Lecture 19:

Course Summary + Graphics at Stanford Today

Interactive Computer Graphics
Stanford CS248, Winter 2019

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 (Hanrahan)

EE267: "Virtual Reality", focuses on display, tracking hardware for VR (Wetzstein)

CS448V: "Computational Video Manipulation," seminar on recent research results on manipulating video (Agrawala)

CS448M: "Making Making Machines for Makers", exploration of systems for making things using computer-aided design and manufacturing (Hanrahan, James)

FALL

CS348K: "Visual Computing Systems", principles of creating efficient parallel systems for computational photography, 3D graphics, and deep learning for vision (Fatahalian)

CS146: "Computer Game Design", make your own games in Unity (James)

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)

CS 205L: "Continuous Mathematical Methods with an Emphasis on Machine Learning" (Fedkiw)

*** red = new courses**

CS348B (Spring, Hanrahan)

- 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

Ron Fedkiw

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

Ron Fedkiw

Segmentation Masks
of Stereo Footage

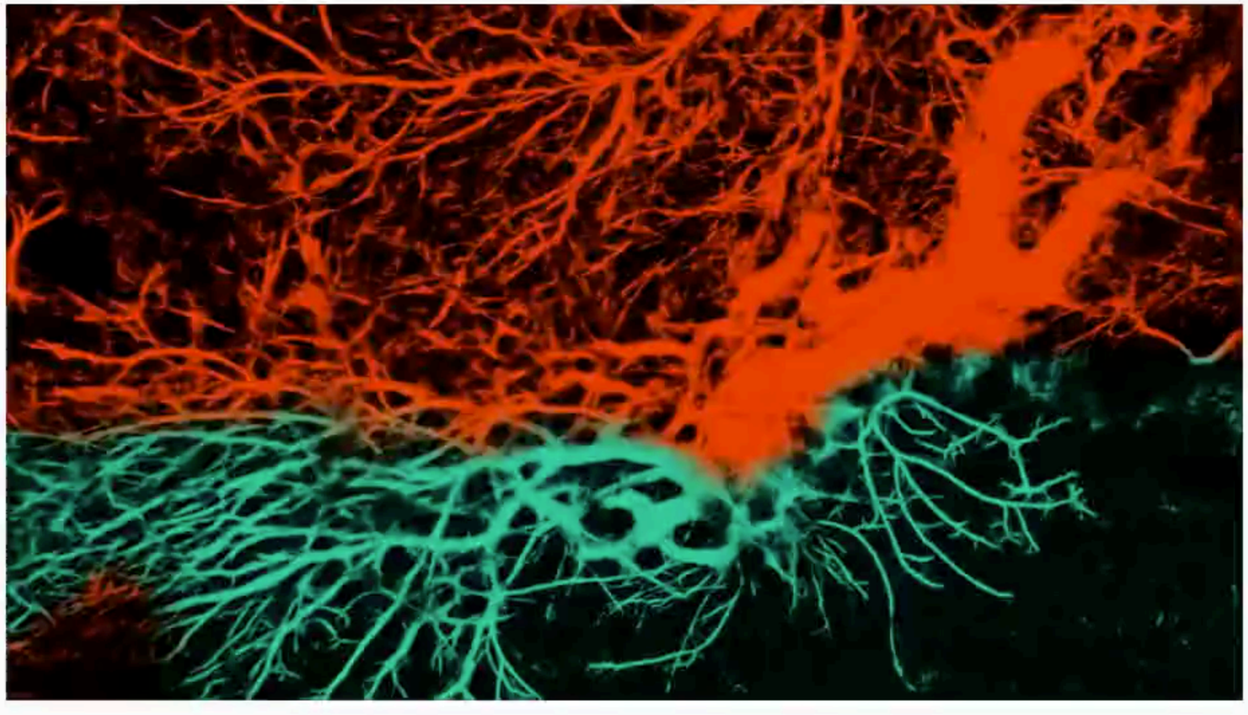

Left Camera

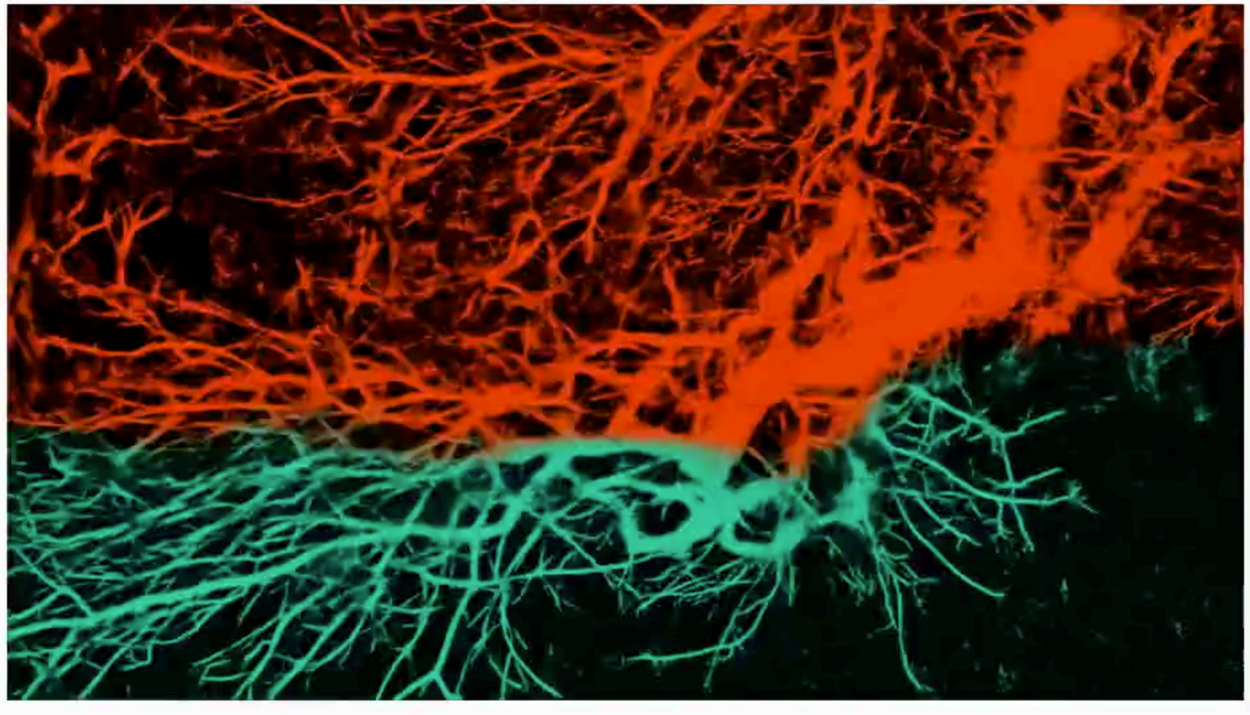

Right Camera

Ron Fedkiw

Maneesh Agrawala

- Many current projects on video editing and manipulation

Visual Rhythm and Beat [Davis et al.]

