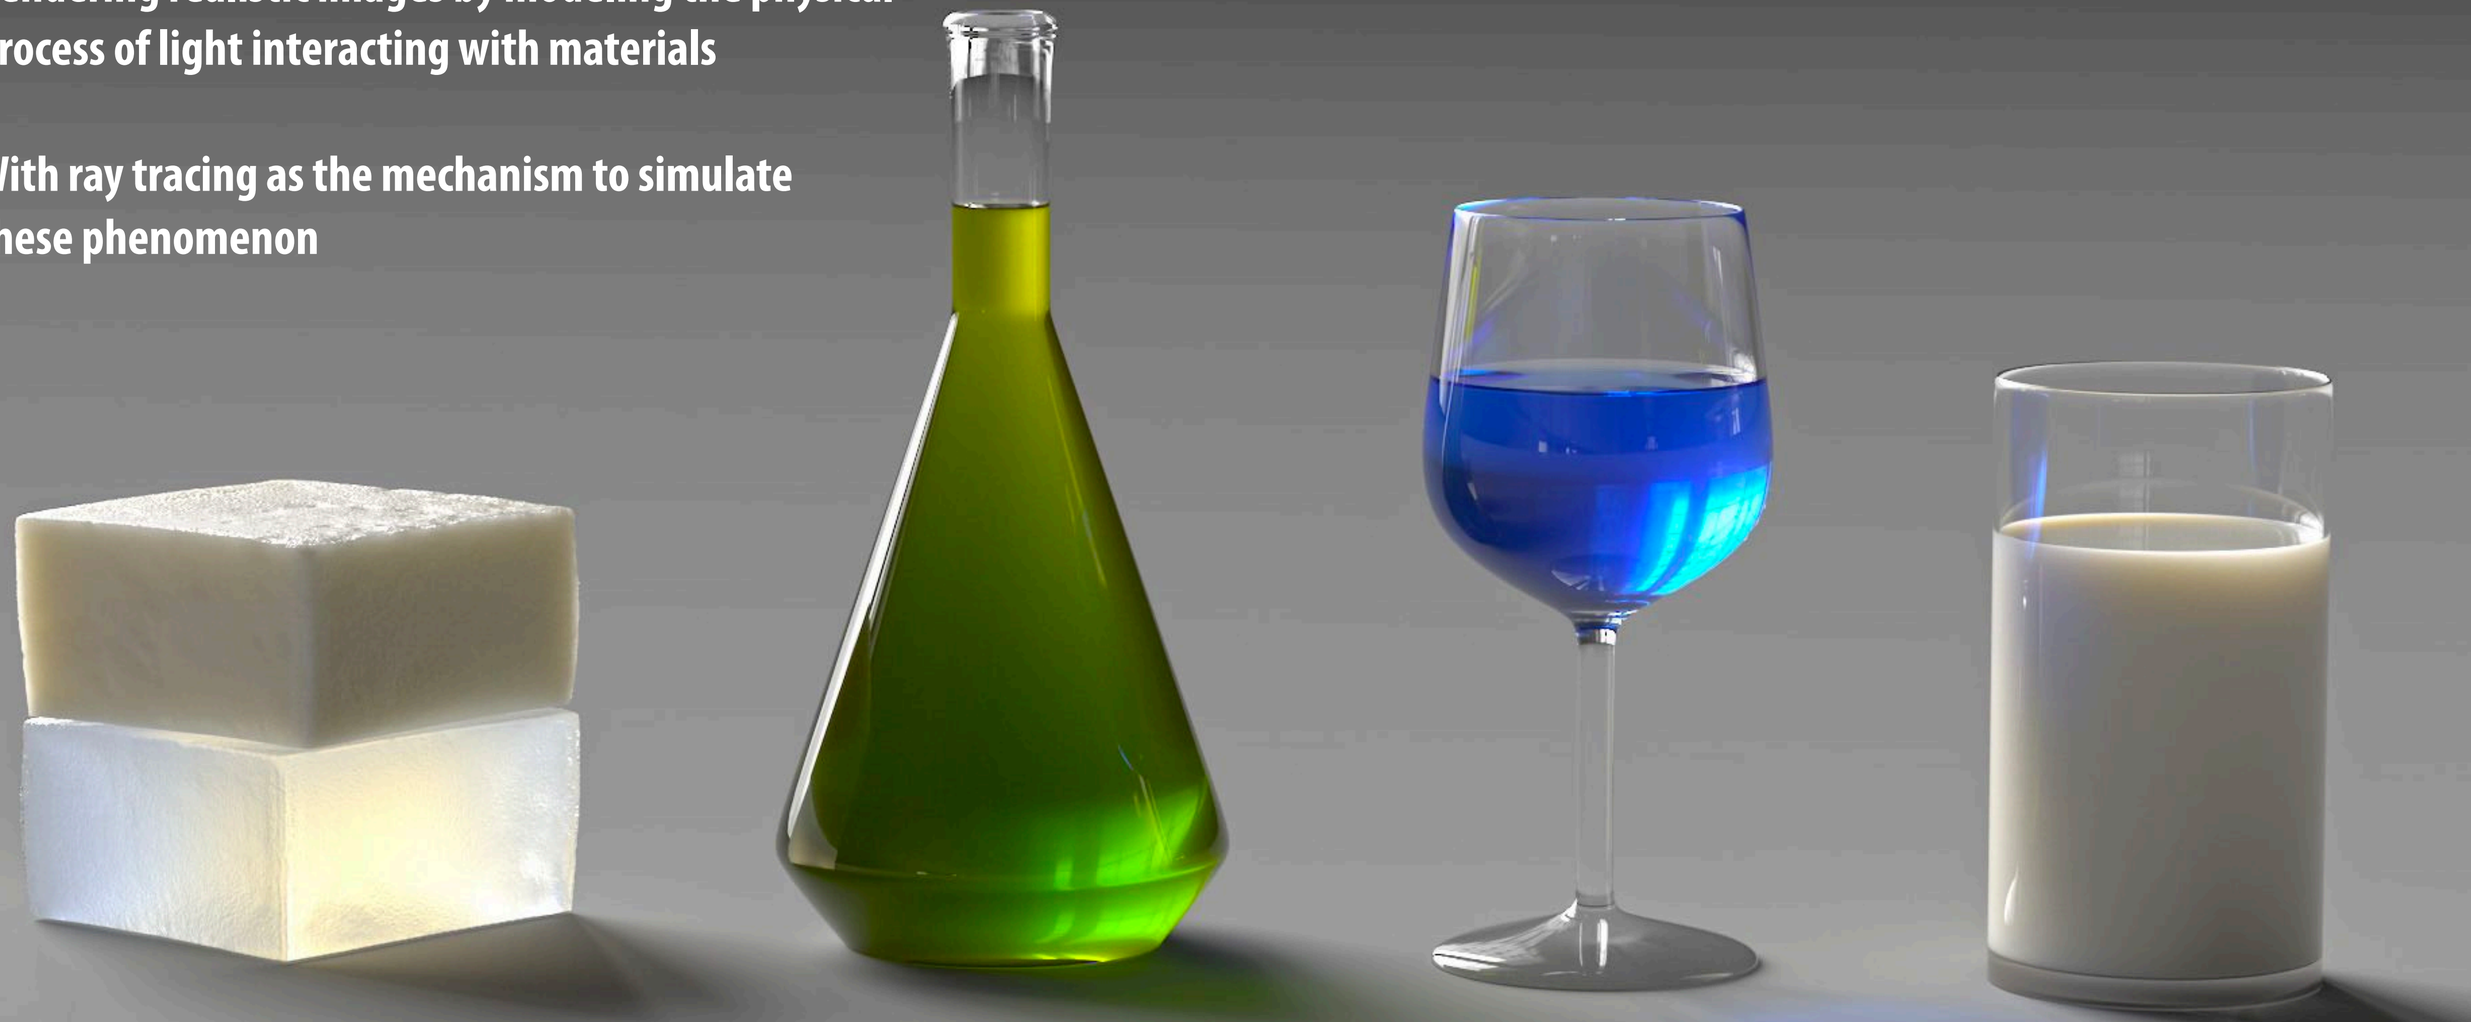

Graphics Research at Stanford Today

Maneesh Agrawala

Maneesh Agrawala

Many current projects on video editing and manipulation

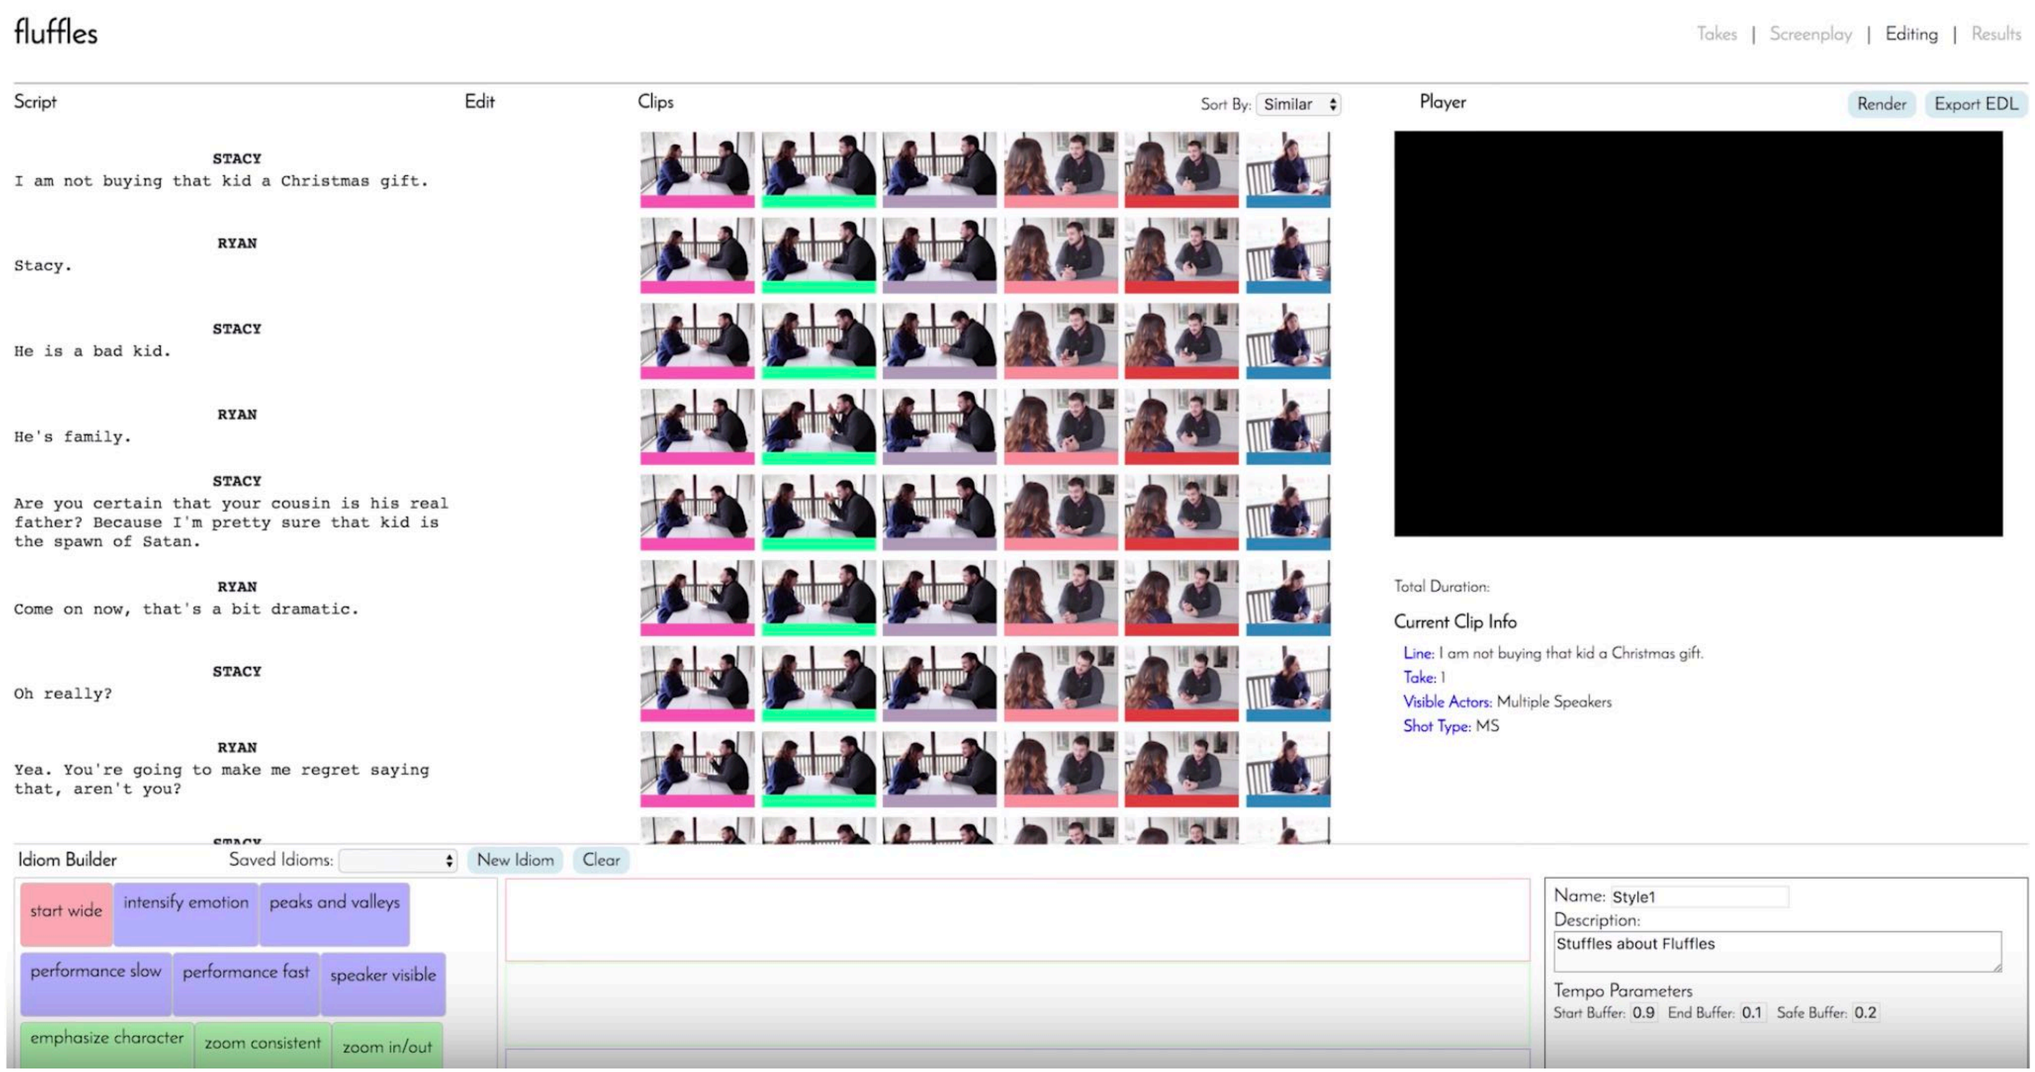

Maneesh Agrawala

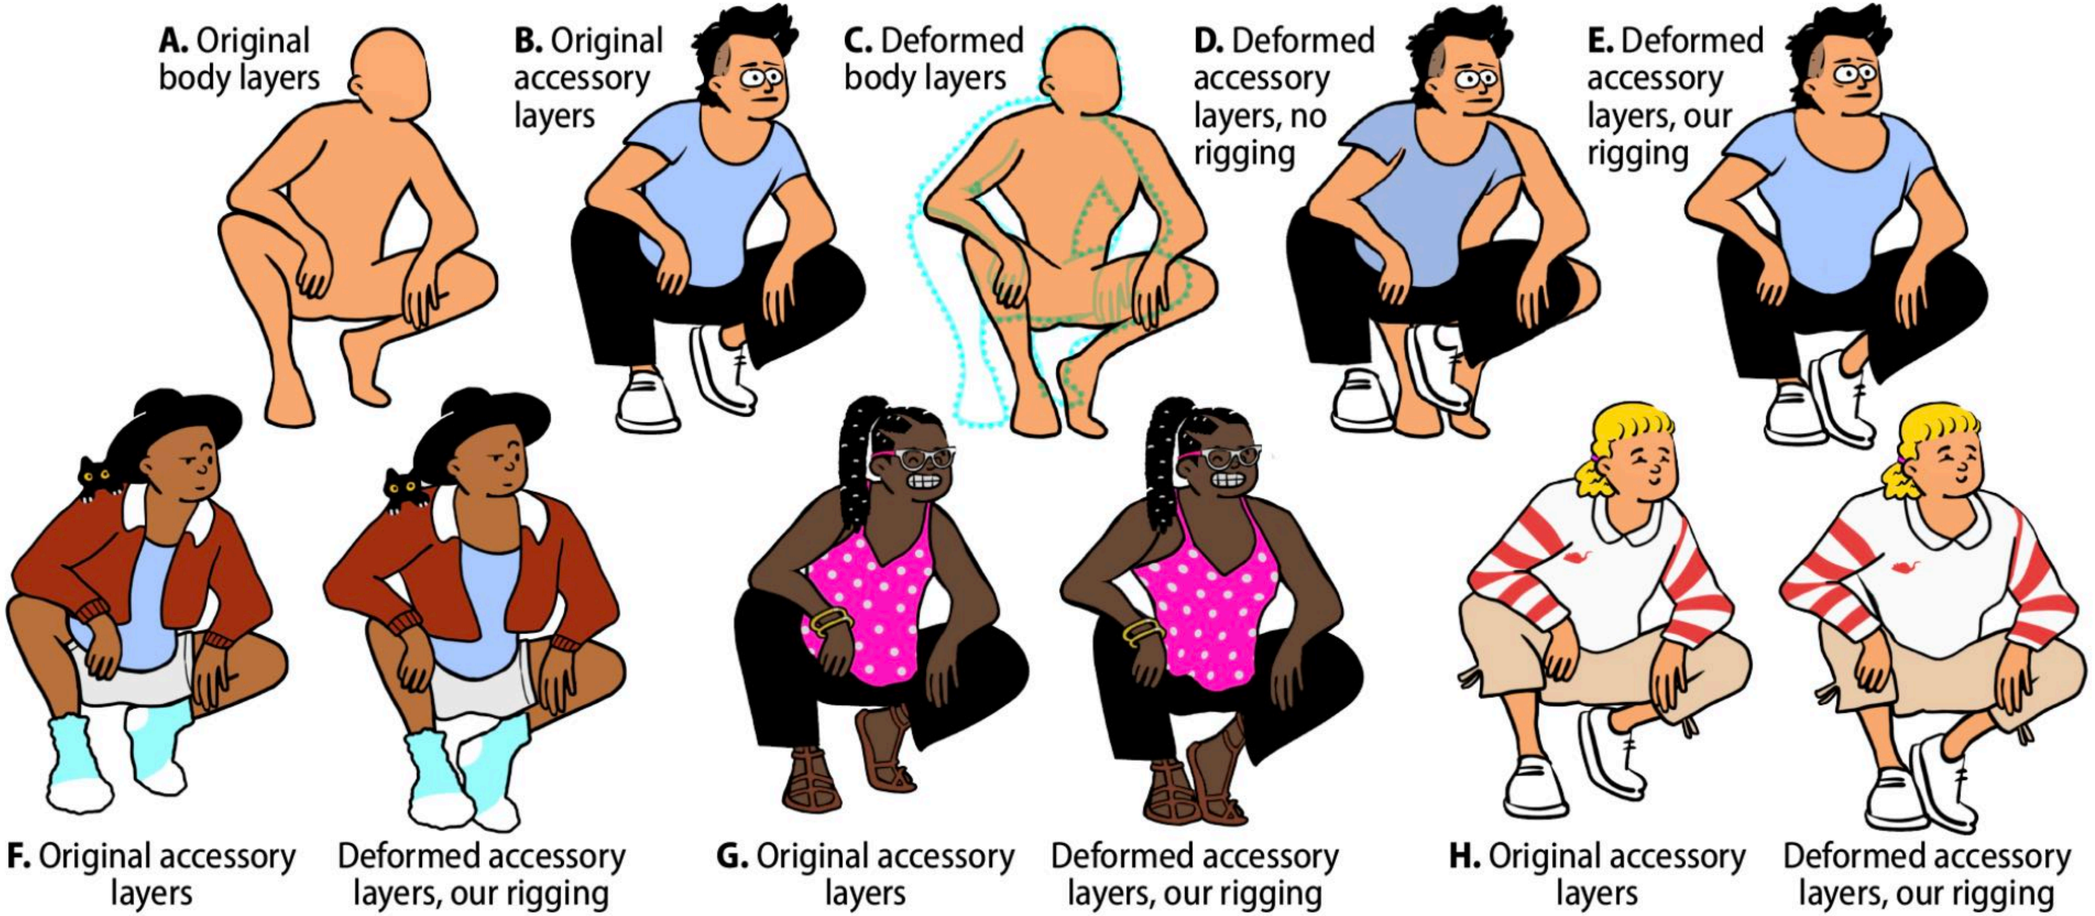