Maneesh Agrawala

- Many current projects on video editing and manipulation

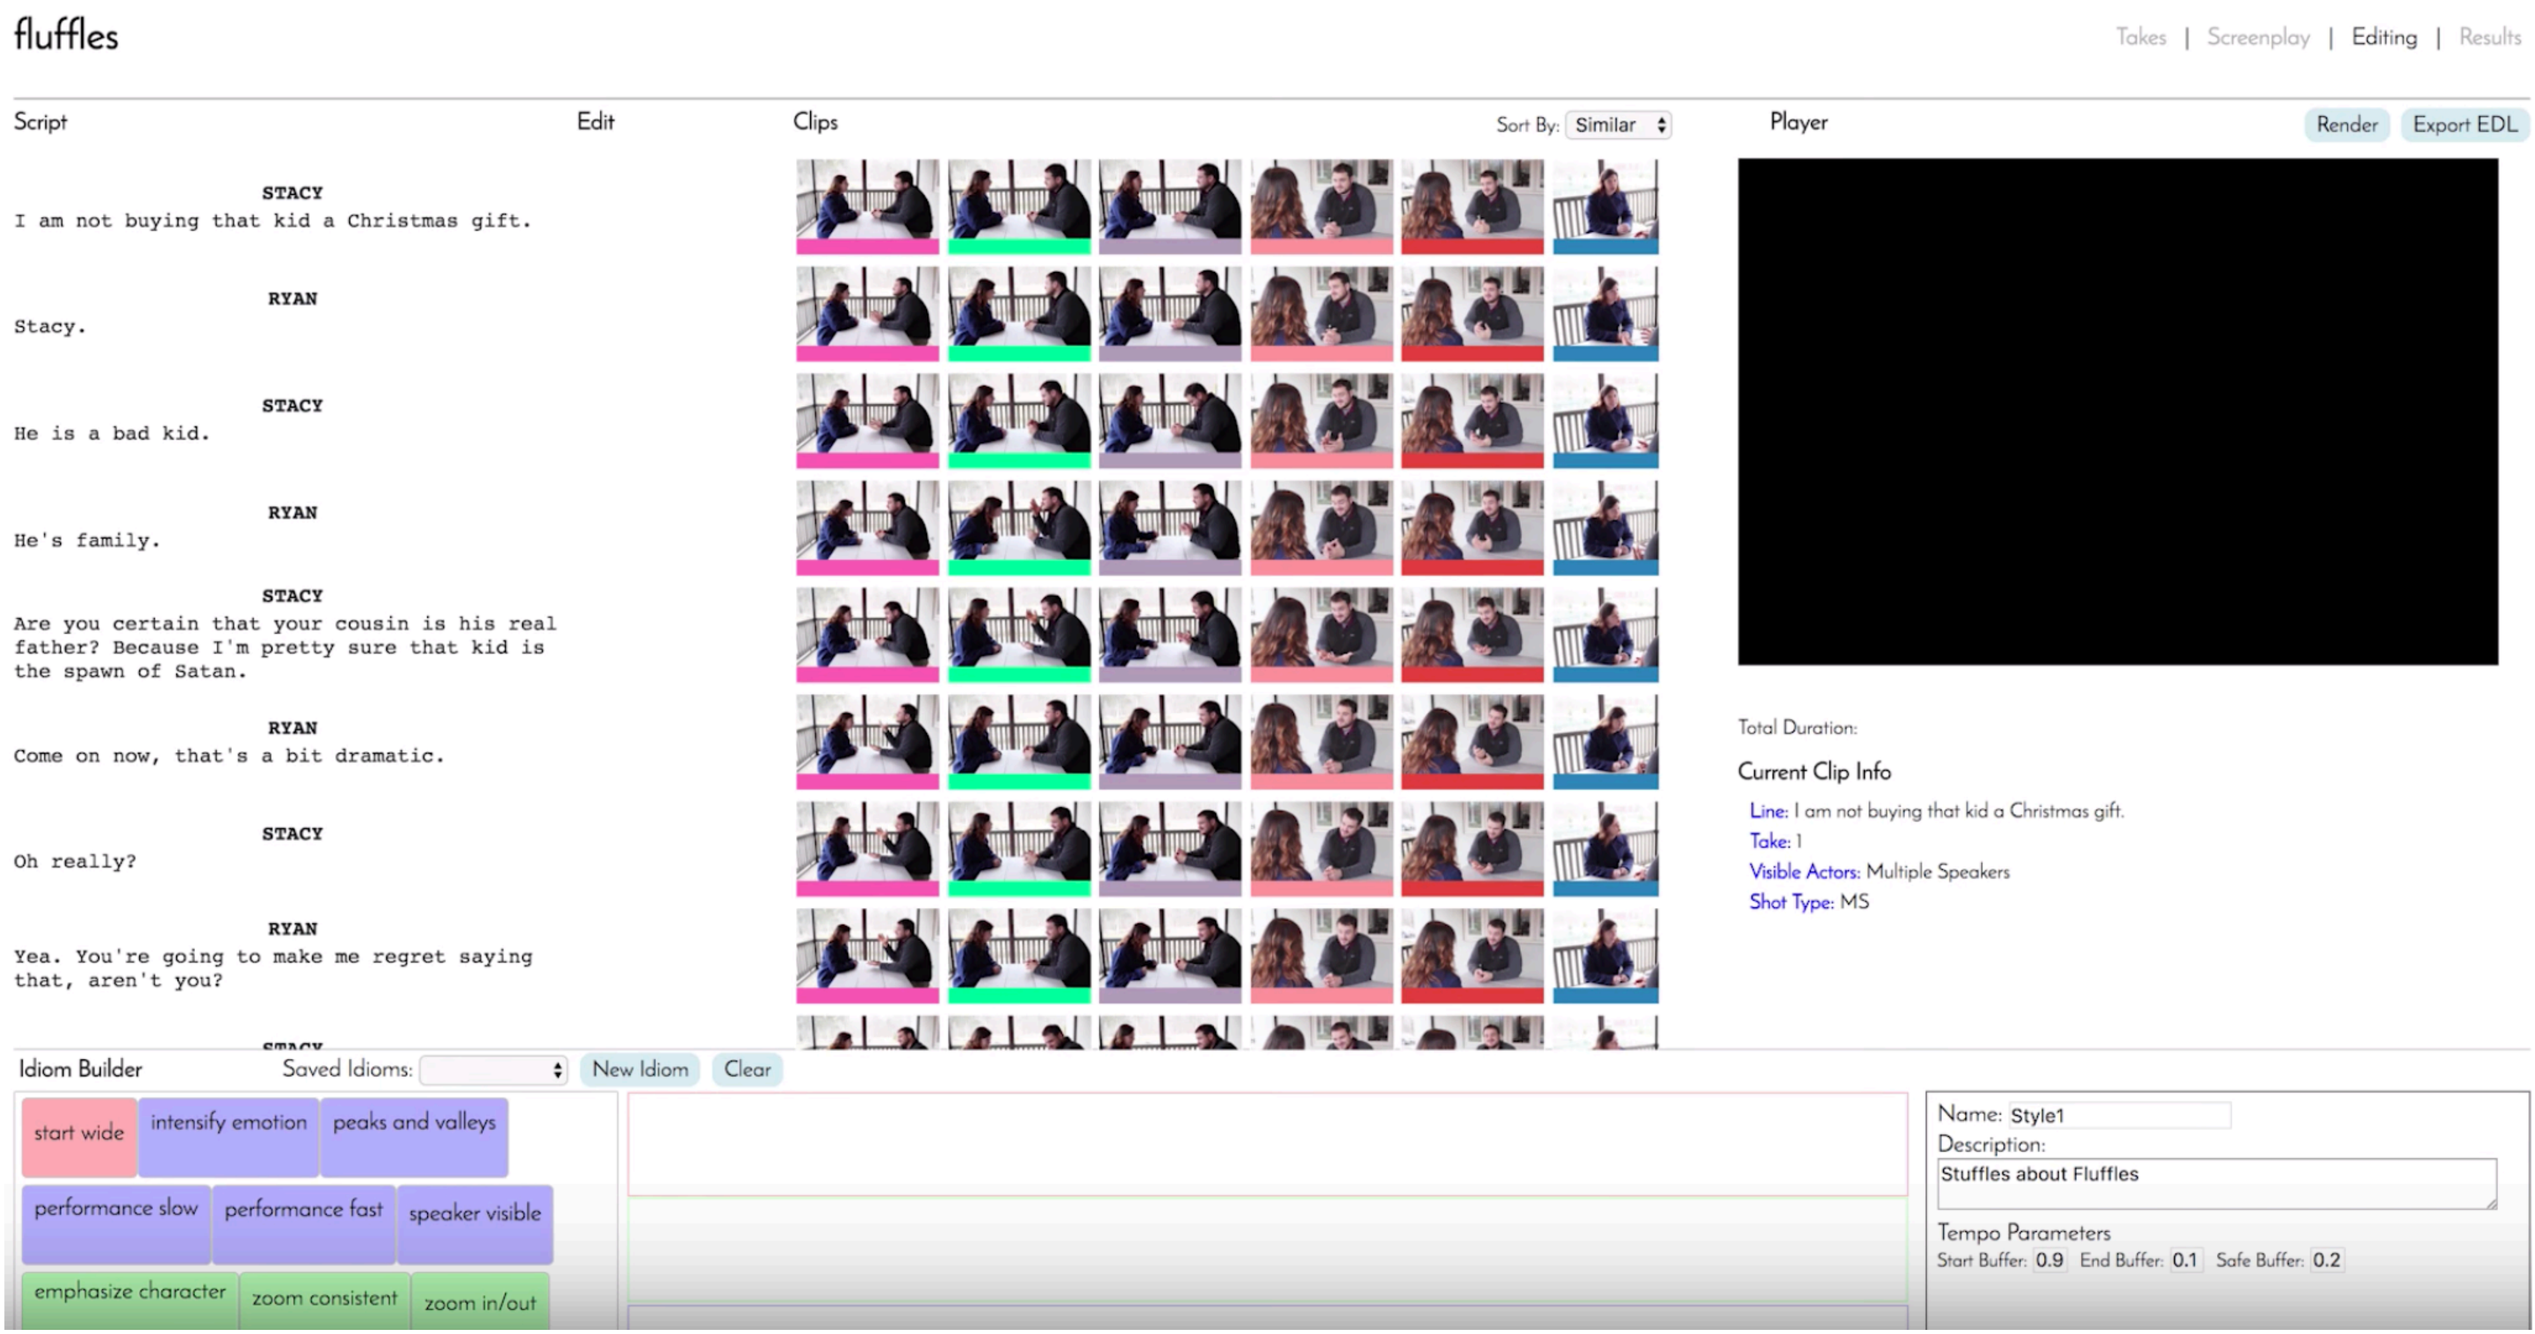

Doug James

- **Physically based simulation**

Example: Pouring Faucet

Frequency-domain radiation
[Langlois et al. 2016]

Time-domain radiation
[Our approach]