Ron Fedkiw

- Simulation techniques (often) targeted at film and game production
- Now exploring use of machine learning to augment or improve physical simulations

Ron Fedkiw

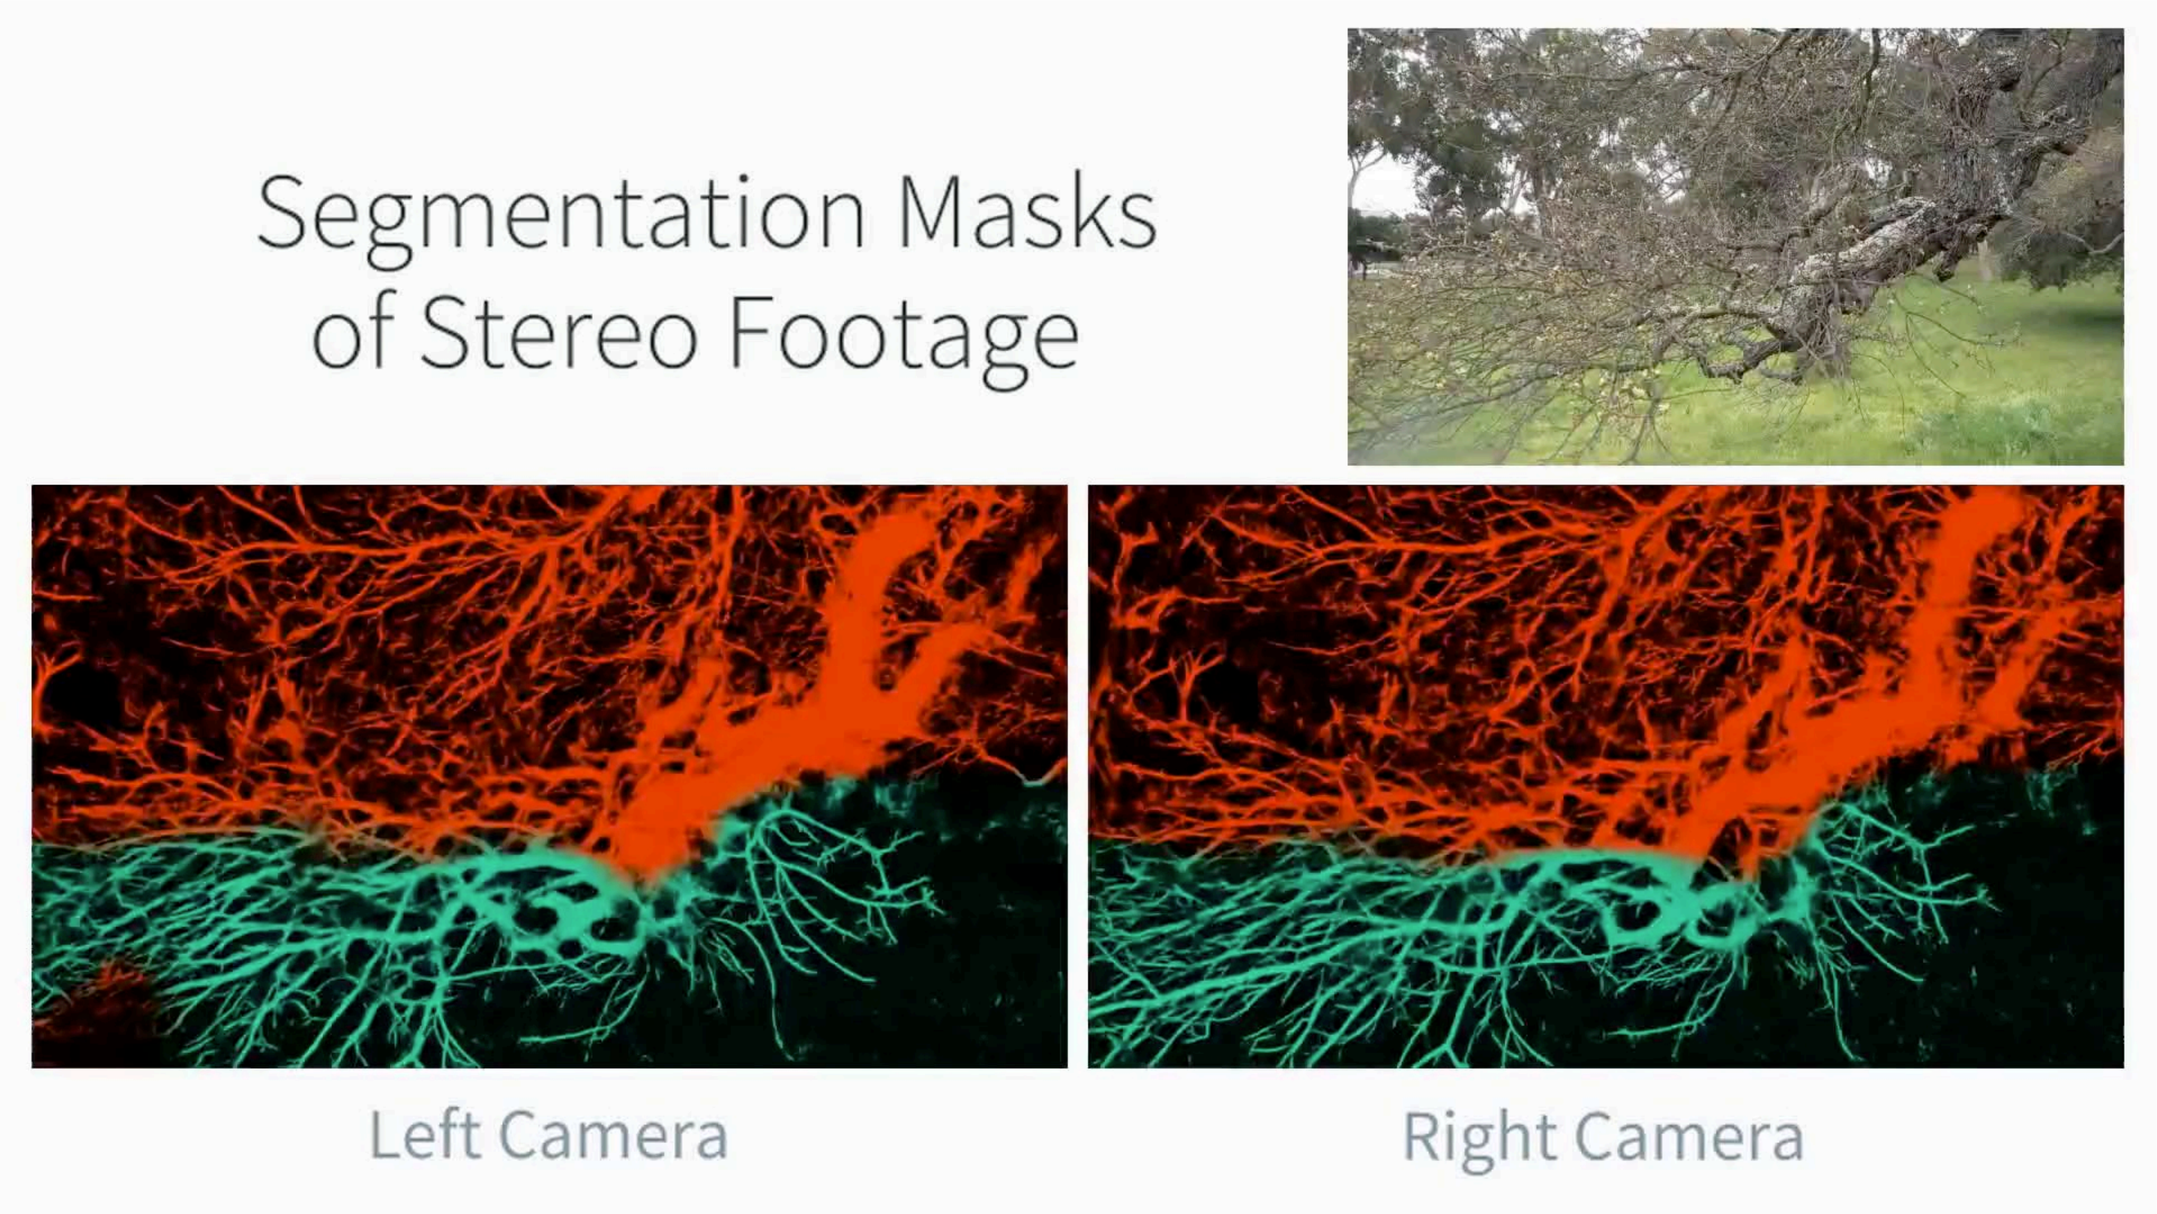
Segmentation Masks
of Stereo Footage
Left Camera
Right Camera