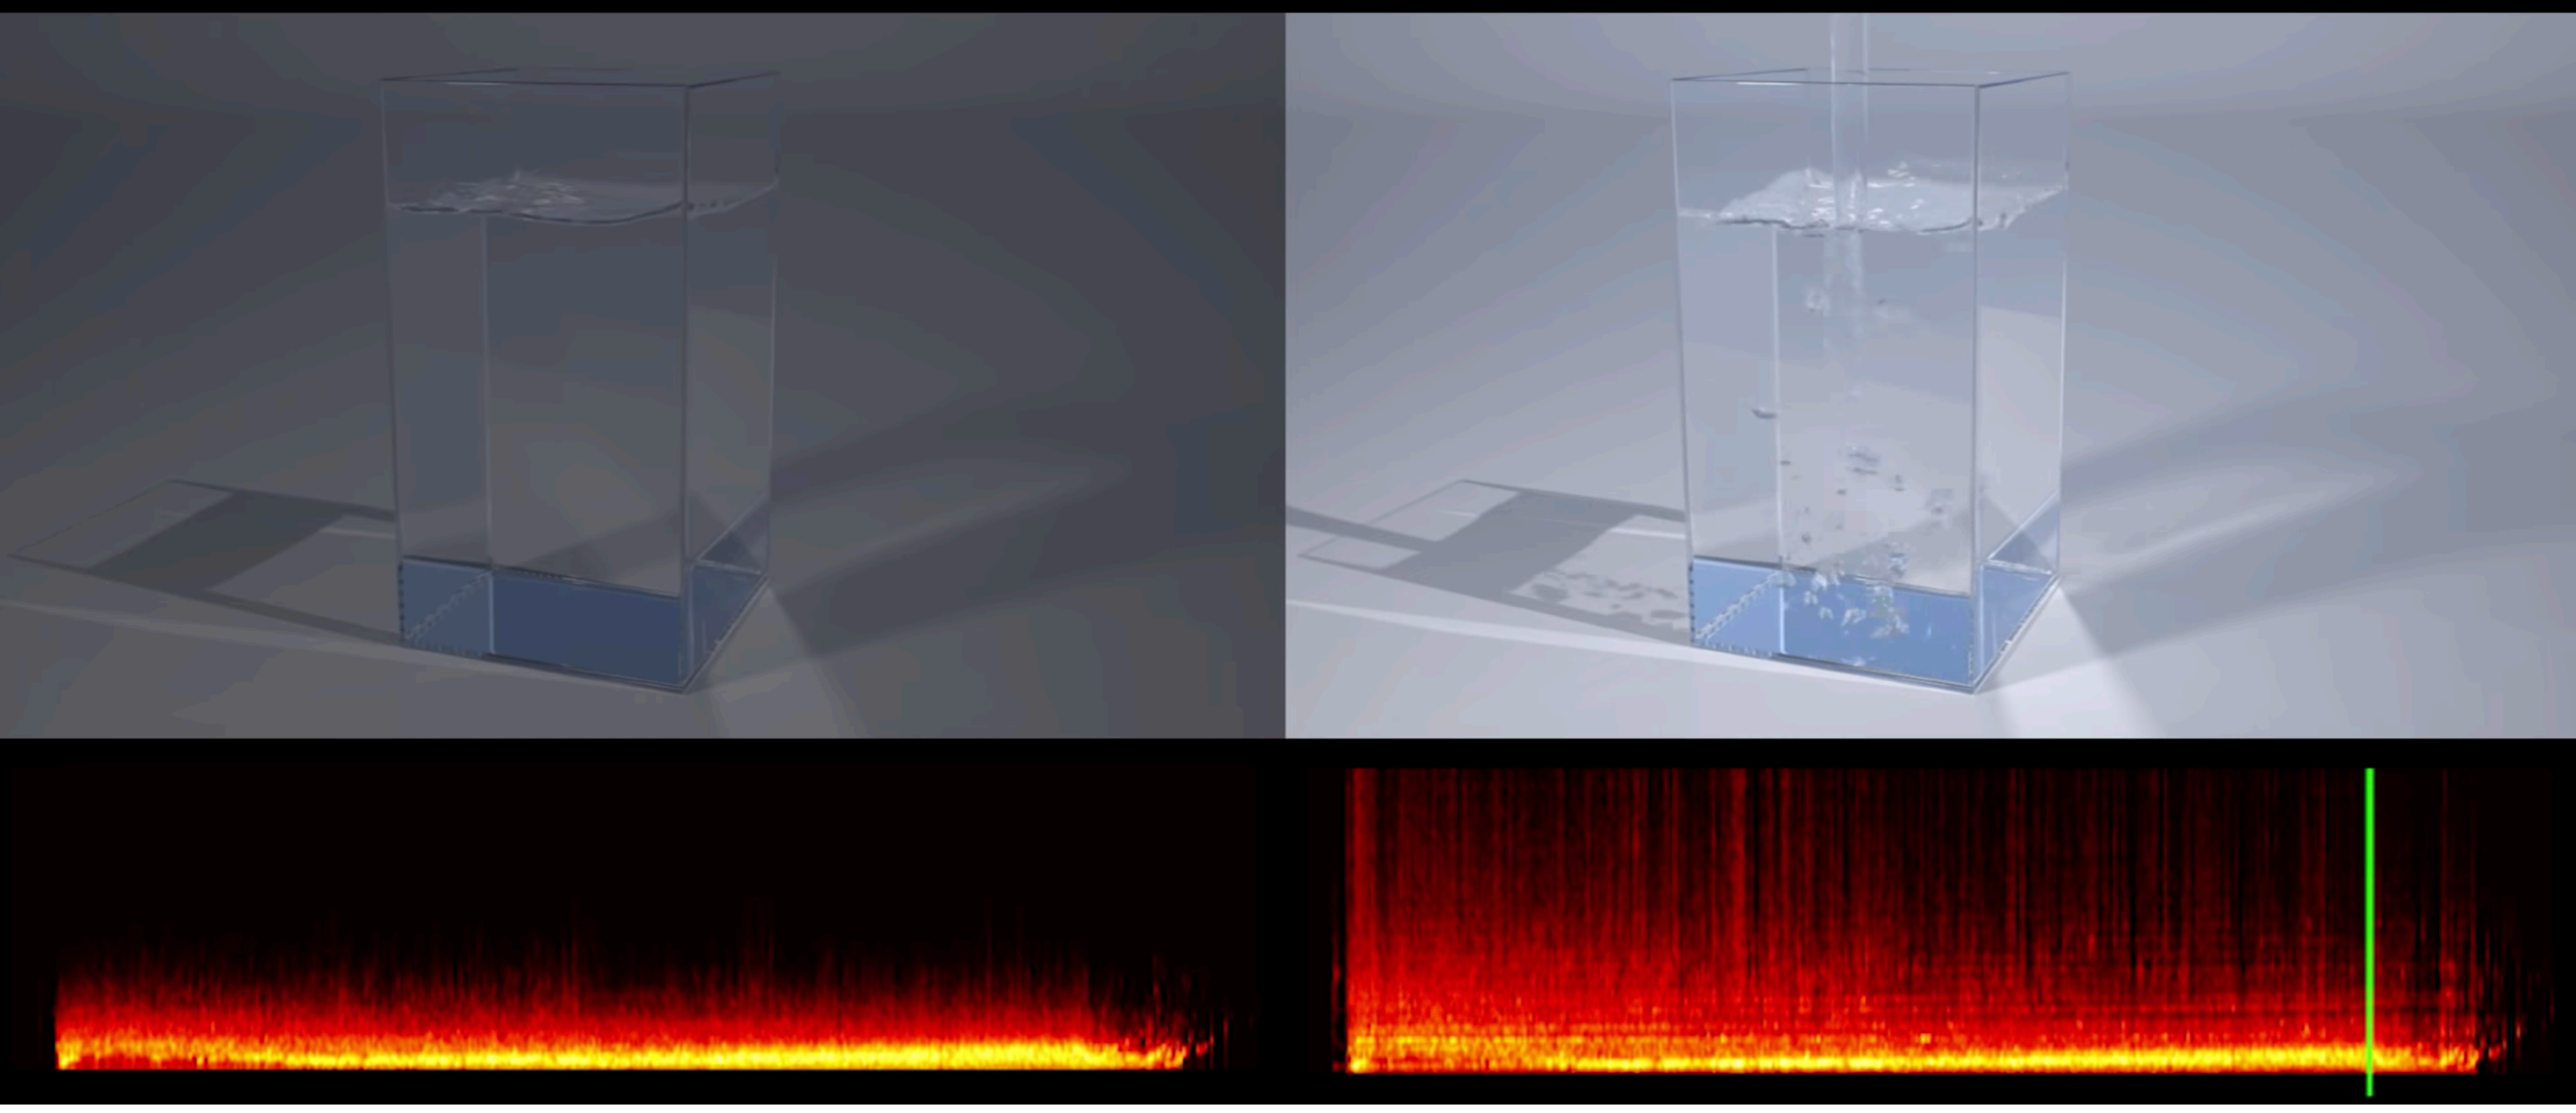

Doug James

- Physically based simulation

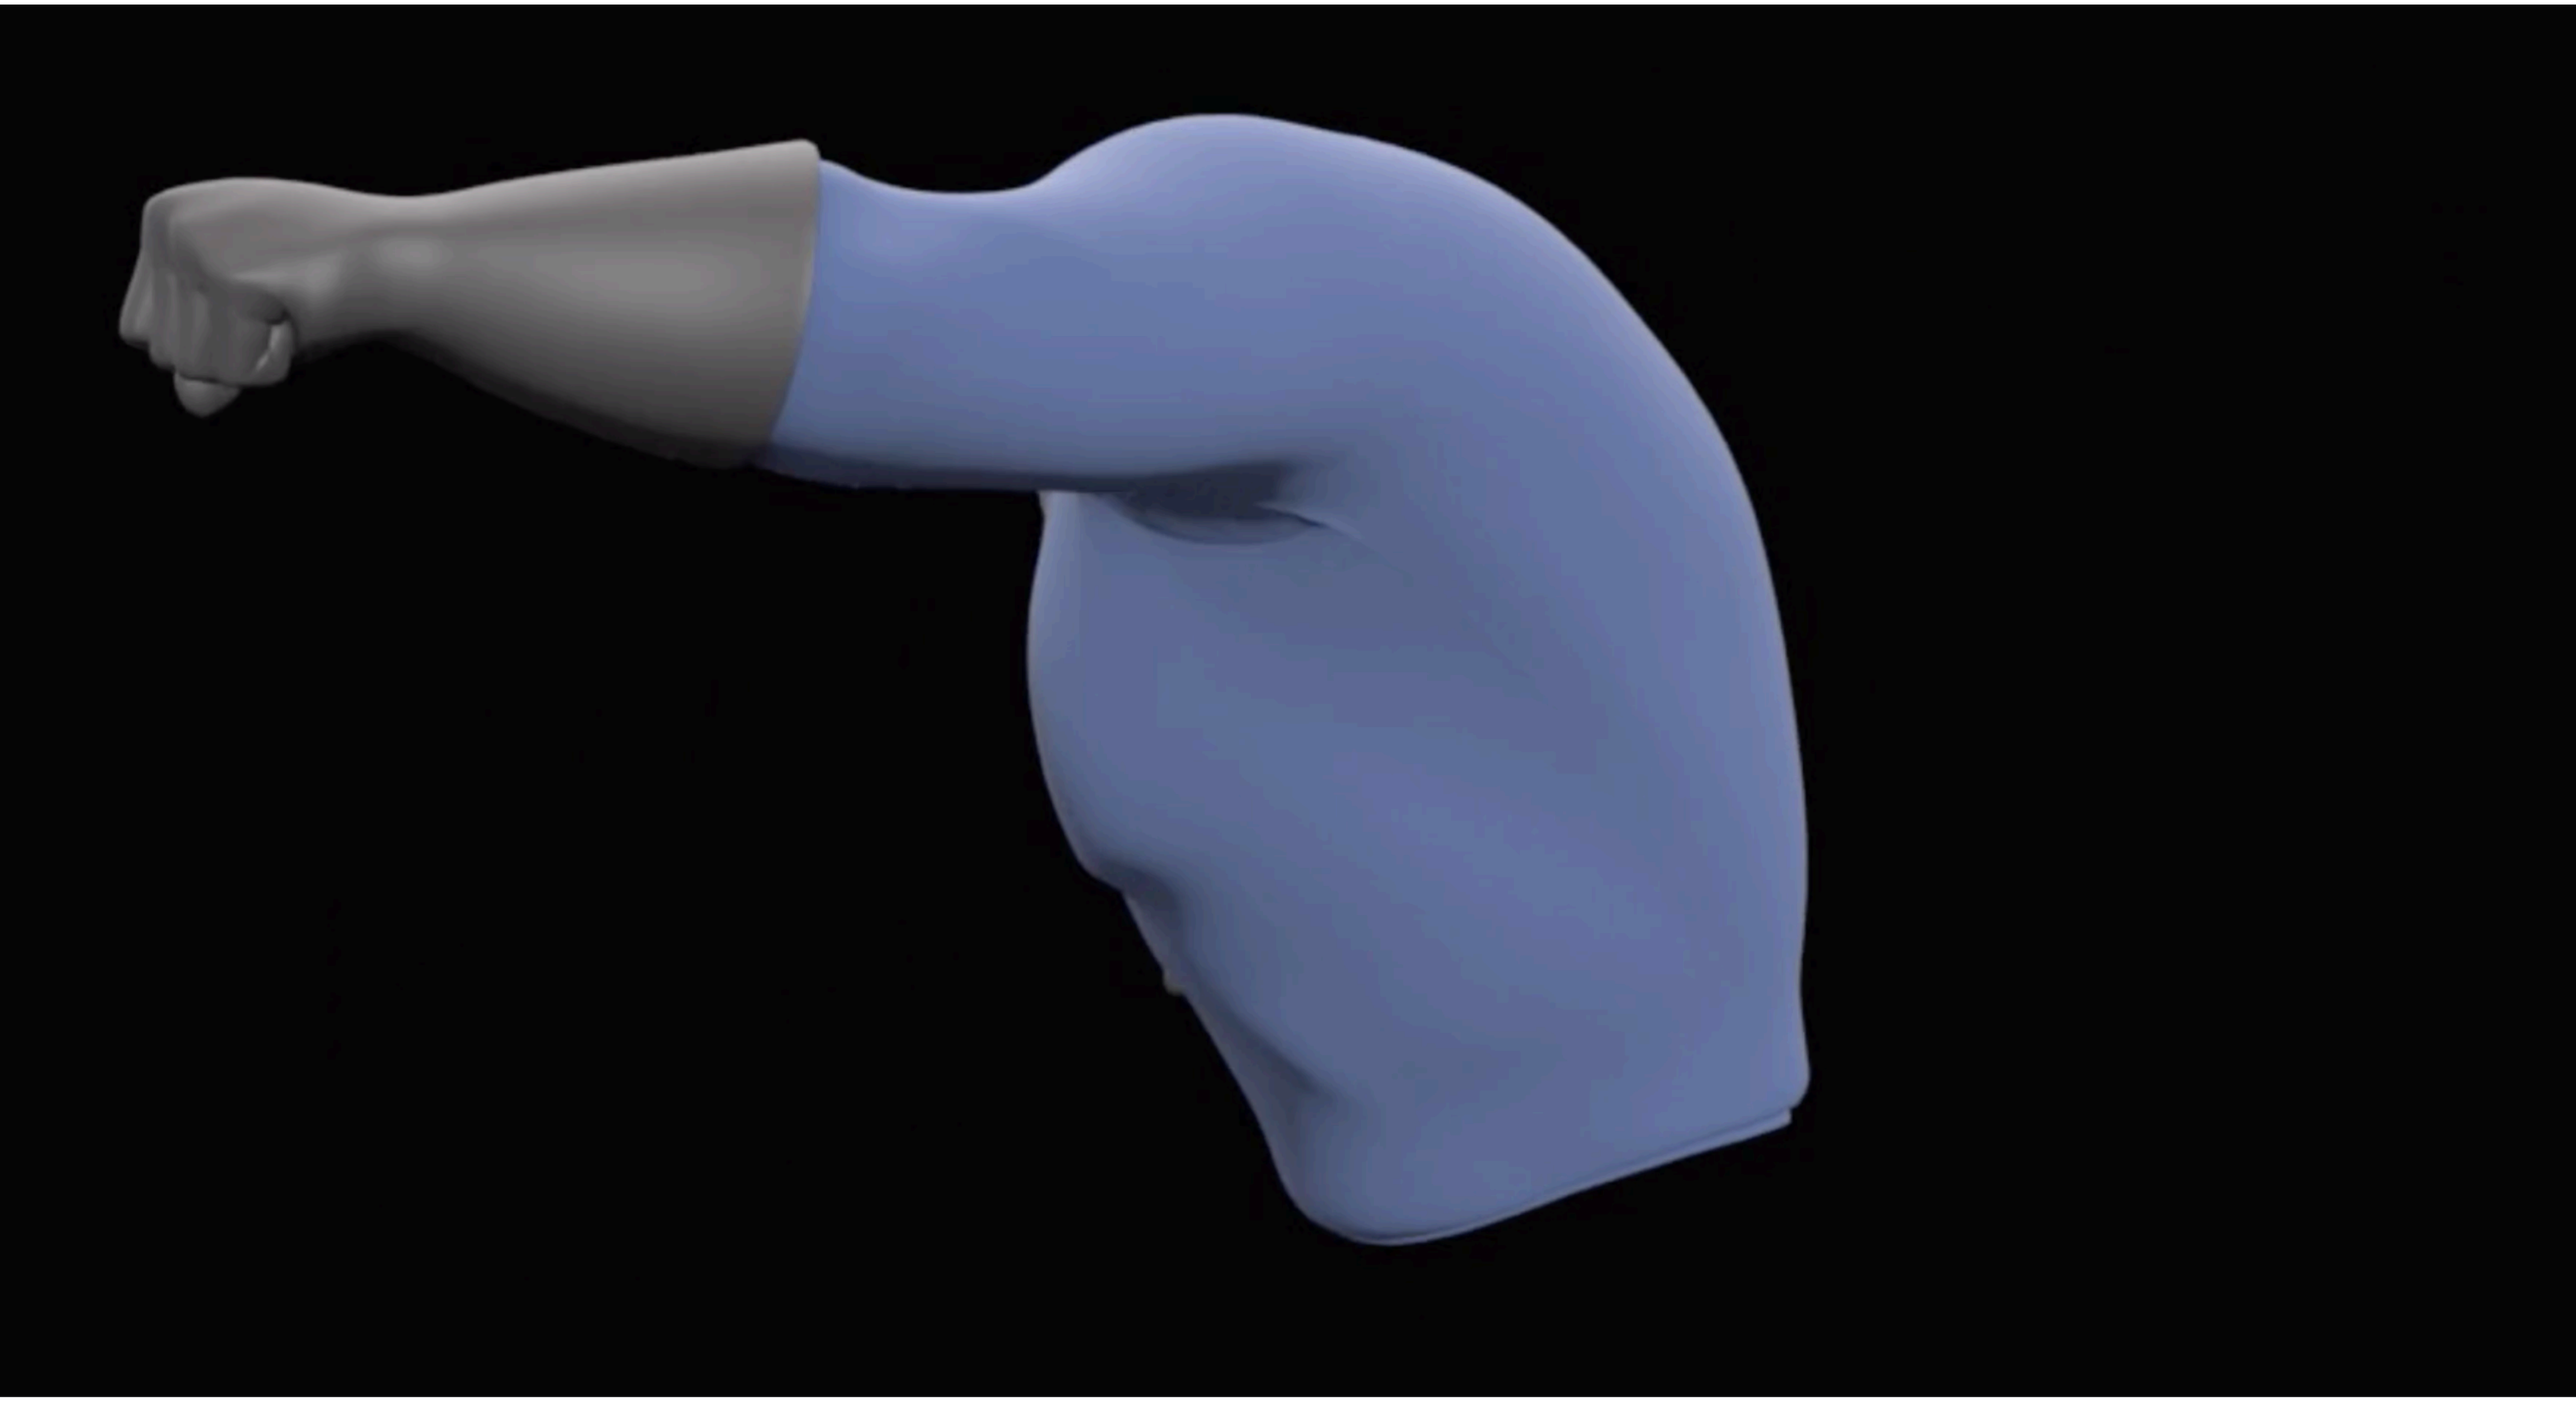

Doug James

- Physically based simulation

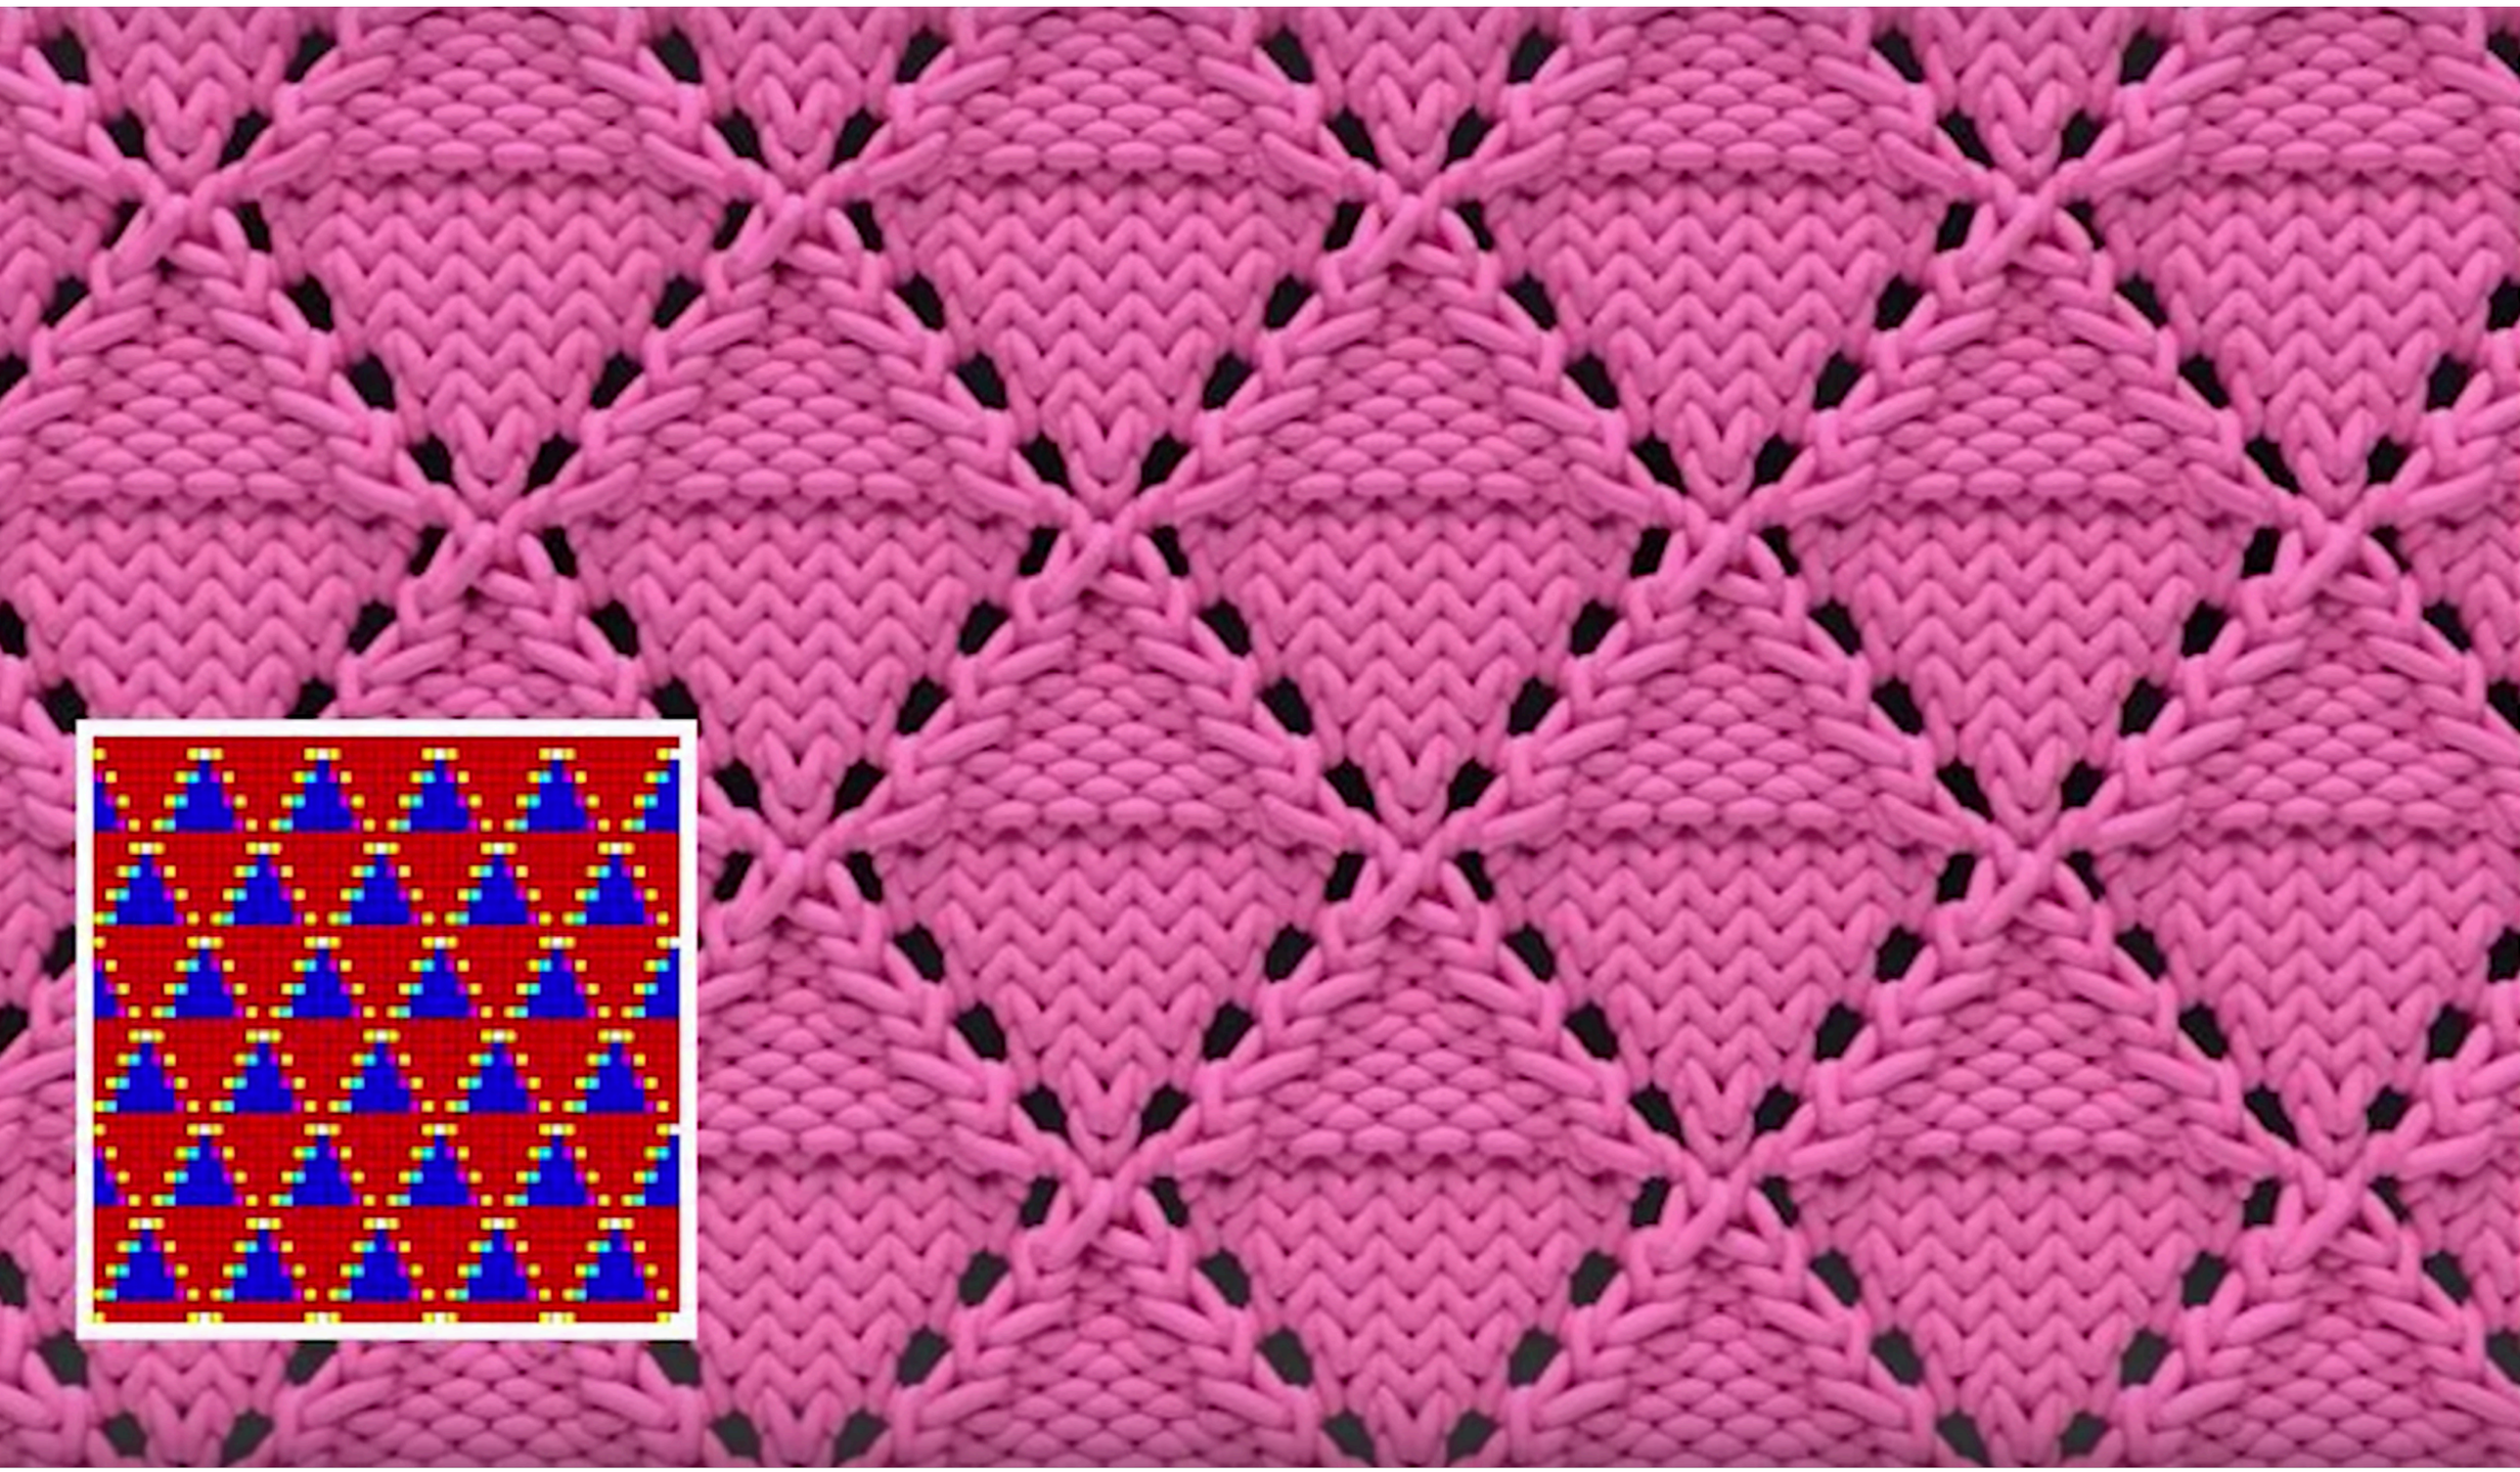

Leo Guibas

- **Geometry processing and analysis**