Ron Fedkiw

Doug James

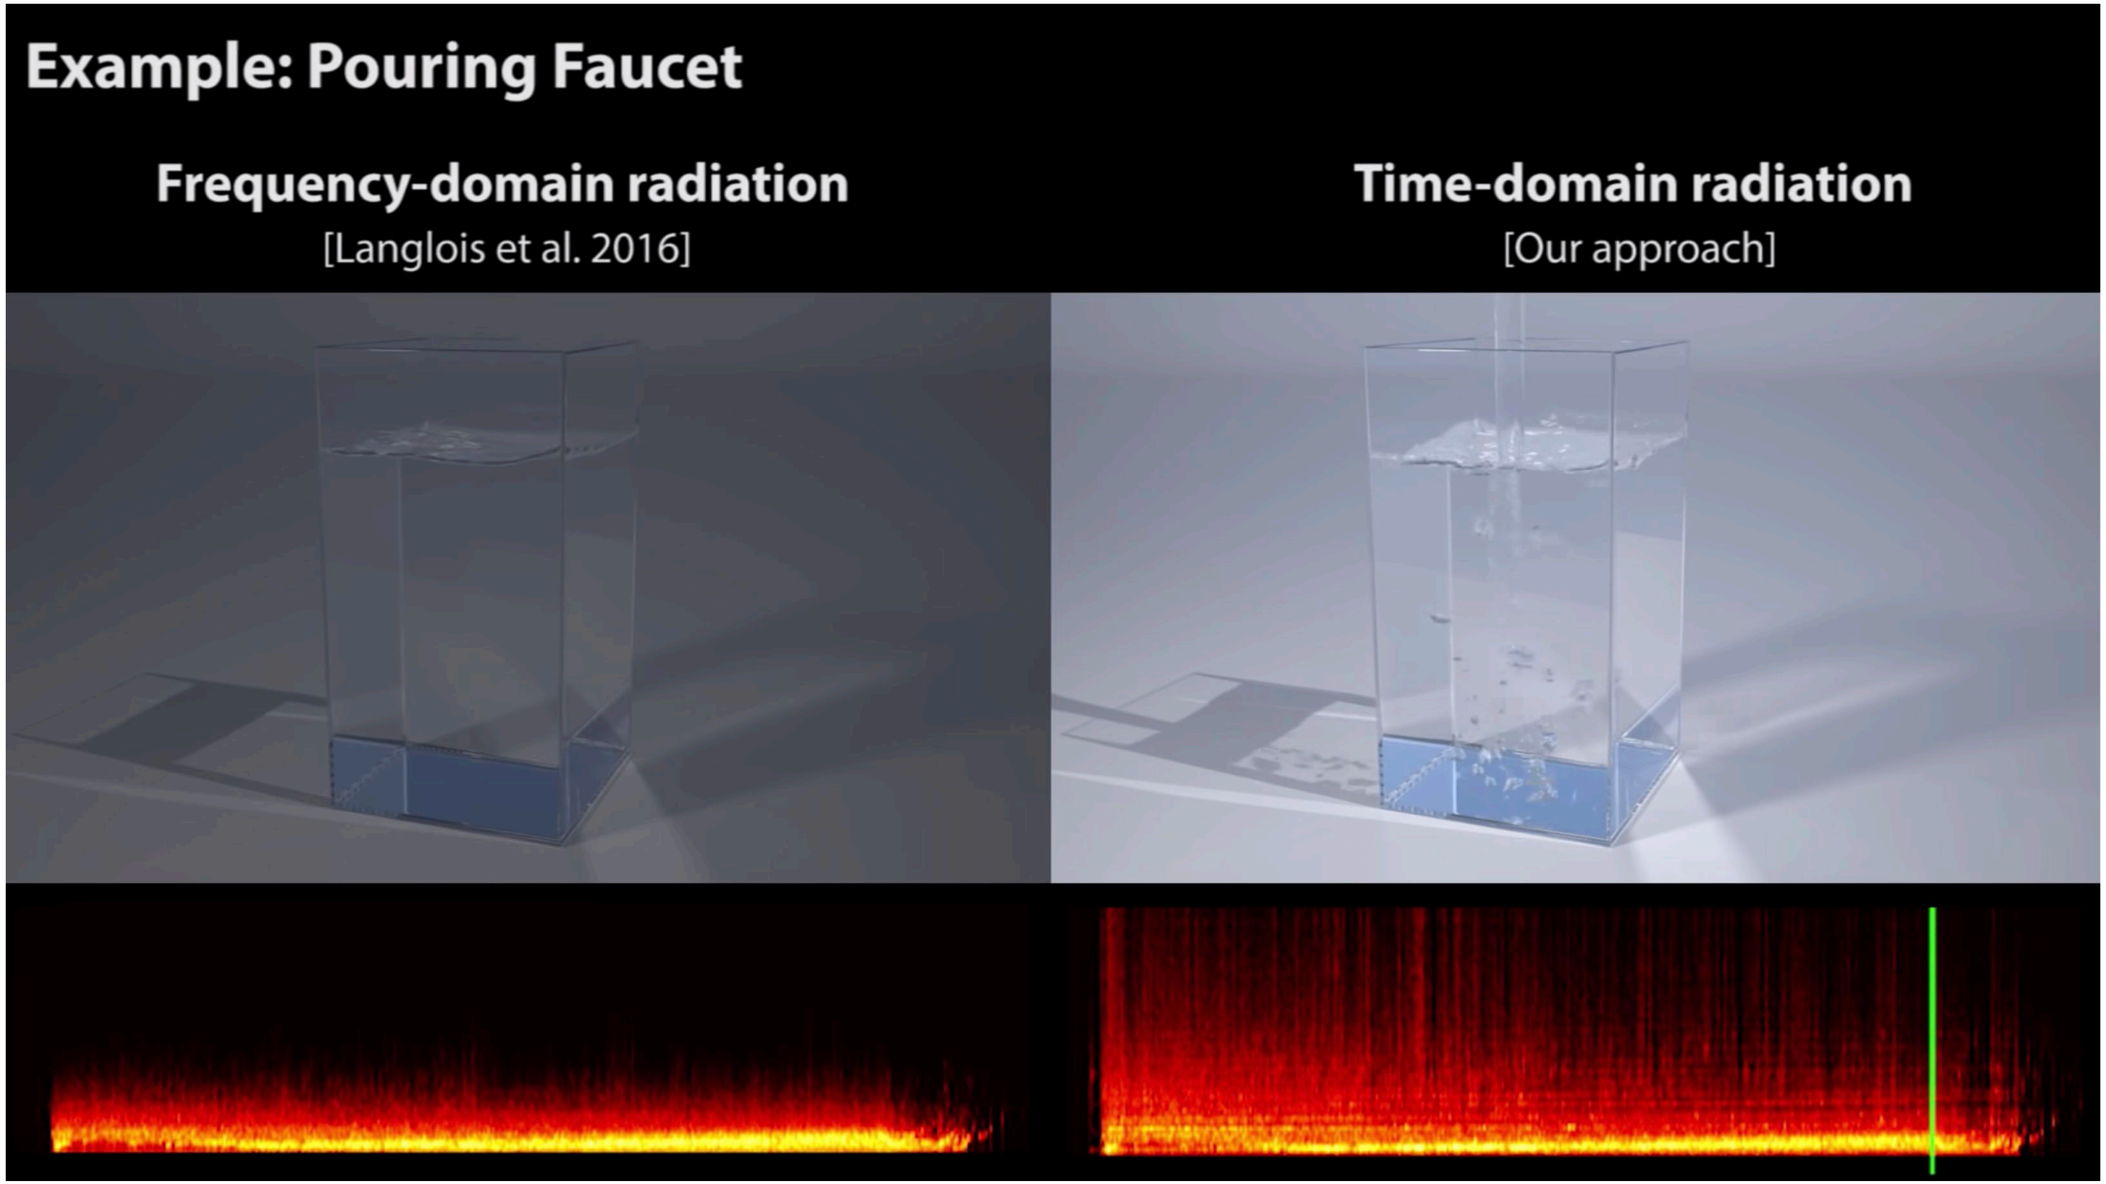
Example: Pouring Faucet
Frequency-domain radiation
[Langlois et al. 2016]
Time-domain radiation
[Our approach]