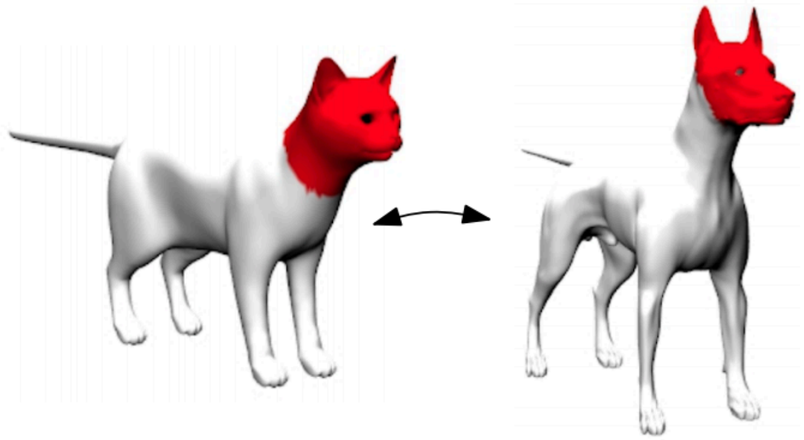

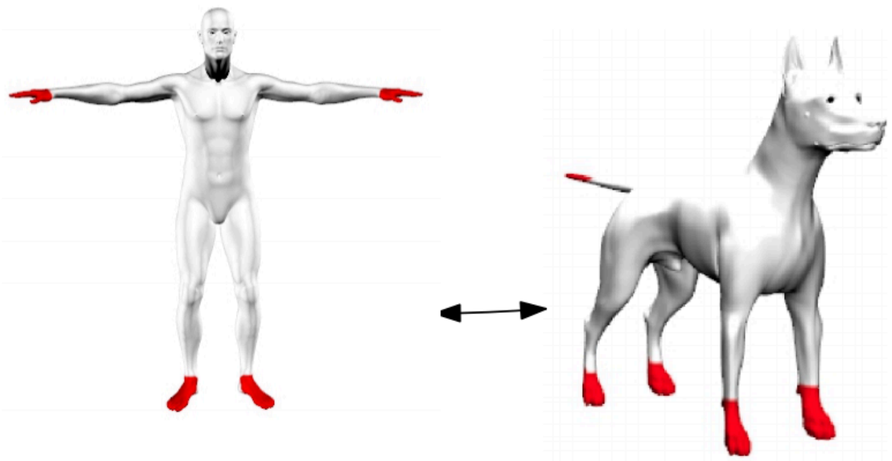

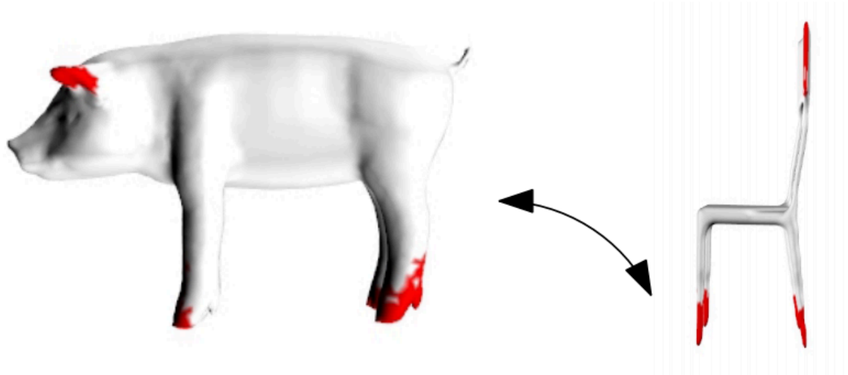

Shape Similarity and Correspondence

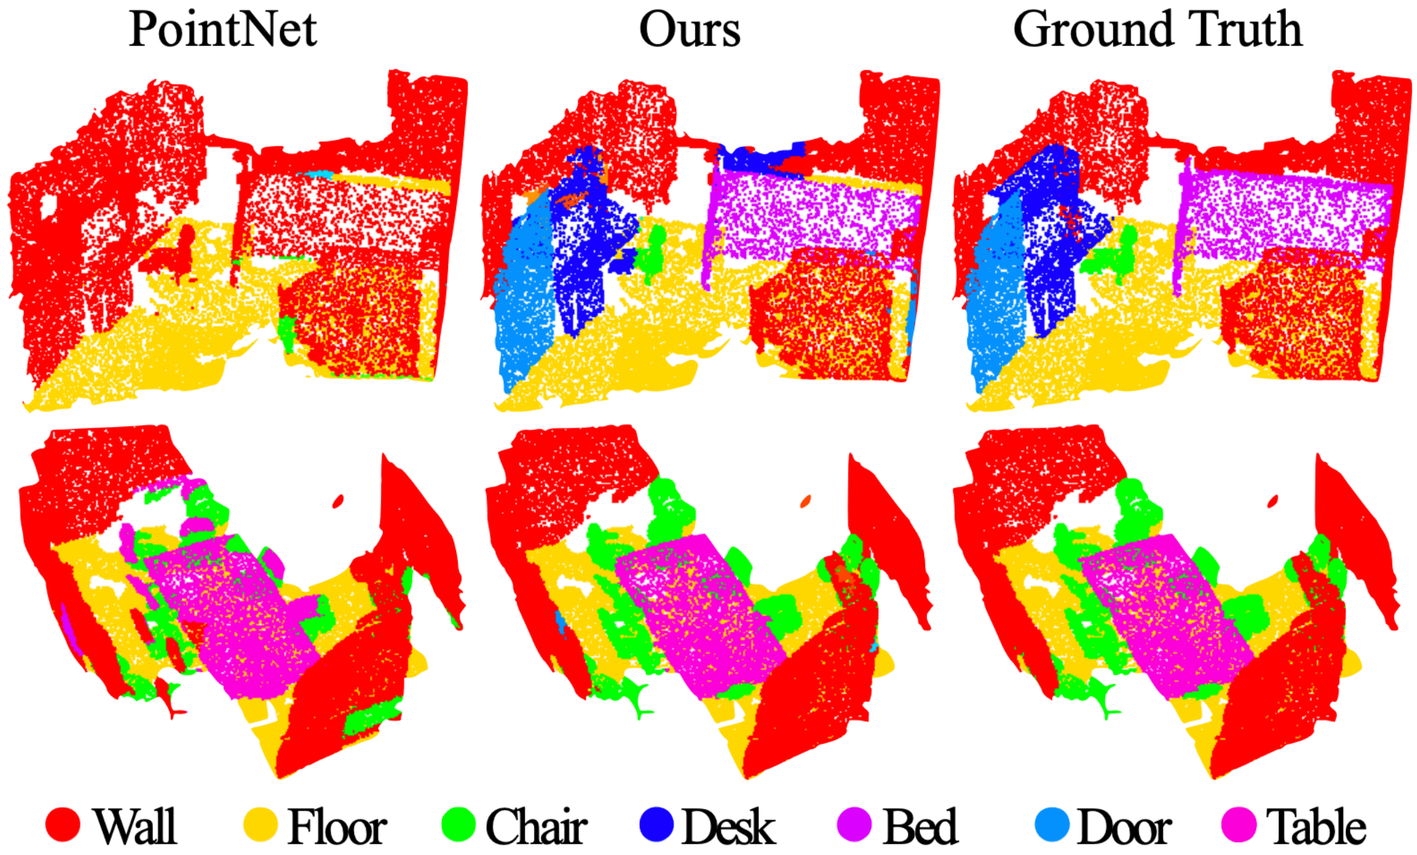

PointNet: Deep Learning on Point Clouds