Doug James

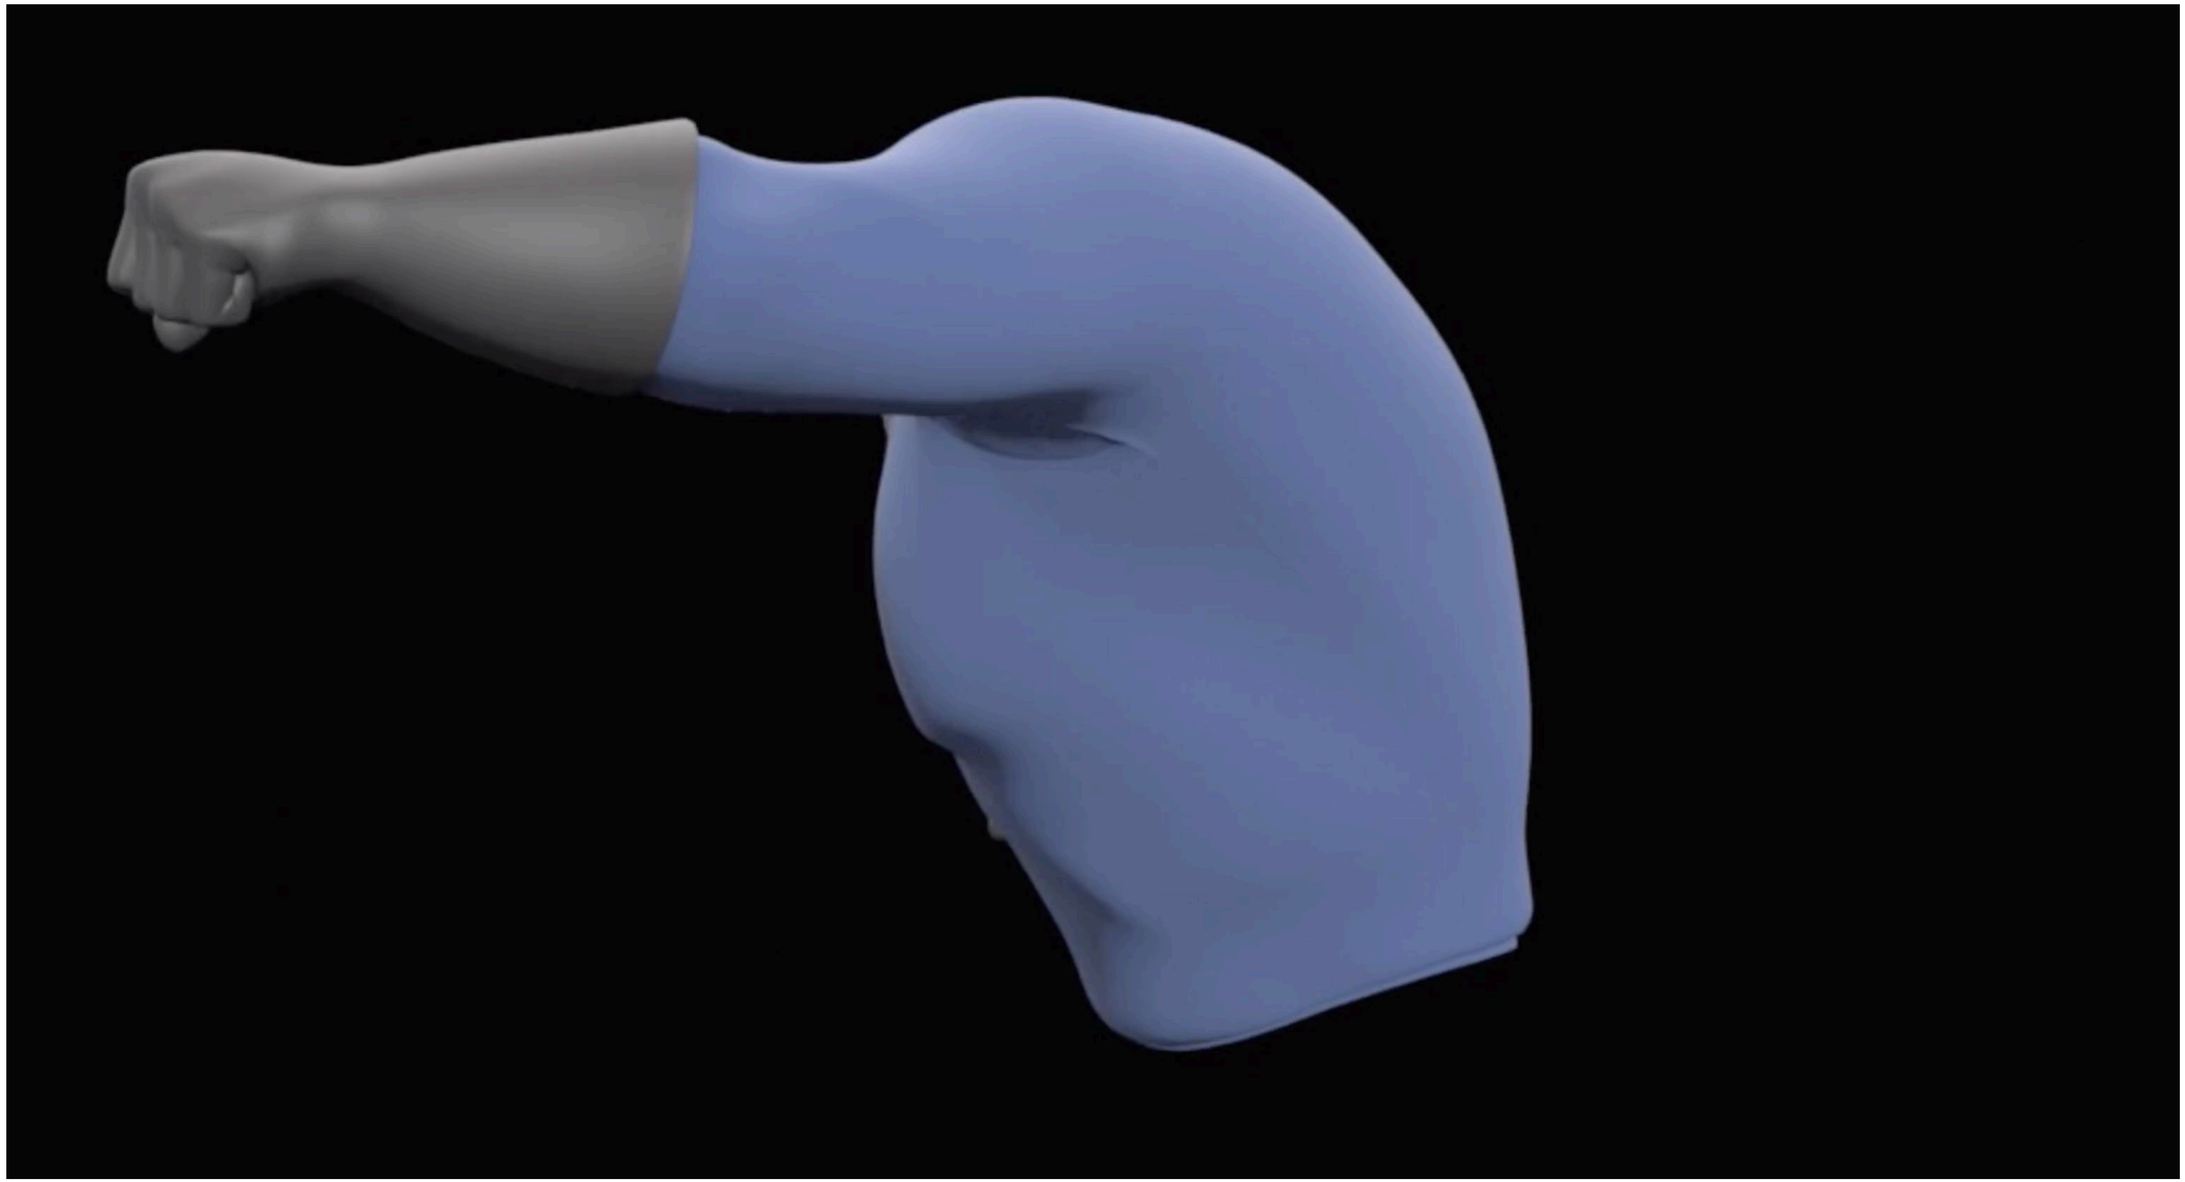

Doug James

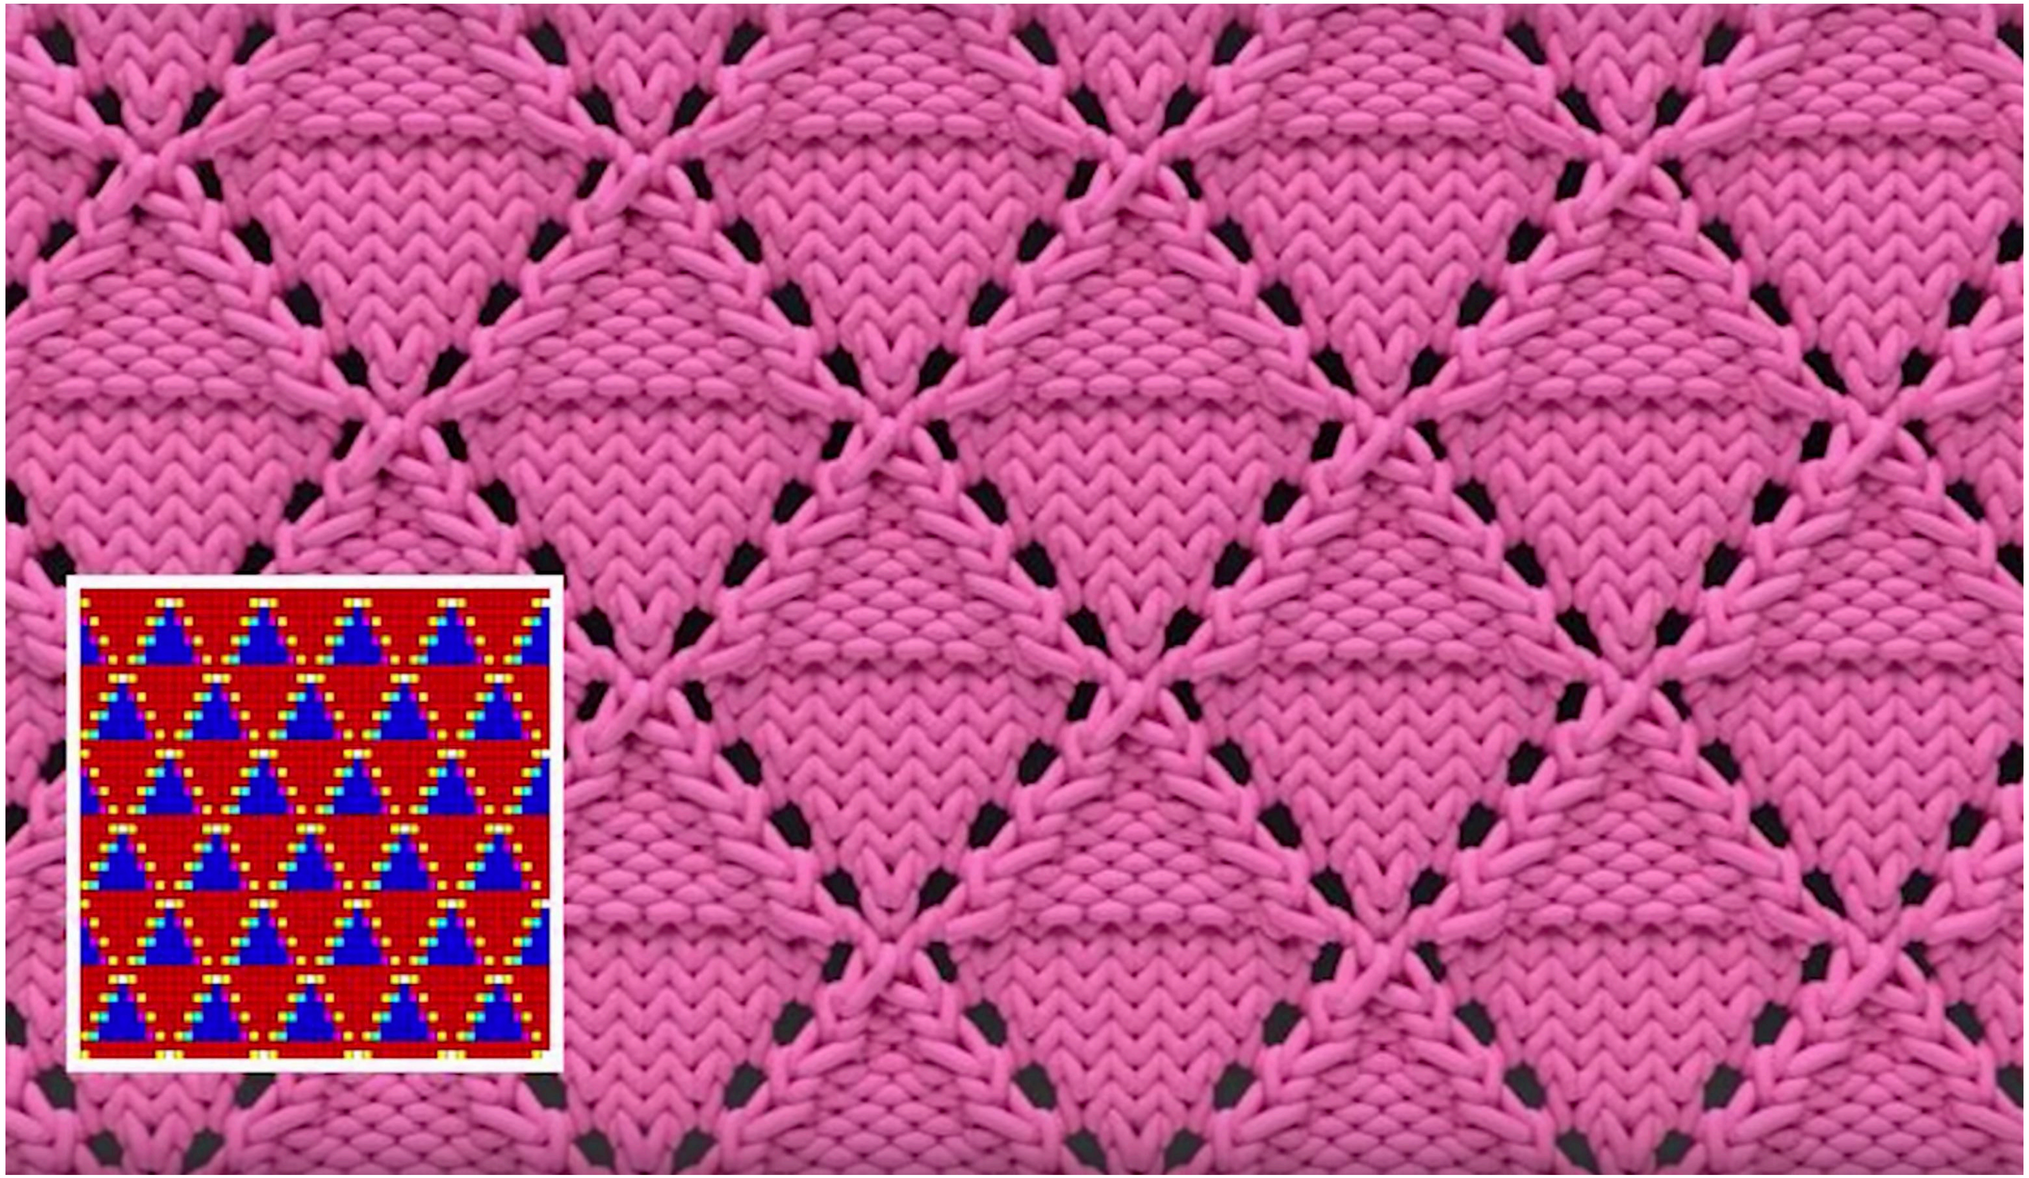

Leo Guibas

Geometry processing and analysis

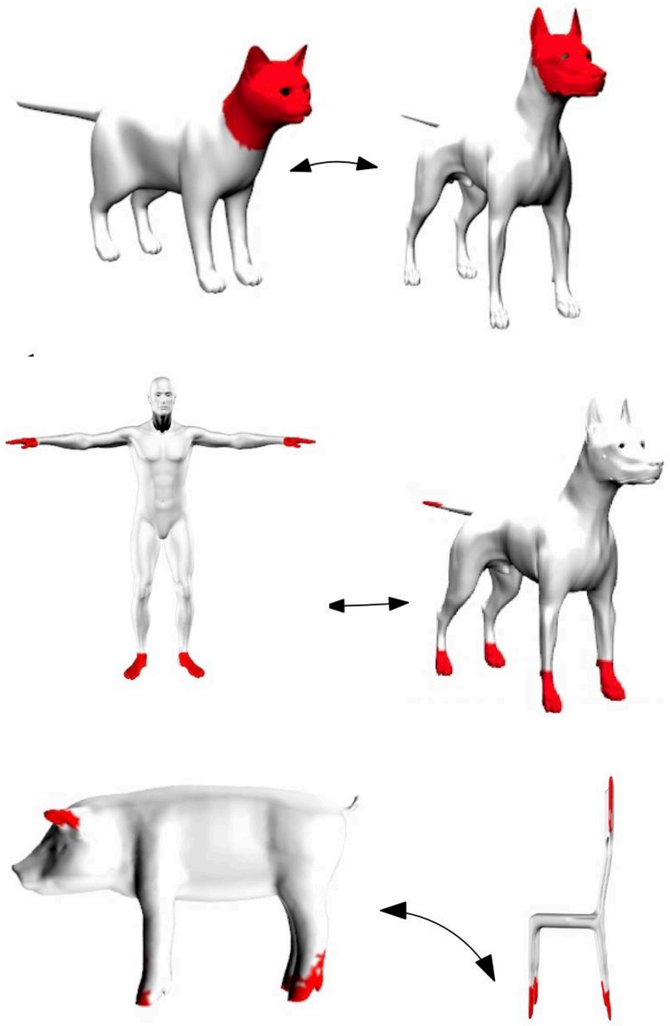

Shape Similarity and Correspondence

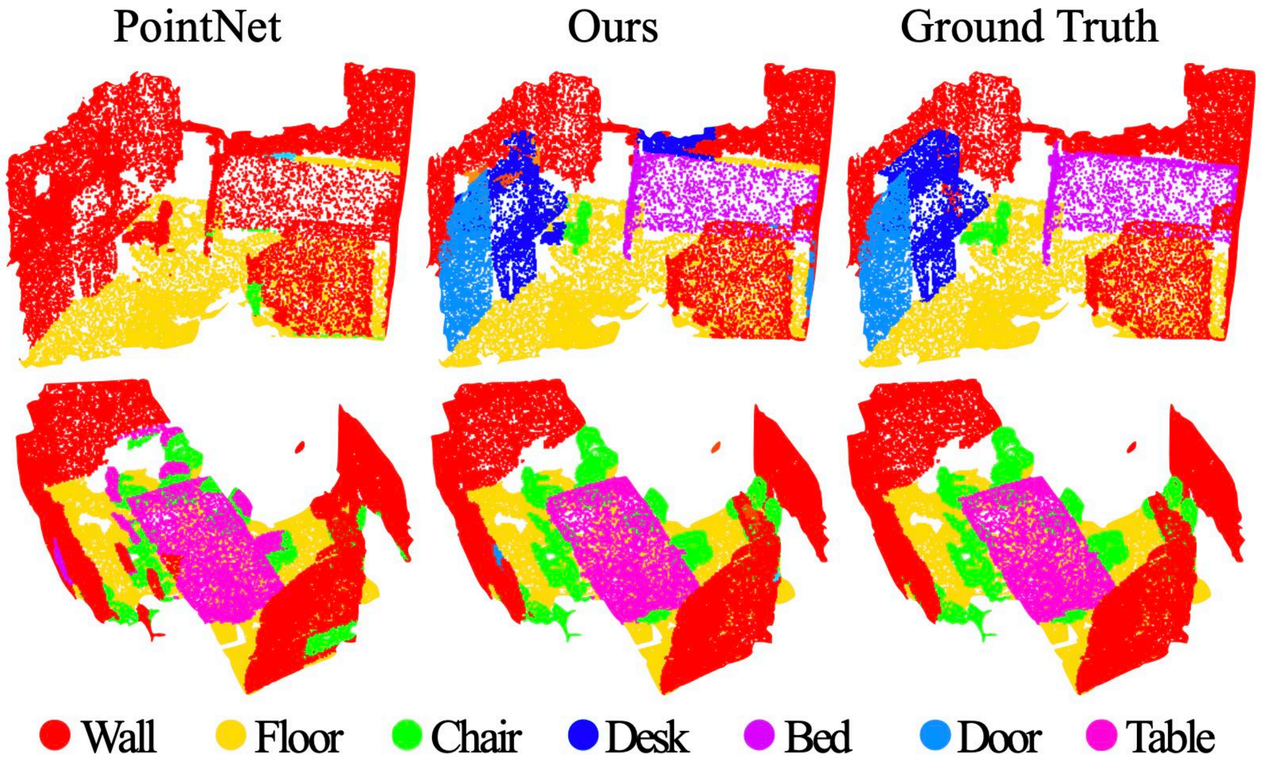

PointNet: Deep Learning on Point Clouds

Karen Liu

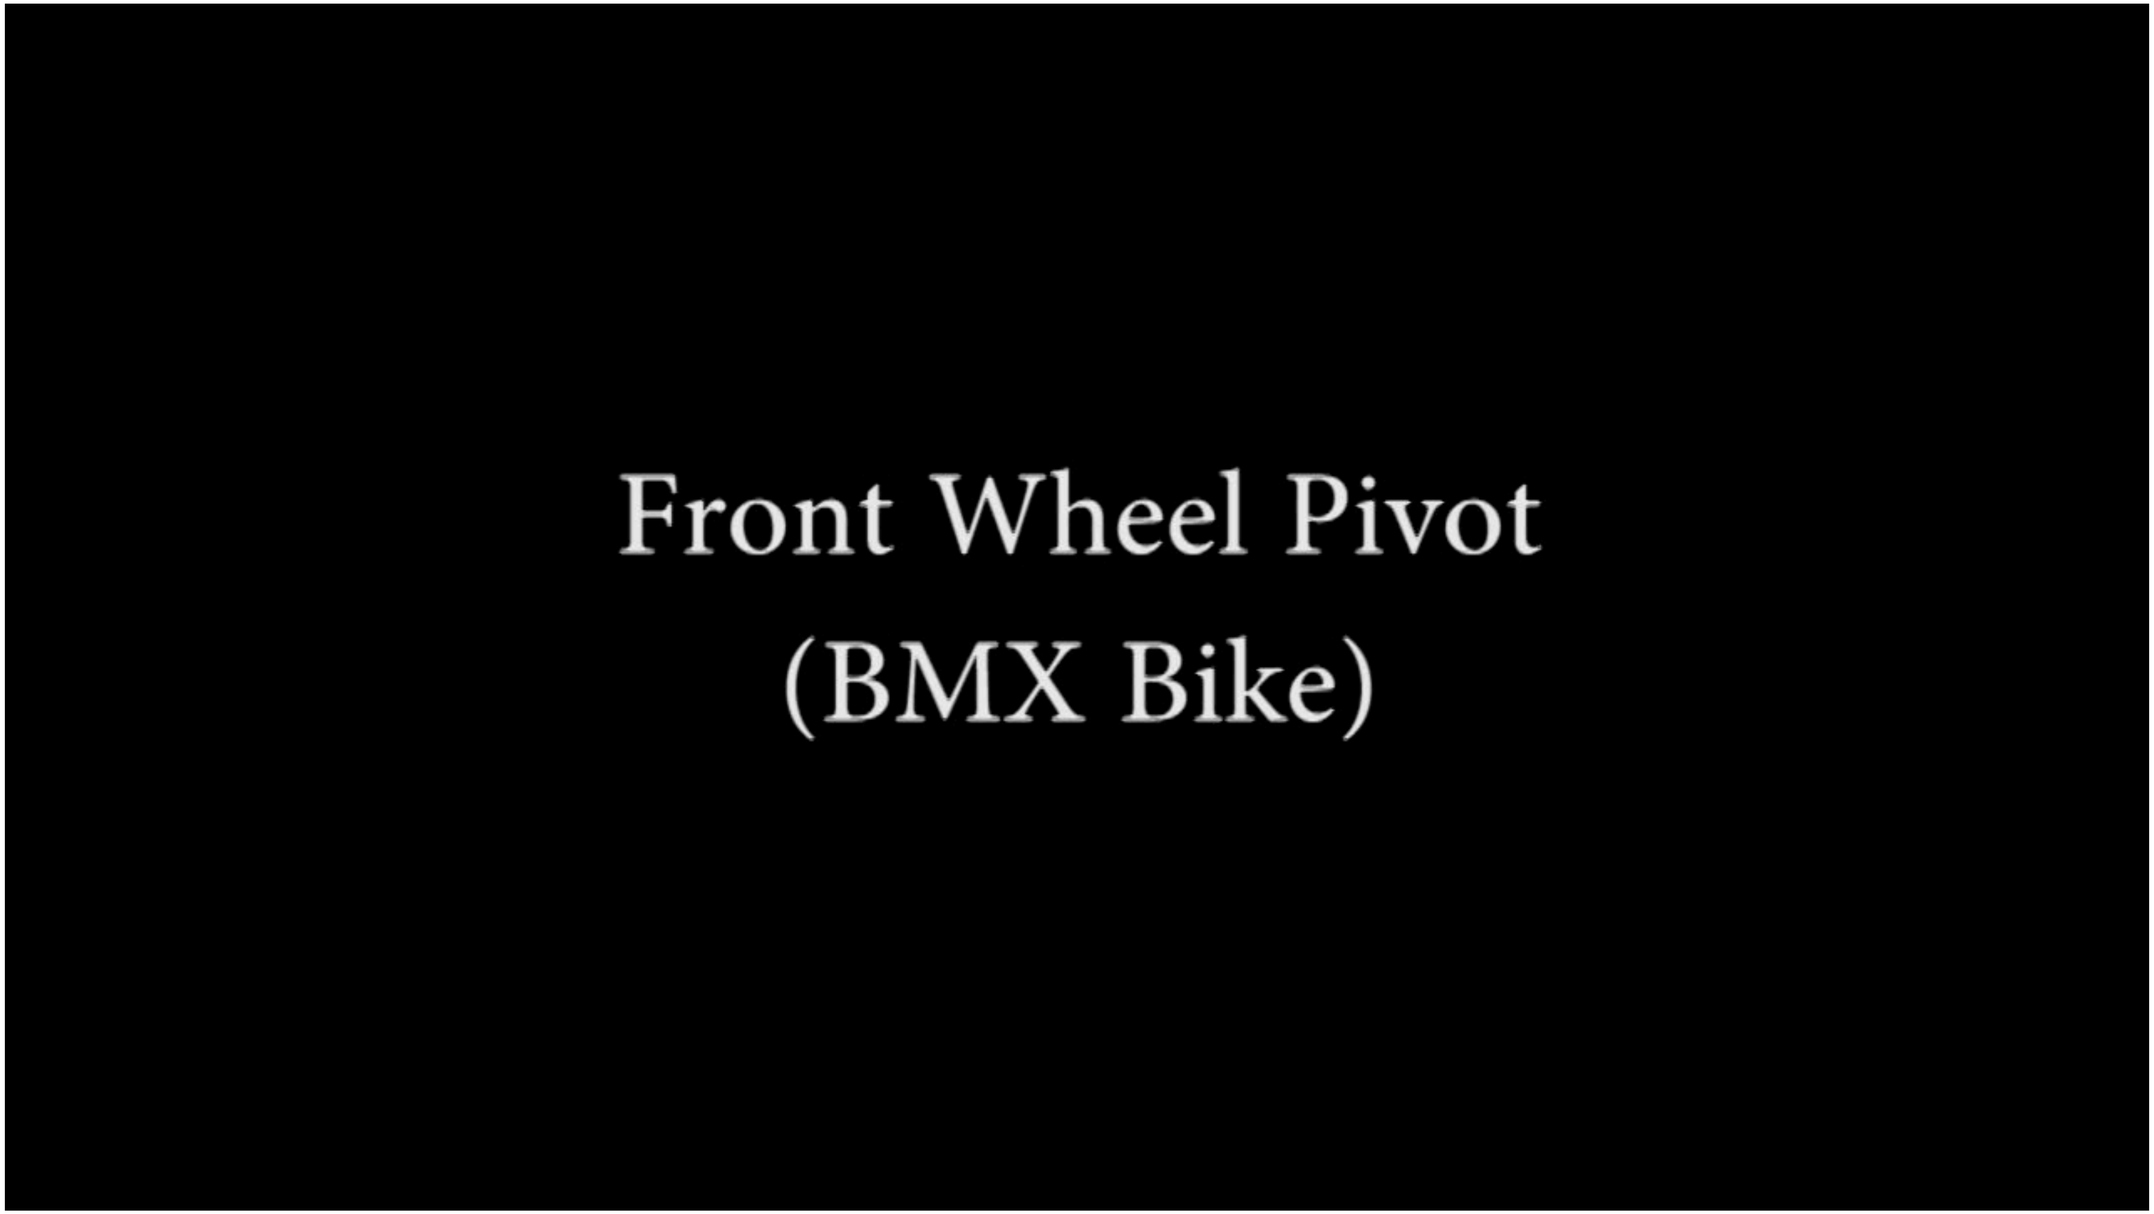

Gordon Wetzstein

Computational imaging and computational displays

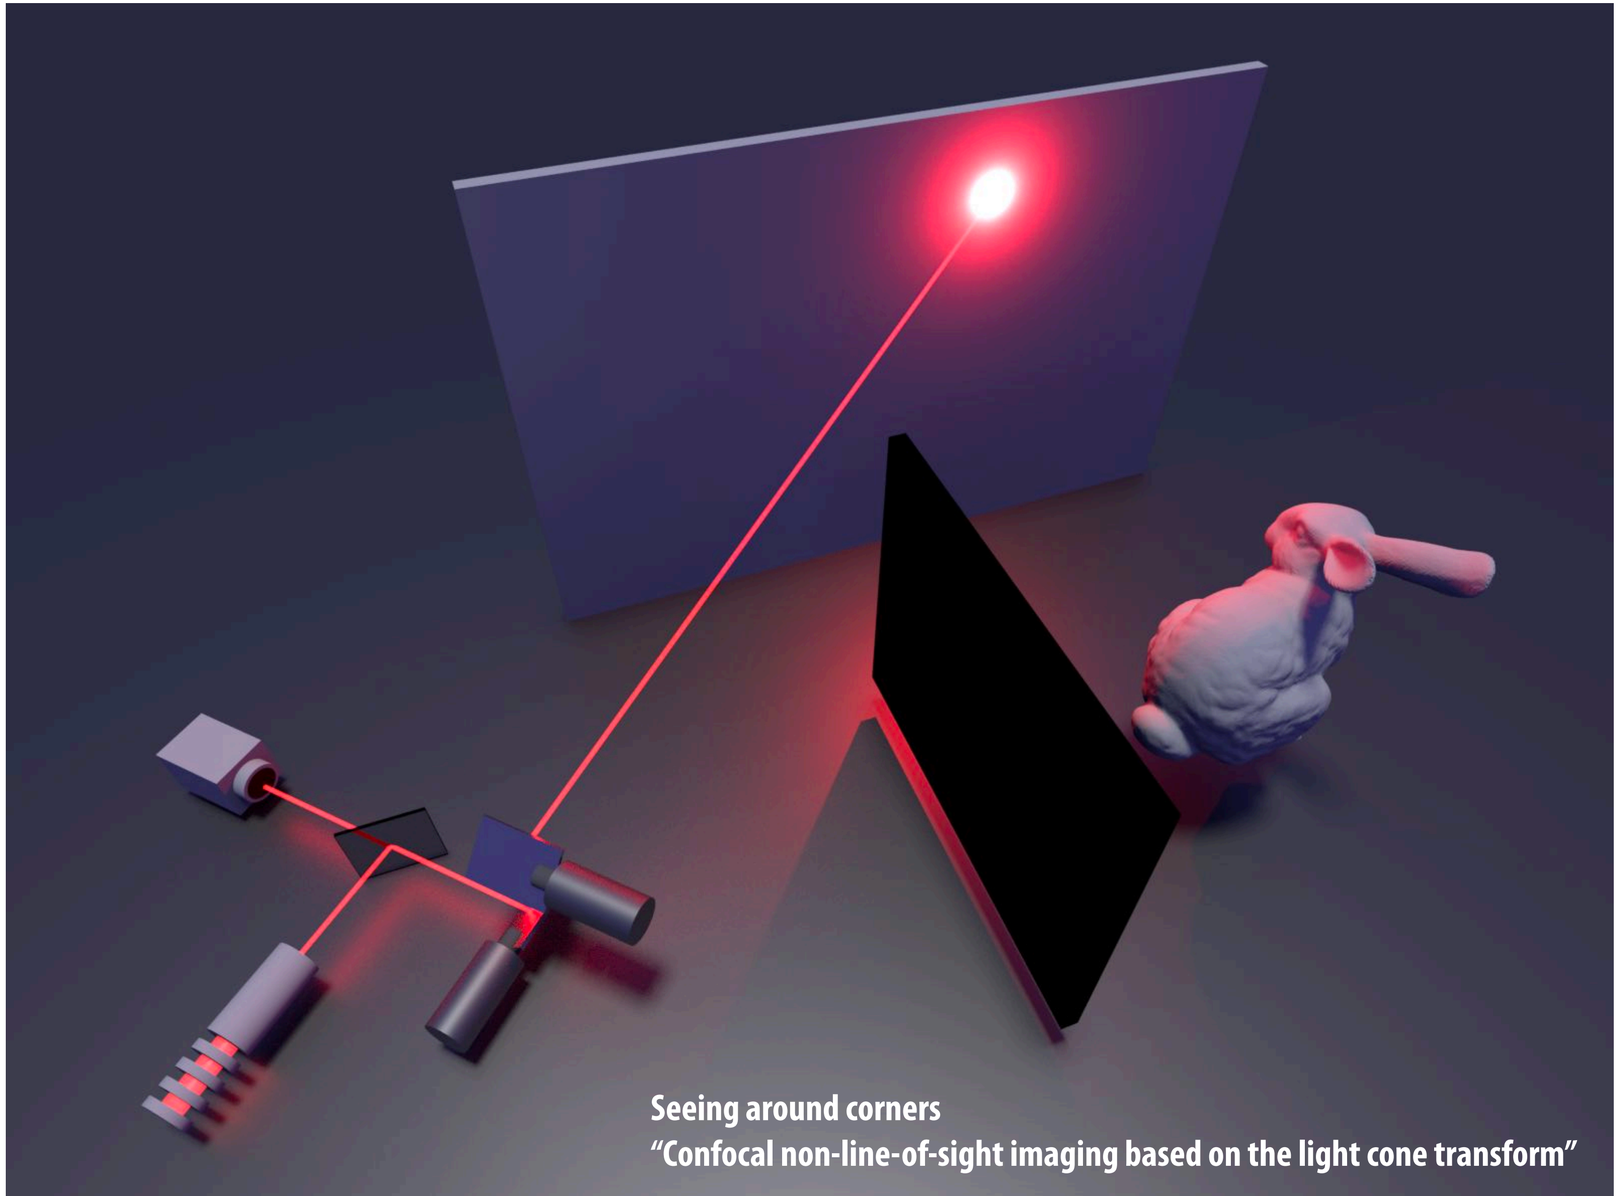

Seeing around corners
"Confocal non-line-of-sight imaging based on the light cone transform"

Gordon Wetzstein

Computational imaging and computational displays

"Hybrid Optical-Electronic Convolutional Neural Networks"

Using carefully designed optics to compute the early layers of a CNN prior to digital processing

Karen Liu

Interests in animation, simulation, and control

Kayvon Fatahalian

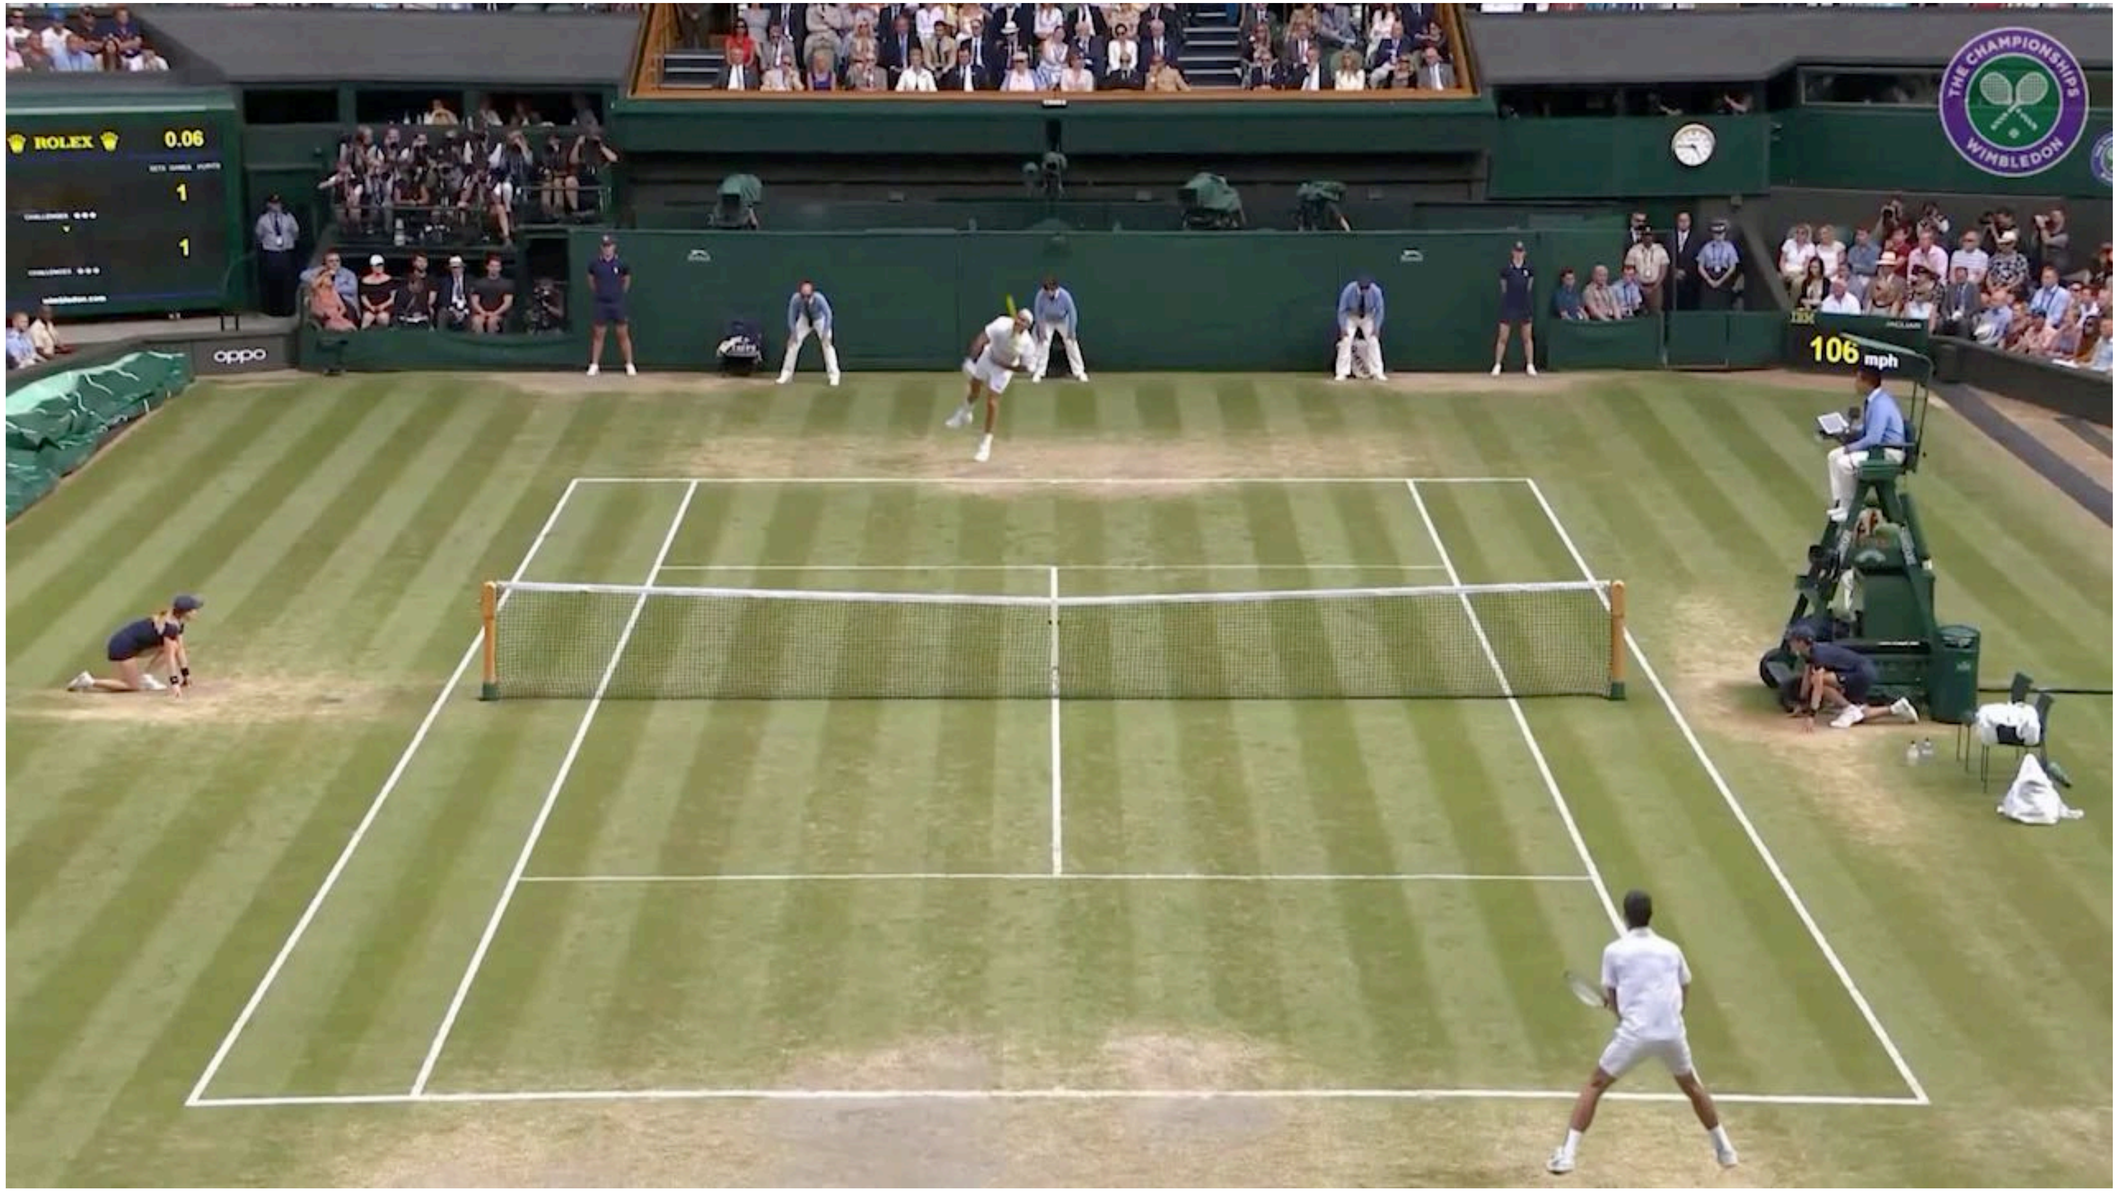

A completely computer generated Wimbledon point.