Gordon Wetzstein

- Computational imaging and computational displays

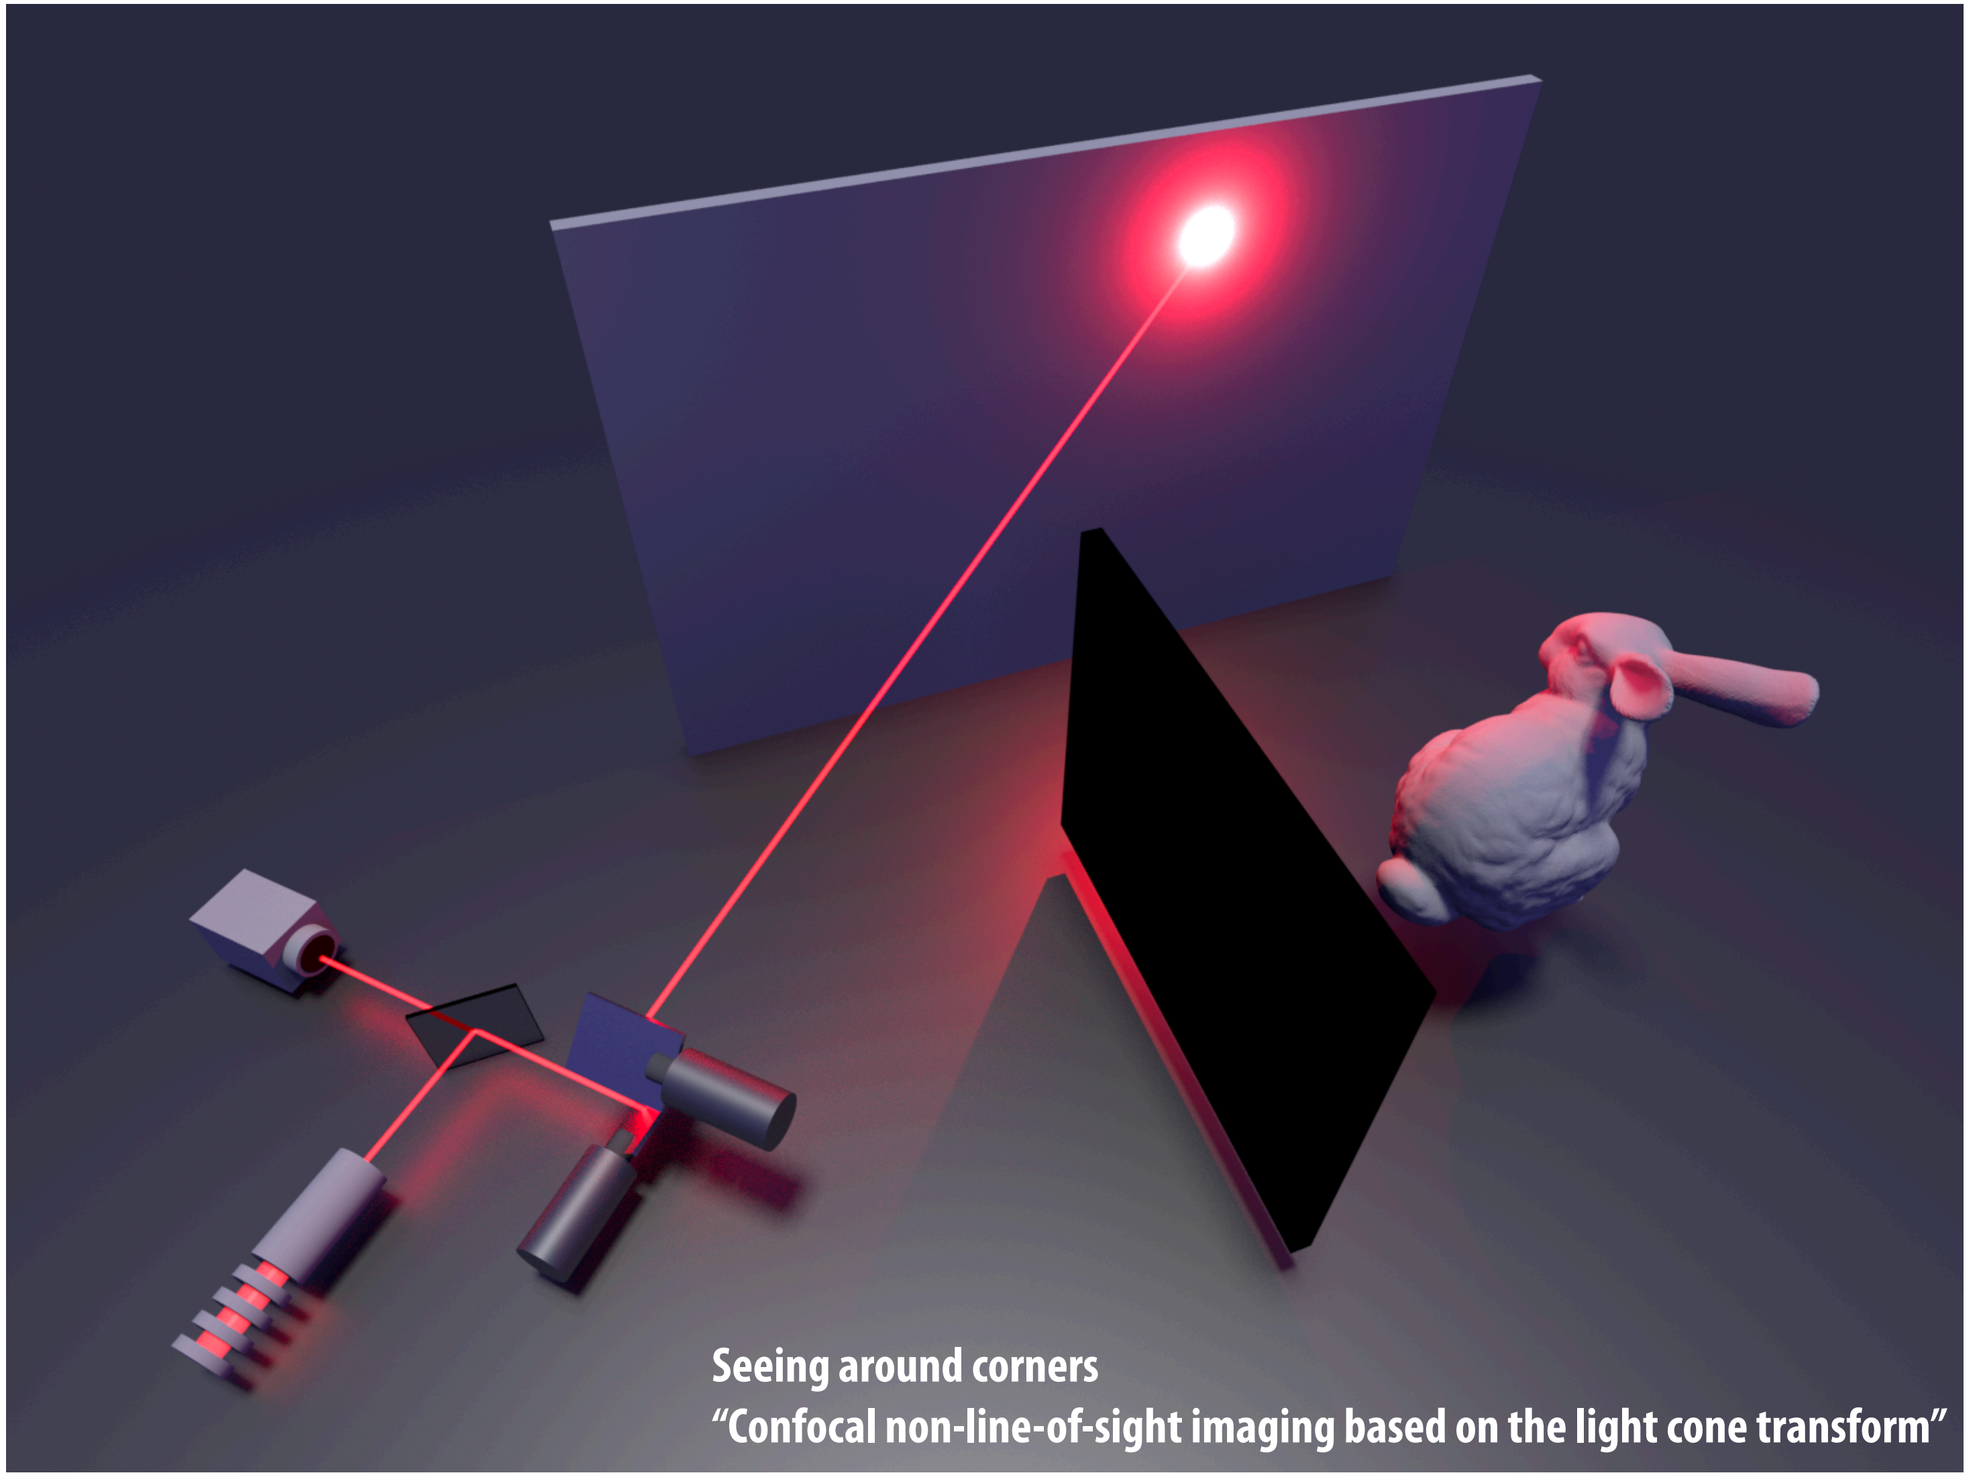

Gordon Wetzstein

- Computational imaging and computational displays

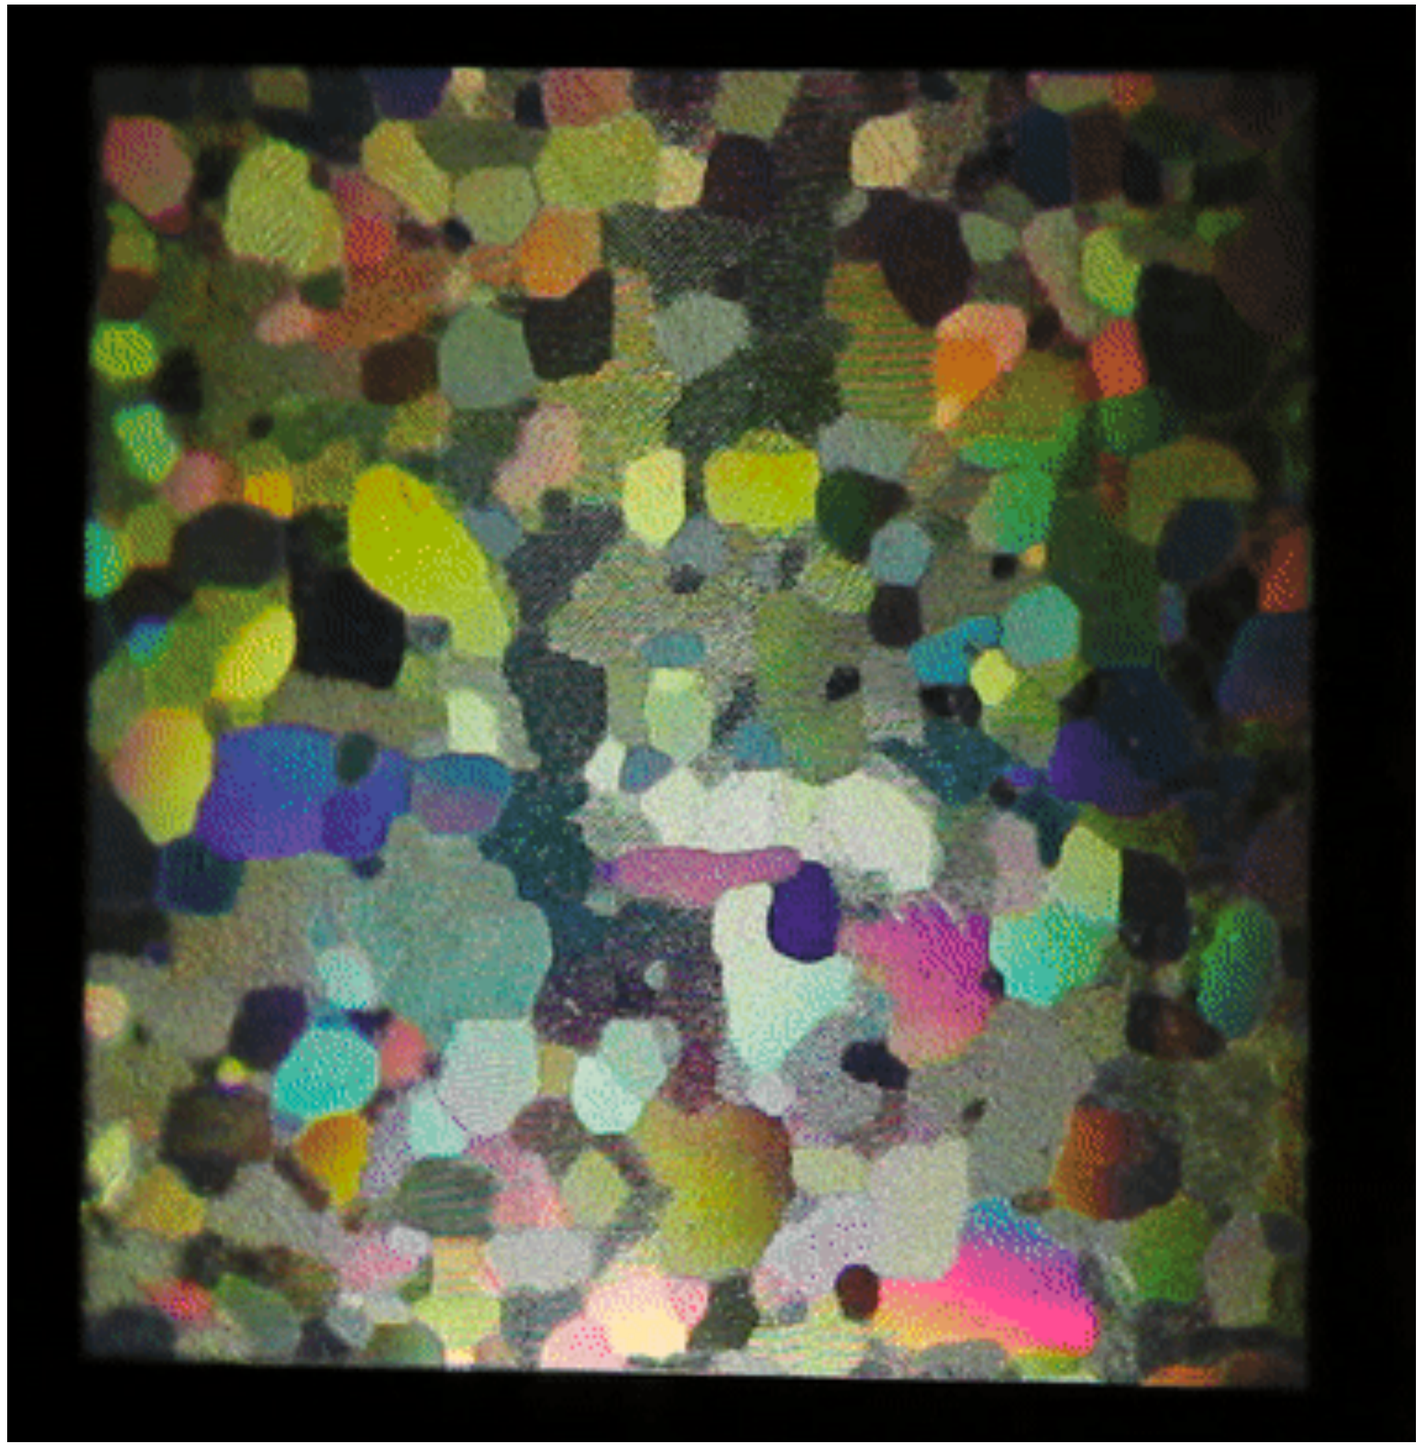

"Hybrid Optical-Electronic Convolutional Neural Networks"

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

Karen Liu

- New Stanford faculty (arriving this summer!)
- Interests in animation, simulation, and control

Karen Liu

Front Wheel Pivot
(BMX Bike)

Kayvon Fatahalian (me)