Kayvon Fatahalian

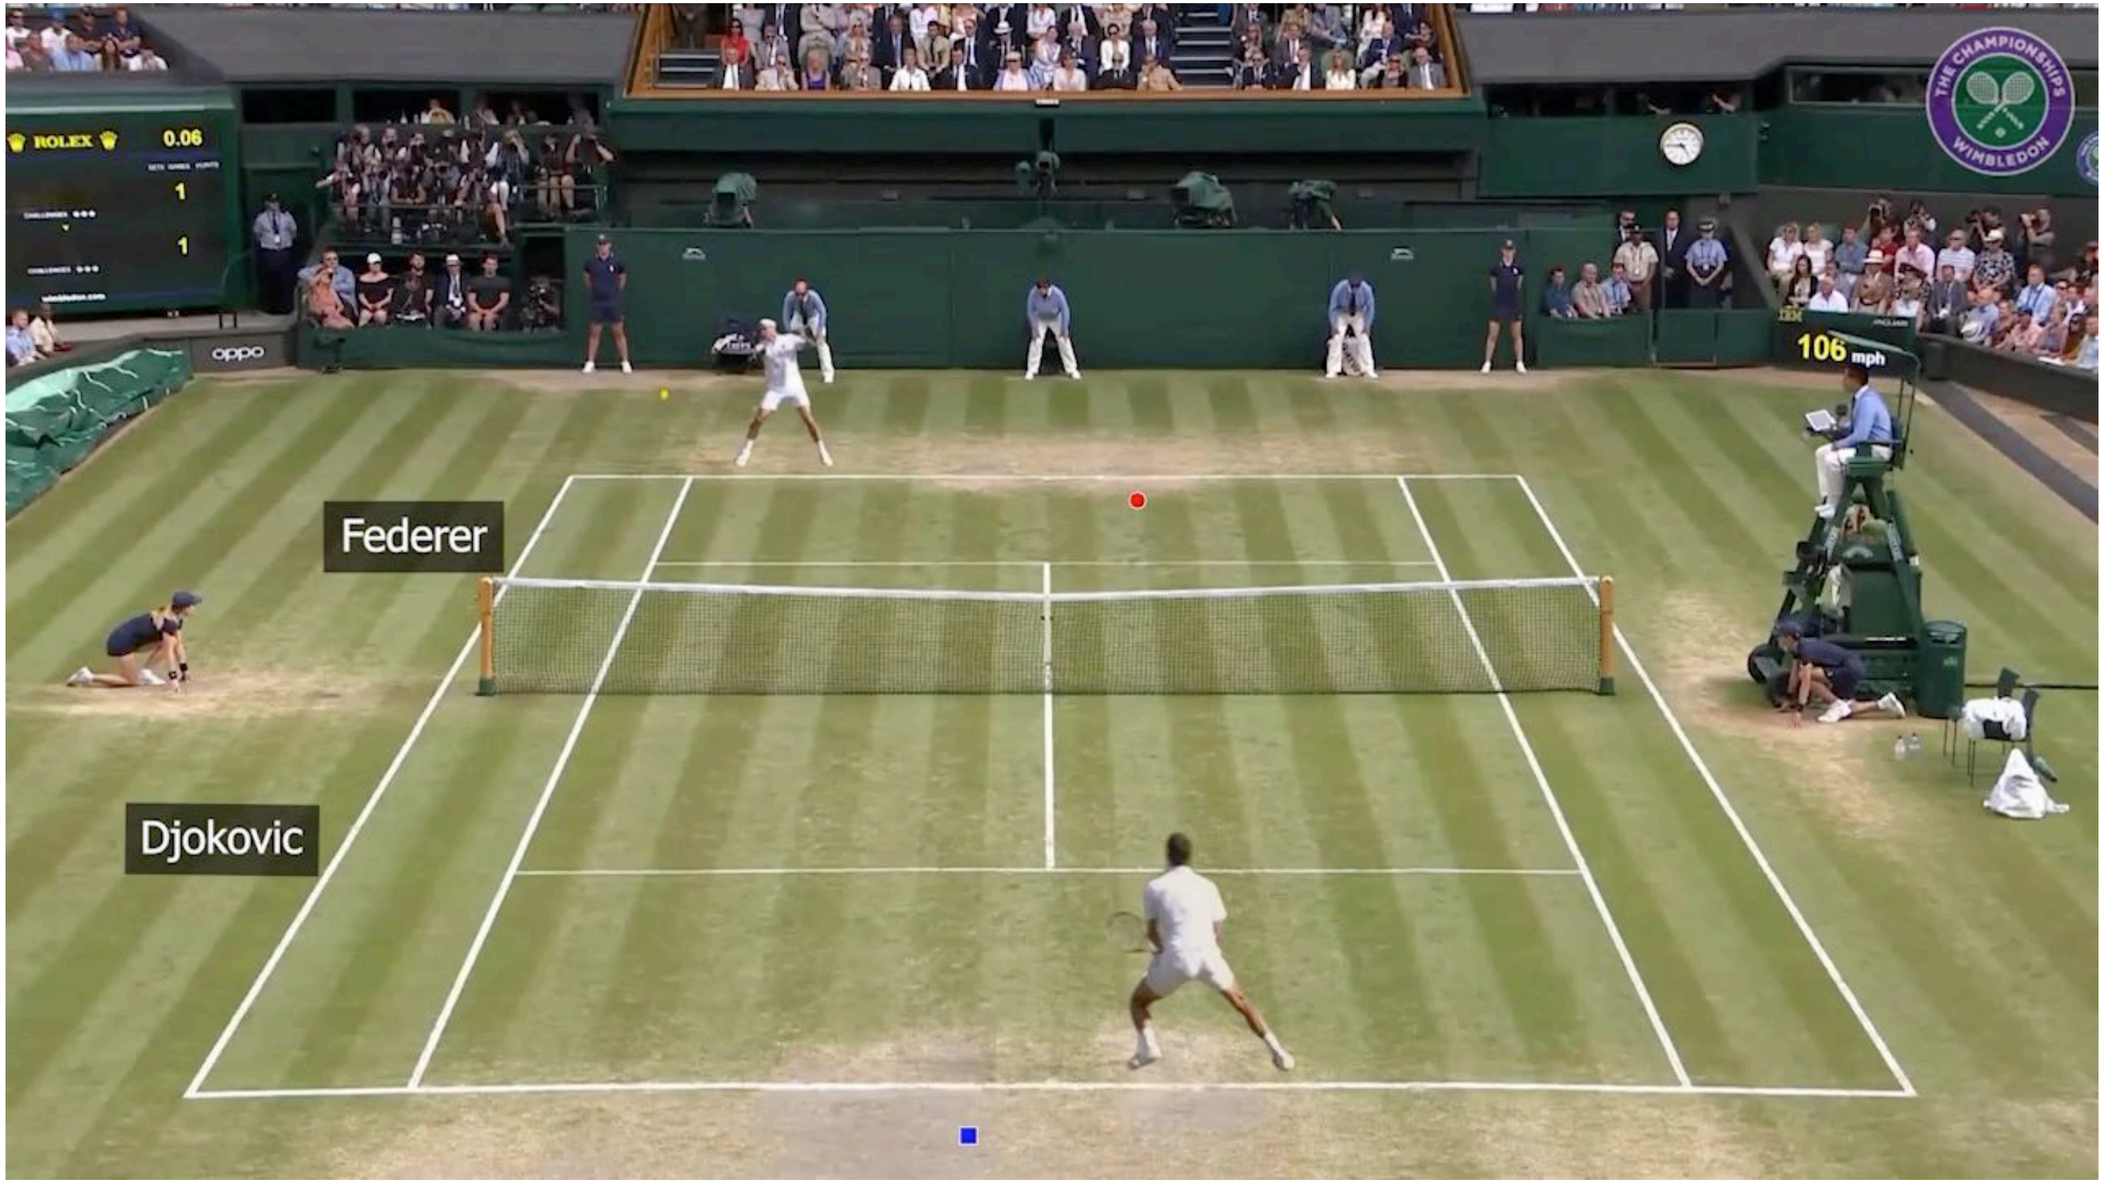

A completely computer generated Wimbledon point.

Example getting-started-in-the-lab project idea

- Take the technology shown on the previous slide and make an interactive "Wimbledon point creator" where folks can visit a web site, click where they want the ball to go, and then download a resulting video.
  - I bet it would be a hit if you could get it done and online before Wimbledon 2022

Ray tracing large scenes using 7,500 AWS lambda cores in the cloud

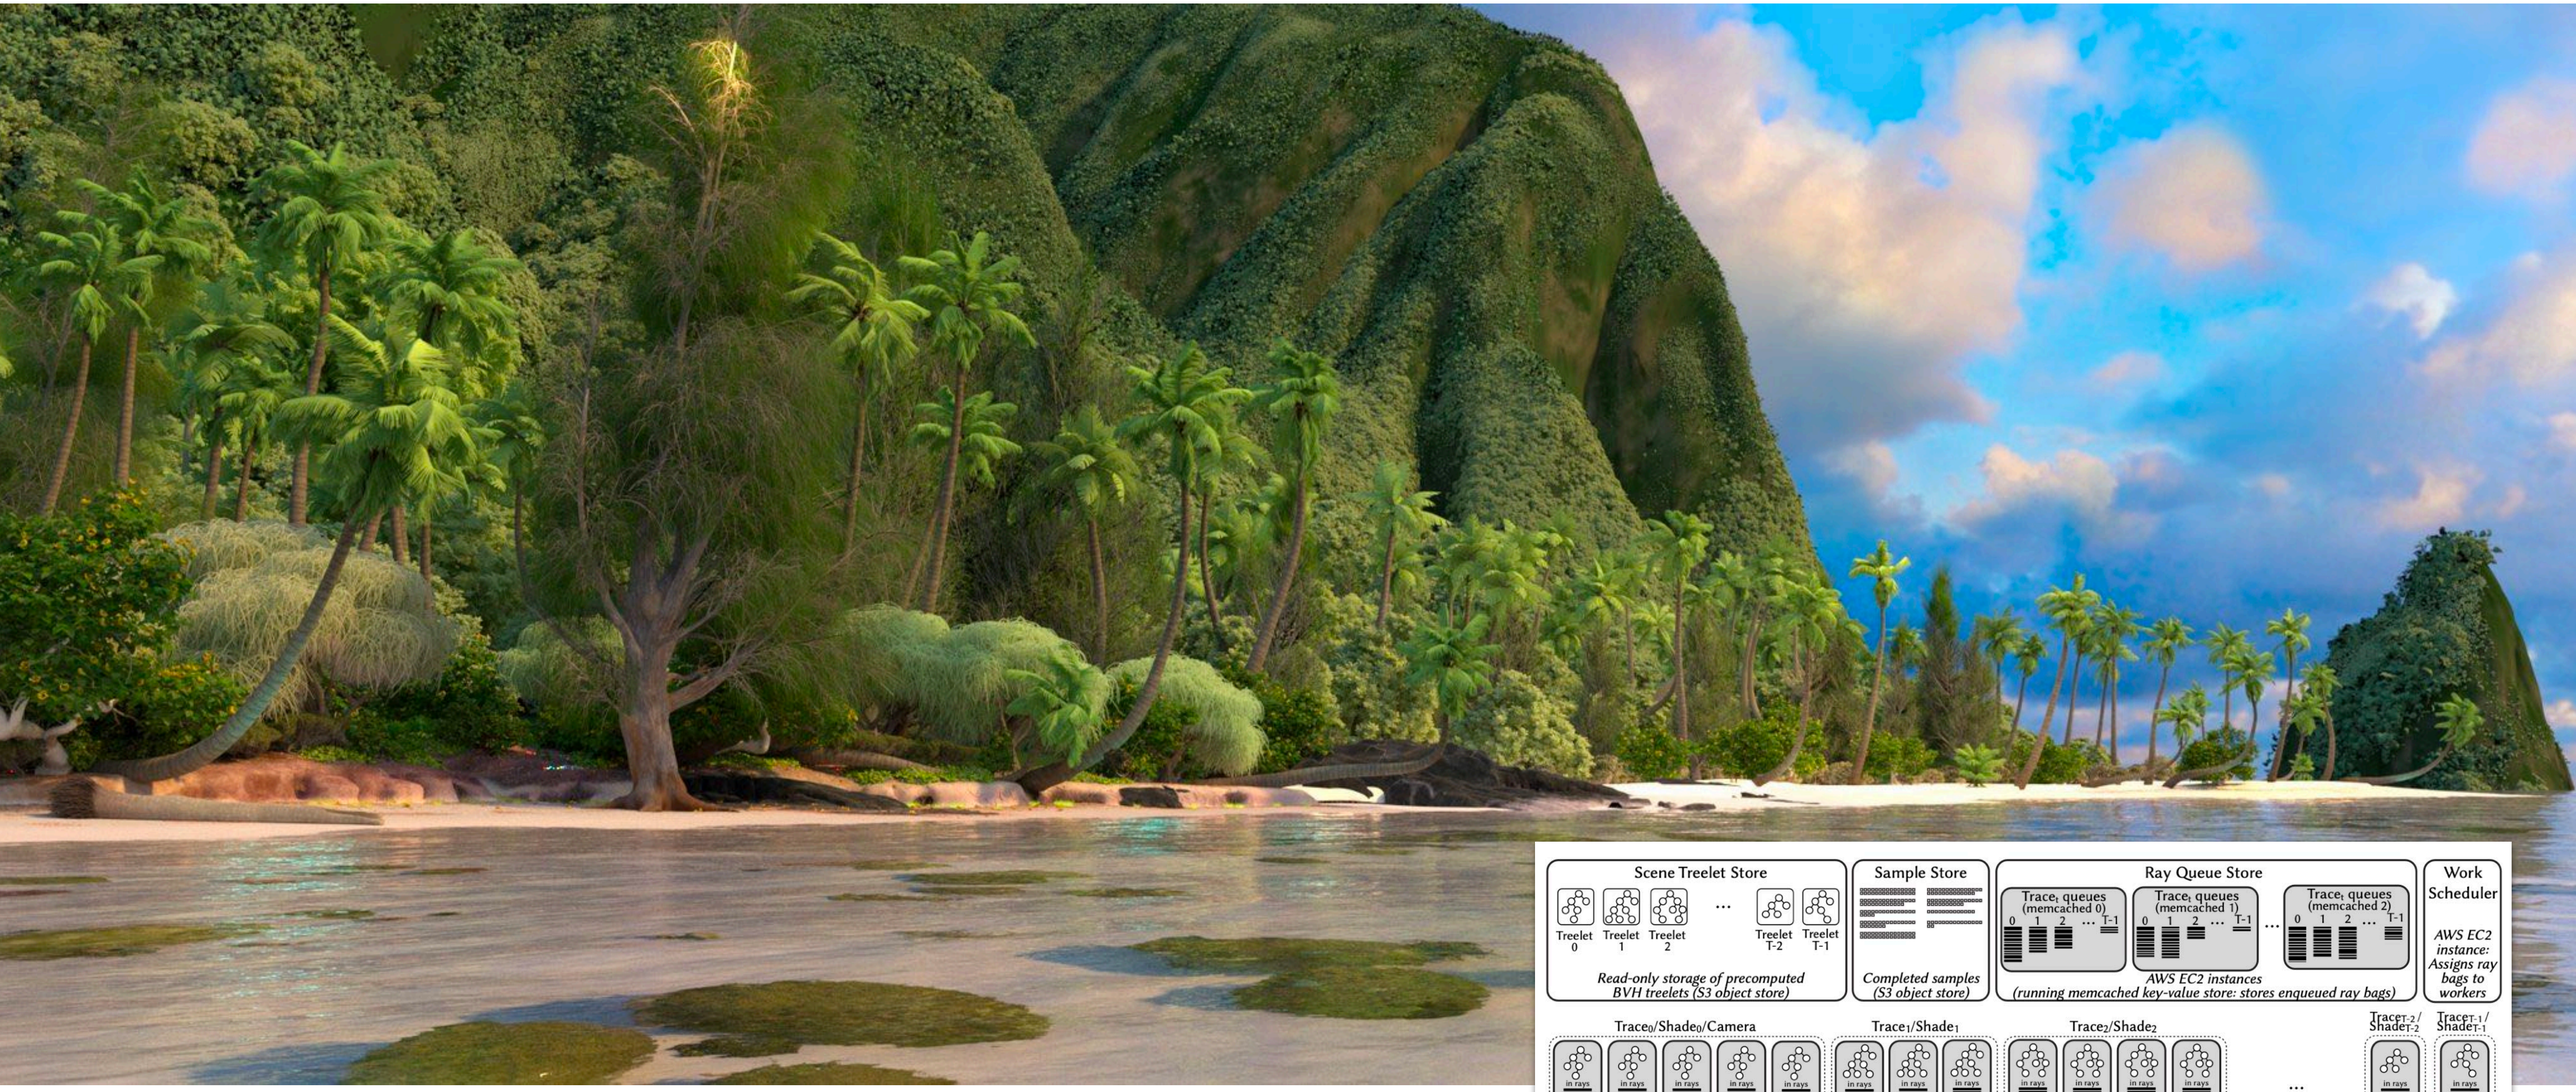

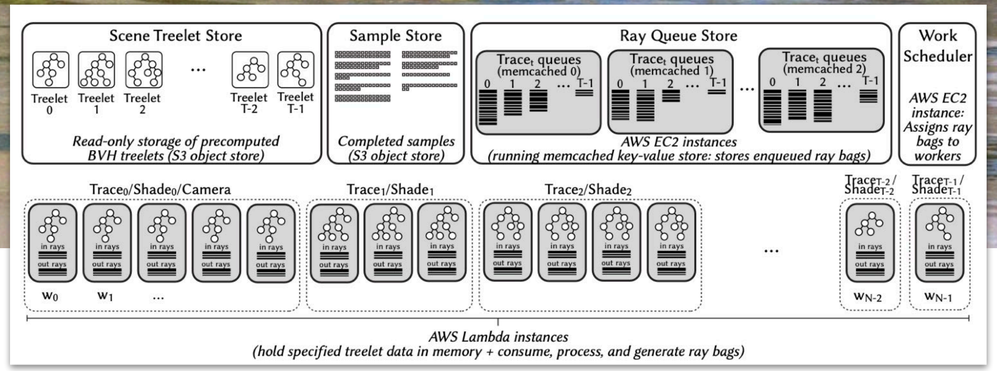

Can we redesign a game engine to achieve much higher performance for ML training workloads?

[Render → DNN inference → DNN train] in 3D scanned environments at 19,200 fps per GPU!
We are now interested in ray tracing simulated environments at similar rates!

Accurate lighting simulation

Example getting-started-in-the-lab project idea

- Work with my Ph.D. student to add a feature (new light type, load new type of geometry, etc) to our Vulcan renderer that's designed for ML use cases

- Other getting started ideas:
  - Write a CSGO bot (Counterstrike) that plays the game and generates training data for ML projects


Parastoo
Nava
she/her
Sharon
Mitchell
Jacob
Maneesh
Kayvon
Jordan
Christina
Terrell
Matthew

Virtual is a new medium.
Let's figure out virtual.

Highly customizable platform support: What are system abstractions and services that should be provided by: "Unity for live virtual and hybrid events"

What are effective social protocols in these environments?
(Both virtual and hybrid)

How can we create tools that help people design virtual interactions:
Example: automatically analyze the design of virtual spaces
to predict social awkwardness? Defy privacy expectations?

Other popular research topics in computer graphics…

Creating physically plausible models

- Via 3D printing, fabrication
- Creatures that locomotes, furniture that stands, etc.

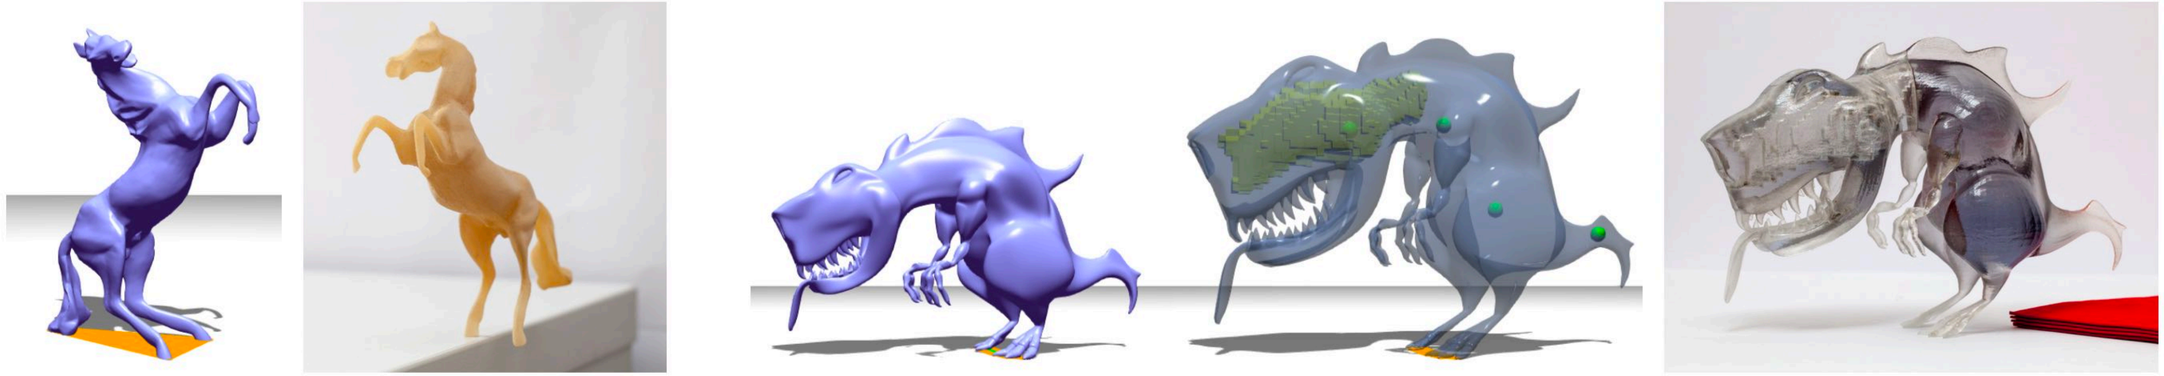

Fabricate models that are balanced to stand

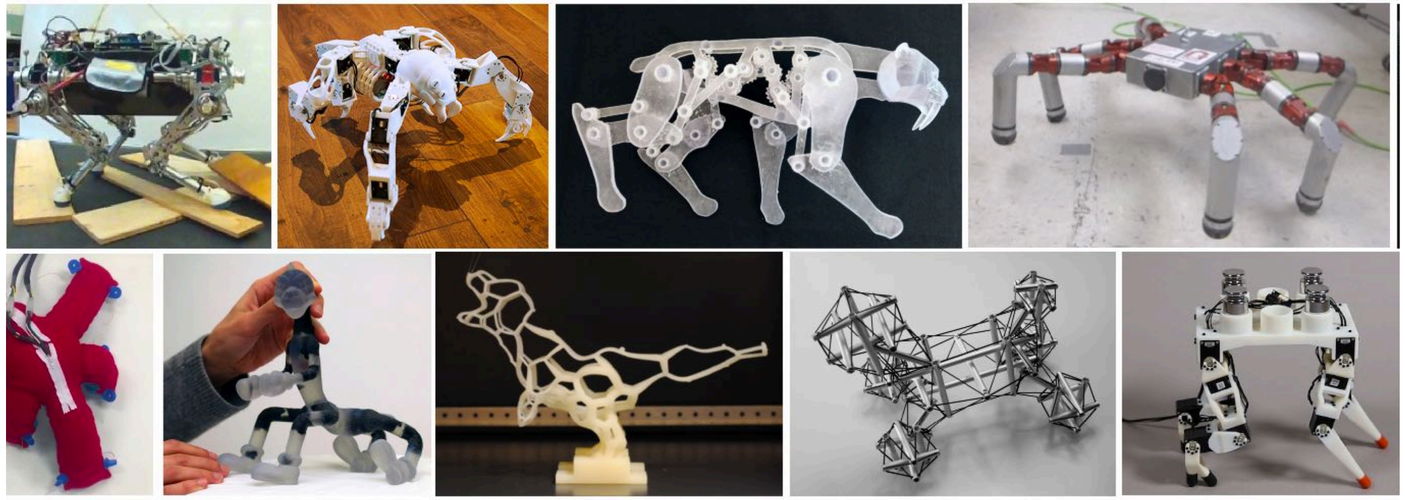

Fabricate robots that can balance and move

Computational photography

Using computation (and increasingly machine learning) to make more aesthetic photographs, simulate behavior of more complex lenses, etc.

Google Pixel 2 Portrait mode

Image credit: Google / Matt Jones (https://ai.googleblog.com/2017/10/portrait-mode-on-pixel-2-and-pixel-2-xl.html)