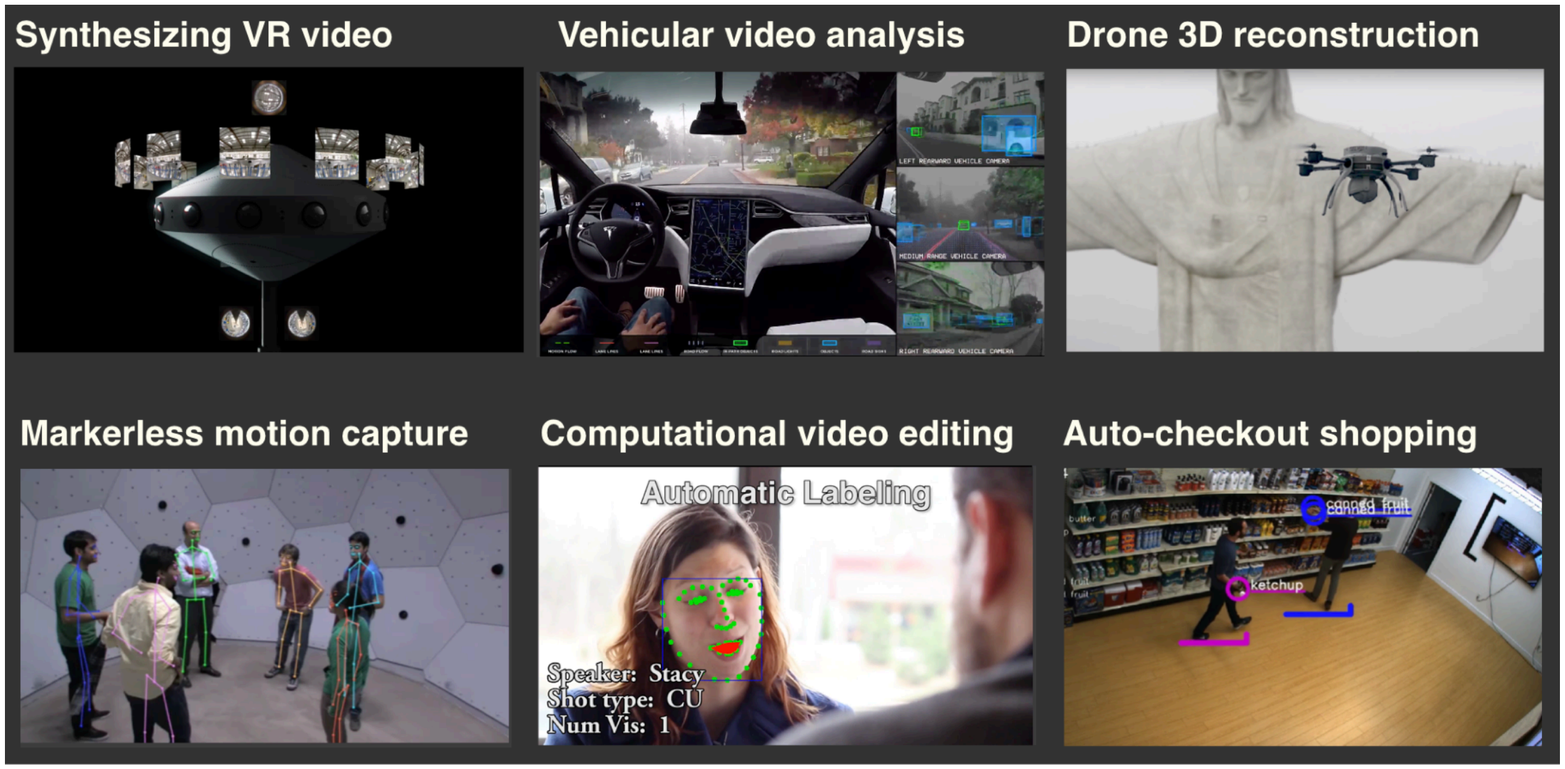

Platforms for scaling modern video processing applications to hundreds of GPUs or thousands of CPUs

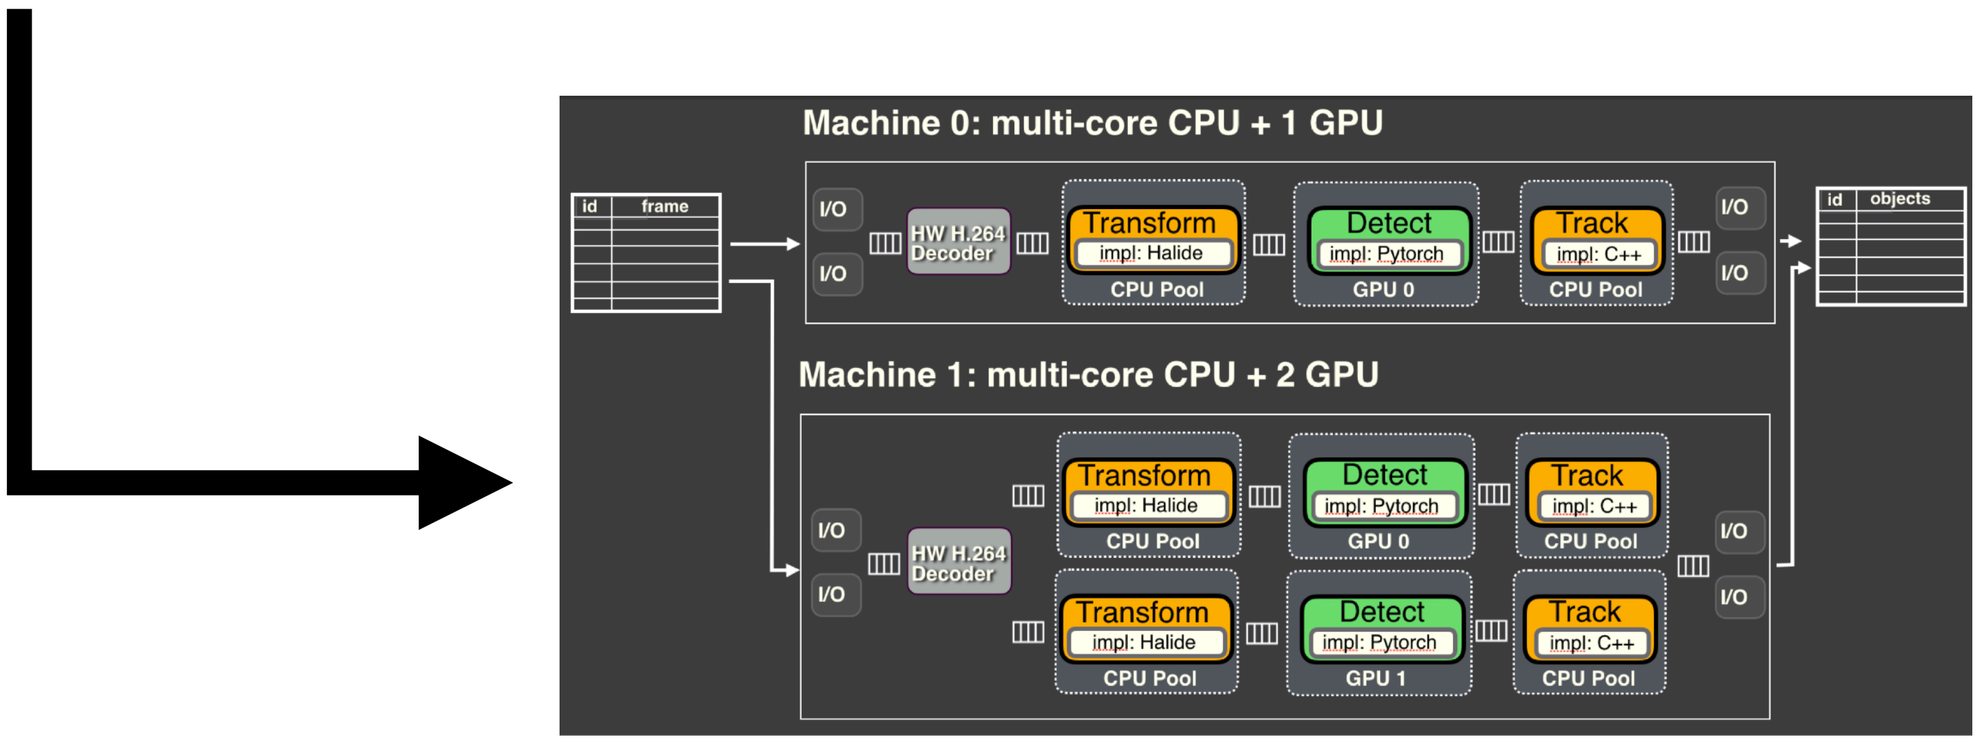

Analyzing video to curate content

All shots of Hermione between Ron and Harry in Harry Potter Films

Analyzing video to curate content

All interviews on TV news

Graphics systems projects

- New shading languages that bring modern programming language features to the GPU

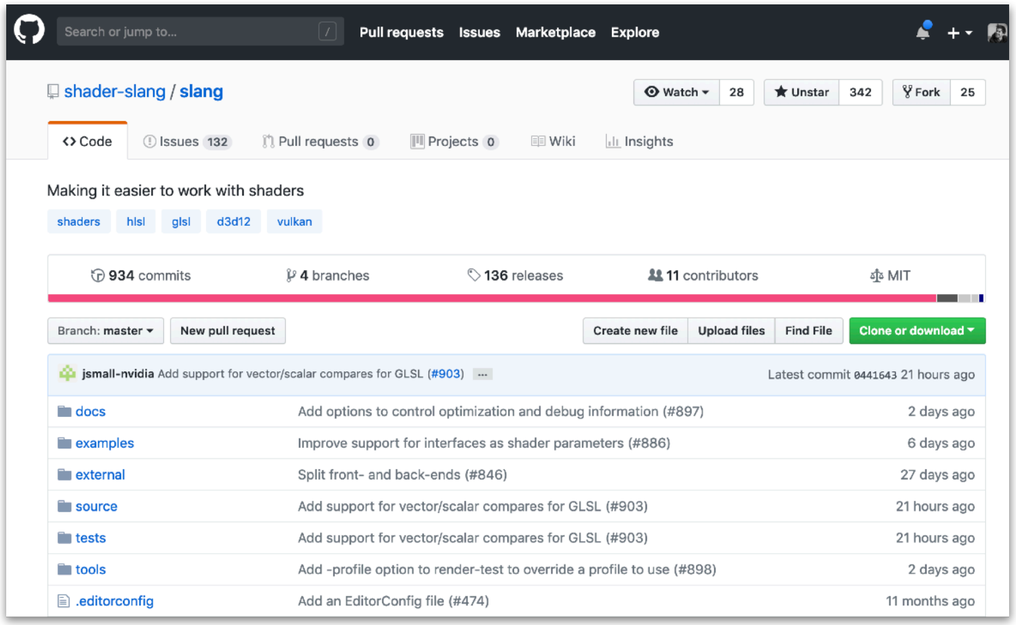

- Compiler techniques for automatically optimizing image processing programs

- DNN model architectures that reduce the cost of image/video processing

Other popular research topics in computer graphics…

Creating physically plausible models

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




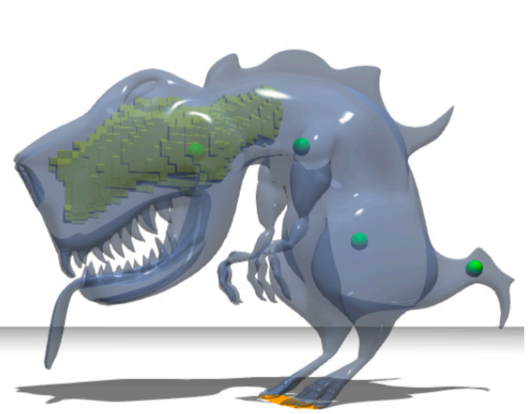
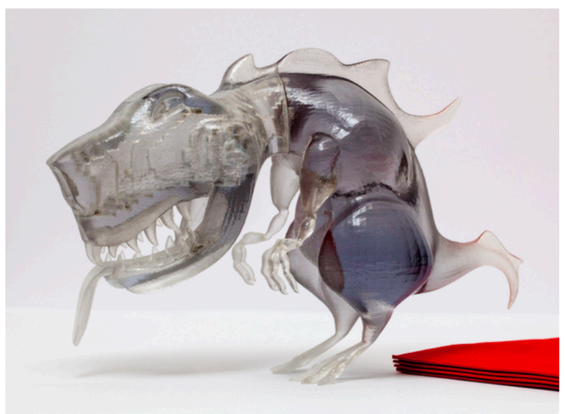