Advanced geometry processing

Fundamental questions about alignment, similarly, symmetry, etc…

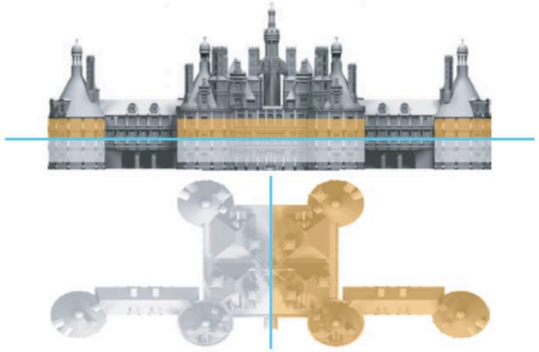

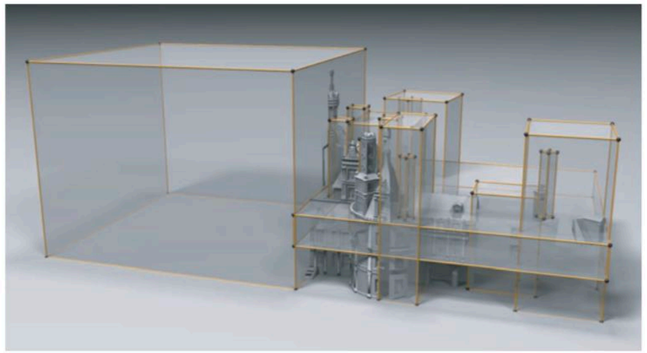

Advanced displays/rendering for VR/AR

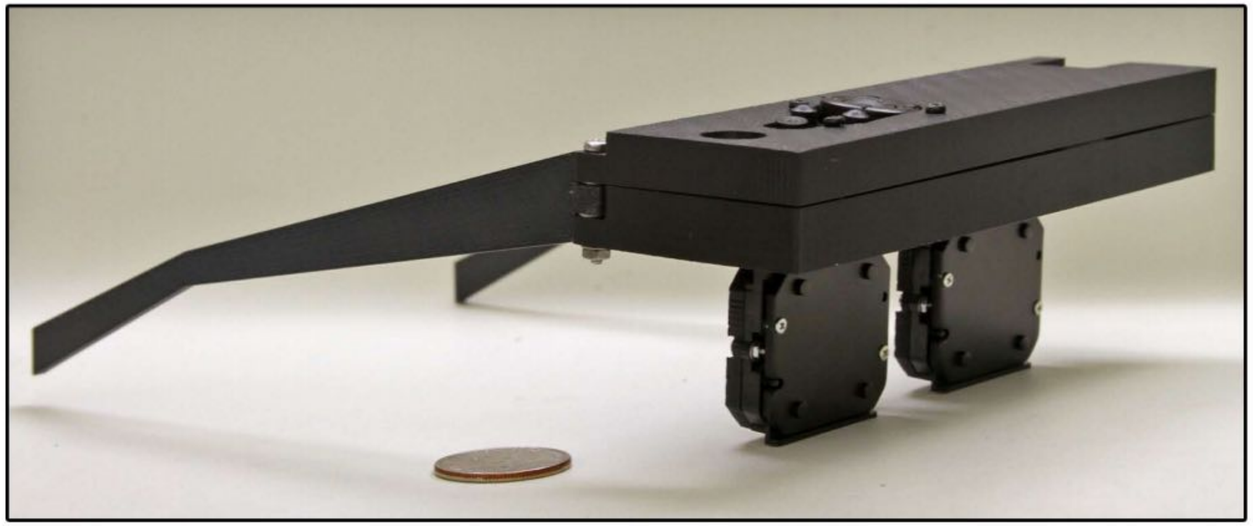

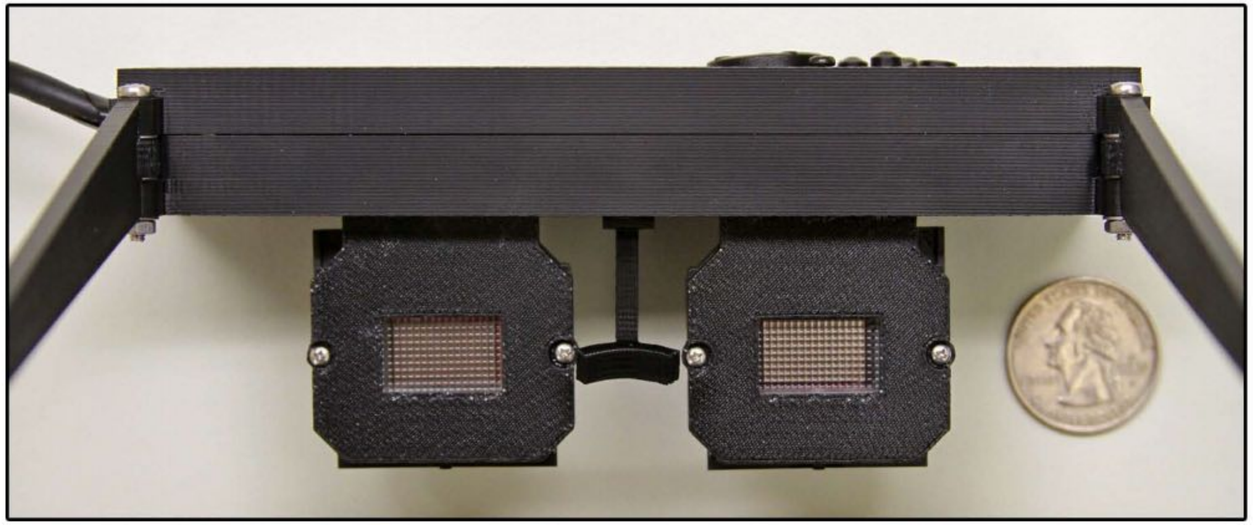

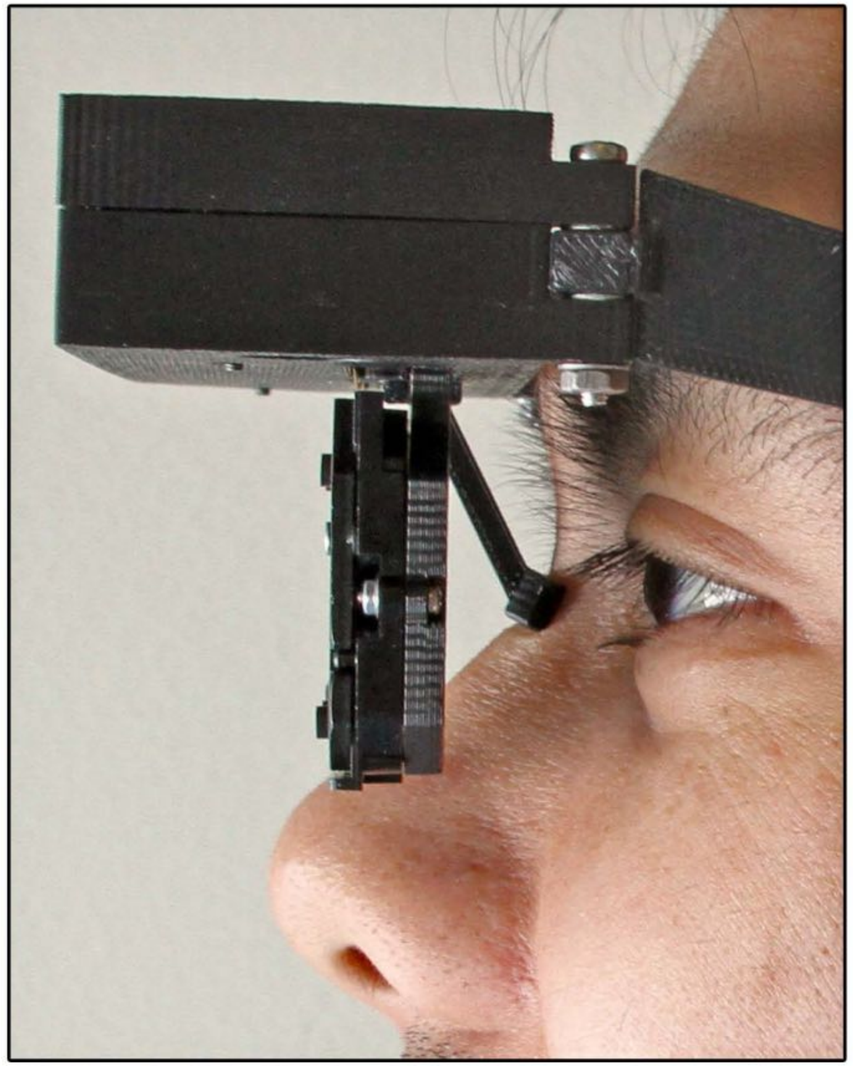

Near eye light field display

AI for lower cost content creation and capture

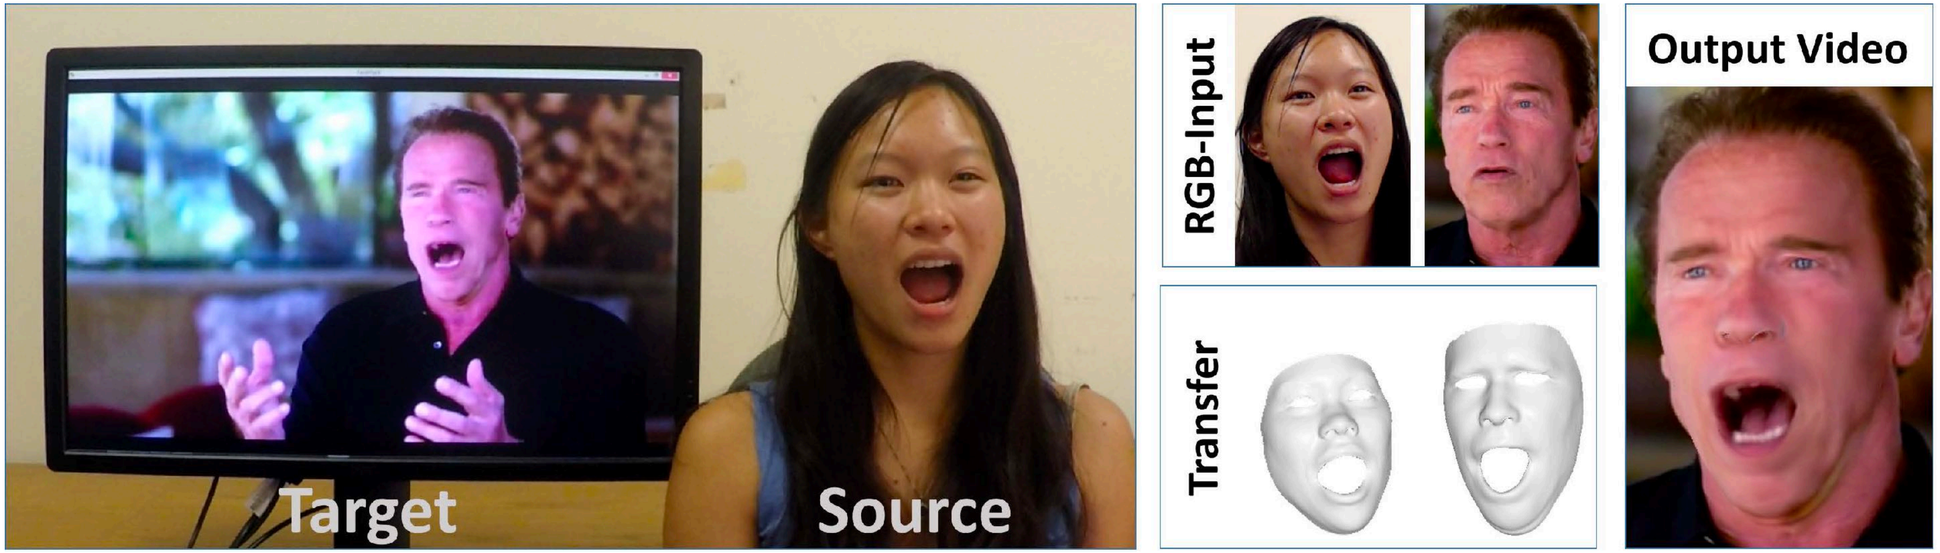

Manipulating actors by performance capture

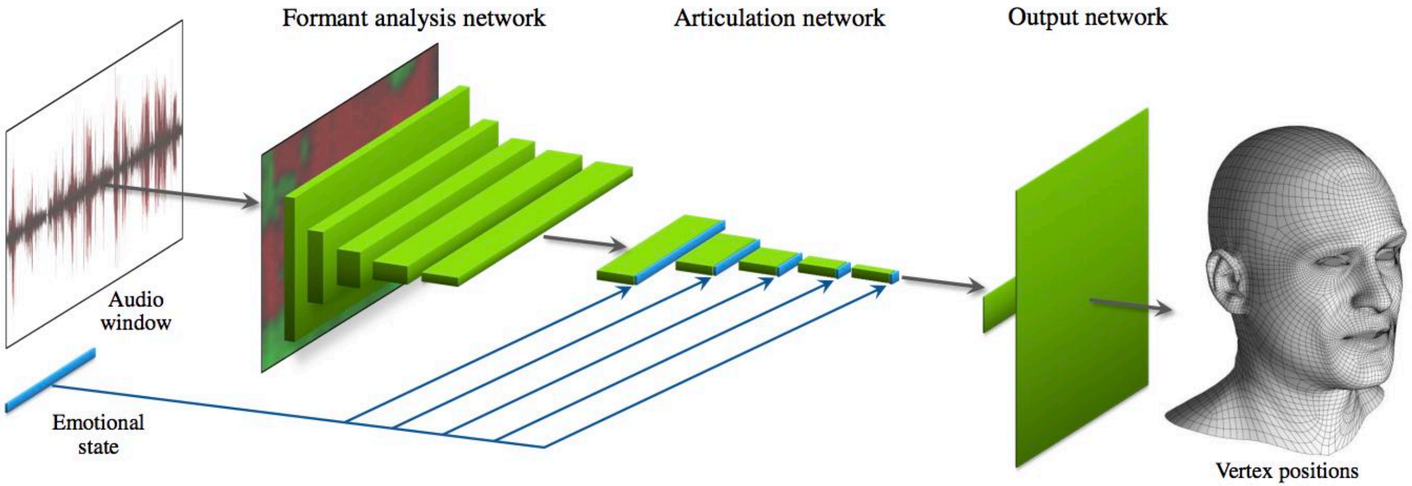

Audio input to mesh animation

The other direction: graphics helping machine learning

Input labels

Synthesized image

Pix2pixHD

Grand Theft Auto Screenshots

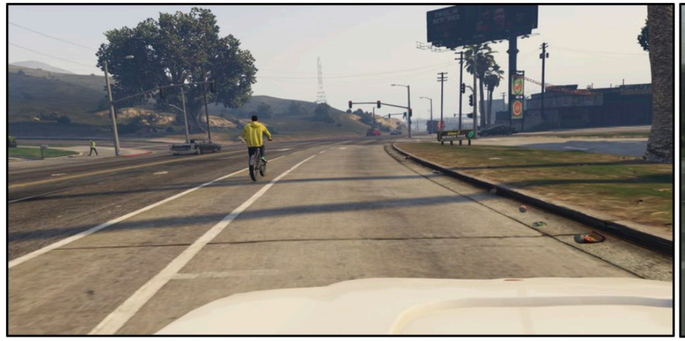

Synthesized "photorealistic" image

A fun resource

Ke-sen Huang's famous site with all the SIGGRAPH papers!

http://kesen.realtimerendering.com/

SIGGRAPH 2020 papers on the web

Page maintained by Ke-Sen Huang. If you have additions or changes, send an e-mail.

Information here is provided with the permission of the ACM

Note that when possible I link to the page containing the link to the actual PDF or PS of the preprint. I prefer this as it gives some context to the paper and avoids possible copyright problems with direct linking. Thus you may need to search on to find the actual document.

ACM Digital Library: ACM Transactions on Graphics (TOG) Volume 39, Issue 4 (July 2020) Proceedings of ACM SIGGRAPH 2020

ACM Digital Library (DOI) Link for the paper **Paper Abstract** **Author Preprint** **Paper Video**

Paper Presentation **Paper Images** **Paper Data** **Demo Program or Source Code** **Related Links**

Changelog

ANIMATION/SIMULATION

A Level-Set Method for Magnetic Substance Simulation
Xingyu Ni (CFCS, Peking University and AICFVE, Beijing Film Academy), Bo Zhu (Dartmouth College), Bin Wang (AICFVE, Beijing Film Academy), Baoquan Chen (CFCS, Peking University and AICFVE, Beijing Film Academy)

A Massively Parallel and Scalable Multi-GPU Material Point Method
Xinlei Wang* (Zhejiang University and University of Pennsylvania), Yuxing Qiu* (University of California, Los Angeles and University of Pennsylvania), Stuart Slattery (Oak Ridge National Laboratory), Yu Fang, Minchen Li (Universit Pennsylvania), Song-Chun Zhu, Yixin Zhu (University of California, Los Angeles), Min Tang (Zhejiang University), Dinesh Manocha (University of Maryland), Chenfanfu Jiang (University of Pennsylvania) (*equal contribution)

A Model for Soap Film Dynamics with Evolving Thickness
Sadashige Ishida, Peter Synak* (IST Austria), Fumiya Narita (Unaffiliated), Toshiya Hachisuka (The University of Tokyo), Chris Wojtan (IST Austria) (*joint first authors)

A Practical Octree Liquid Simulator With Adaptive Surface Resolution
Ryoichi Ando (National Institute of Informatics), Christopher Batty (University of Waterloo)

A Scalable Approach to Control Diverse Behaviors for Physically Simulated Characters
Jungdam Won, Deepak Gopinath, Jessica Hodgins (Facebook AI Research)

A System for Efficient 3D Printed Stop-Motion Face Animation **(TOG Paper)**
Rinat Abdrashitov, Alec Jacobson, Karan Singh (University of Toronto)

Accurate Face Rig Approximation With Deep Differential Subspace Reconstruction
Steven L. Song* (Blue Sky Studios), Weiqi Shi* (Yale University), Michael Reed (Blue Sky Studios) (*Authors contributed equally)

Adaptive Merging for Rigid Body Simulation
Eulalie Coevoet, Otman Benchekroun, Paul G. Kry (McGill University)

An Implicit Compressible SPH Solver for Snow Simulation
Christoph Gissler (University of Freiburg and FIFTY2 Technology), Andreas Henne (FIFTY2 Technology), Stefan Band (University of Freiburg), Andreas Peer (FIFTY2 Technology), Matthias Teschner (University of Freiburg)

AnisoMPM: Animating Anisotropic Damage Mechanics
Joshuah Wolper, Yunuo Chen, Minchen Li, Yu Fang, Ziyin Qu, Jiecong Lu, Meggie Cheng, Chenfanfu Jiang (University of Pennsylvania)

Capturing Subjective First-Person View Shots With Drones for Automated Cinematography
Amirsaman Ashtari (KAIST), Stefan Stevsic (ETH Zurich), Tobias Nageli (ETH Zurich and Tinamu Labs), Otmar Hilliges (ETH Zurich), Jean-Charles Bazin (KAIST)

CARL: Controllable Agent with Reinforcement Learning for Quadruped Locomotion
Ying-Sheng Luo*, Jonathan Hans Soeseno*, Trista Pei-Chun Chen (Inventec Corp.), Wei-Chao Chen (Inventec Corp. and Skywatch Innovation Inc.) (*Joint first authors)

Discussion: graphics jobs

Discussion: how to get involved in graphics at Stanford

- Email your graphics professors and ask to talk to them about independent study
  - Although to be honest… the best intro line is ("I took and loved your 300-level graphics class and did well and want to keep going)

- A common way to get started
  - Hack code to contribute to a Ph.D. student's research project

Why research (or independent study)?

- You will learn <u>way more</u> about a topic than in any class.
- You think your undergrad friends are very smart? Come hang out with Stanford Ph.D. students! (you get to work side-by-side with them and with faculty). Imagine what level you might rise to.
- It's way more fun to be on the cutting edge. Industry might not even know about what you are working on. (imagine how much more valuable you are if you can teach them)
- It widens your mind as to what is possible.

Maybe you might like research and decide you want to go to grad school

Pragmatic comment: Without question, the number one way to get into a top grad school is to receive a strong letter of recommendation from faculty members. You get that letter only from being part of a research team for an extended period of time.

DWIC letter: ("did well in class" letter) What you get when you ask for a letter from a faculty member who you didn't do research with, but got an 'A' in their class. This letter is essentially thrown out by the Ph.D. admissions committee at good schools.

A very good reference

CMU Professor Mor Harchol-Balter's writeup:

"Applying to Ph.D. Programs in Computer Science"

http://www.cs.cmu.edu/~harchol/gradschooltalk.pdf

Thanks for being a great class!

Good luck finishing projects.

Make sure you have fun, that's the point!

And, above all else, do your best to stay healthy, and keep others healthy.