Fabricate models that are balanced to stand

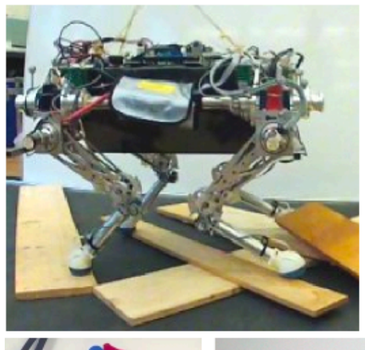
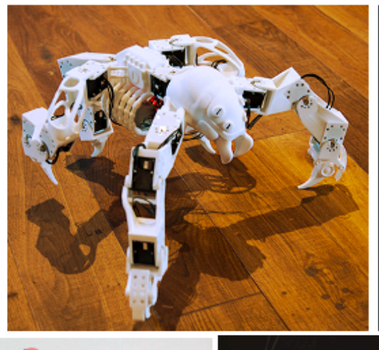
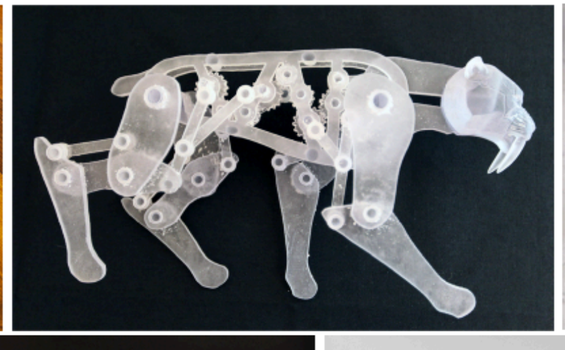
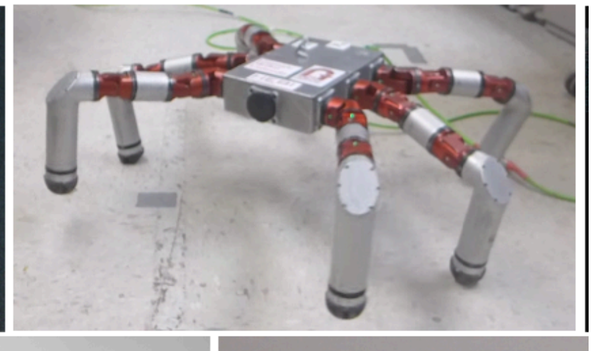

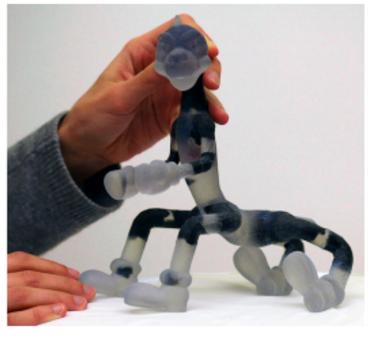
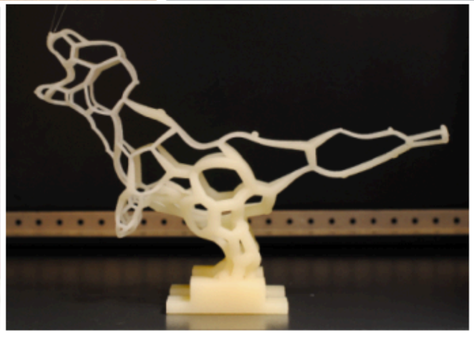
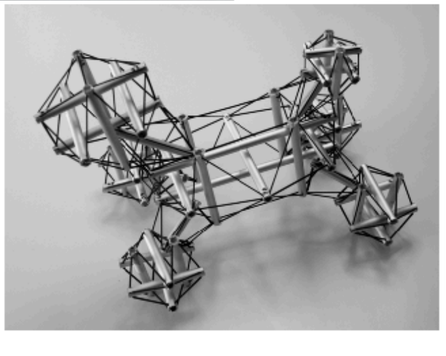
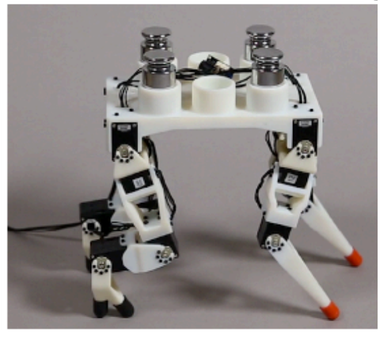

Fabricate robots that can balance and move

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)

Computational photography

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

High Dynamic Range Imaging (HDR)

Advanced geometry processing

Fundamental questions about alignment, similarly, symmetry, etc…

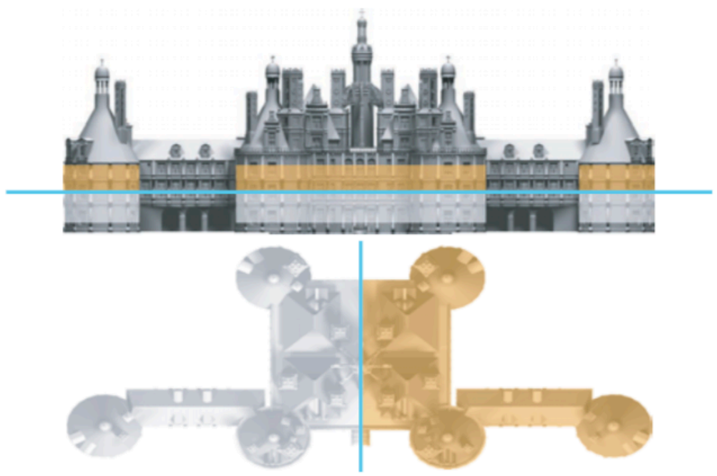

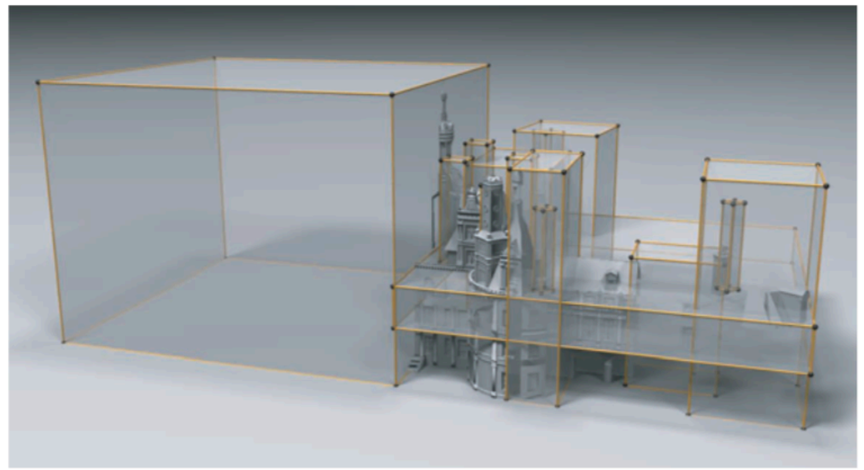

Advanced displays/rendering for VR/AR

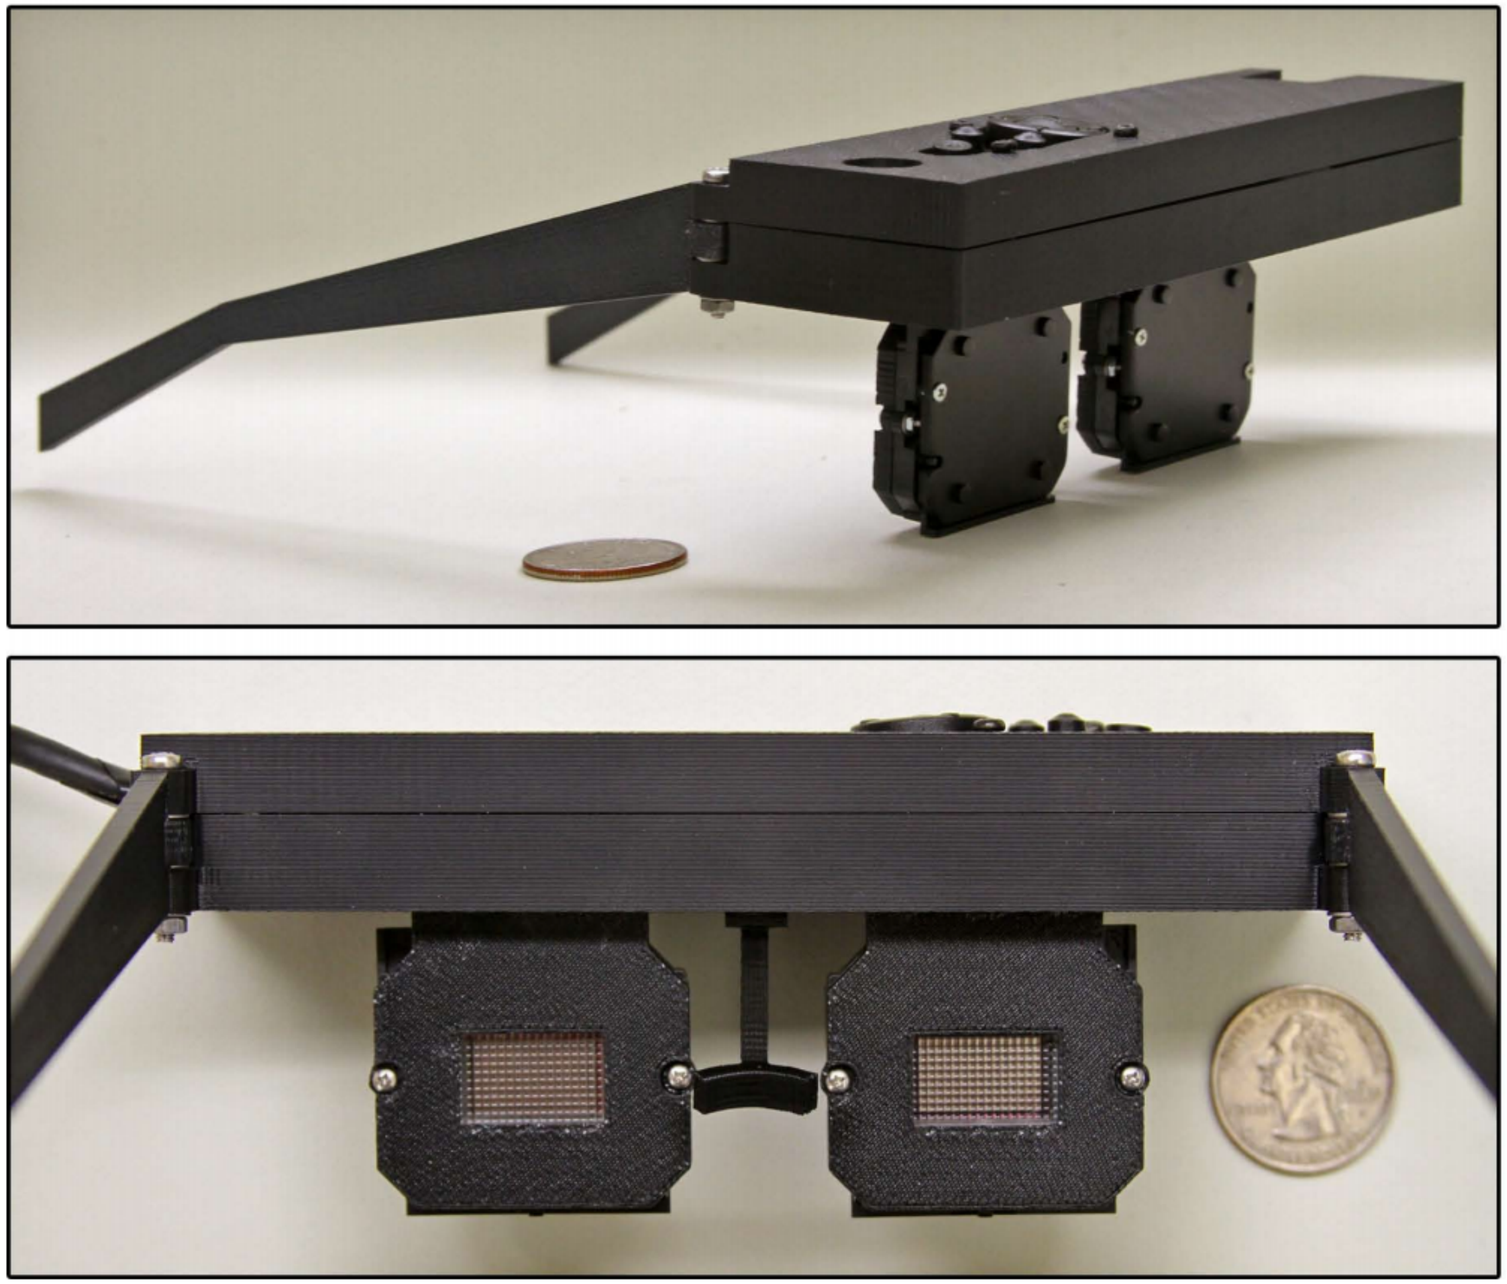

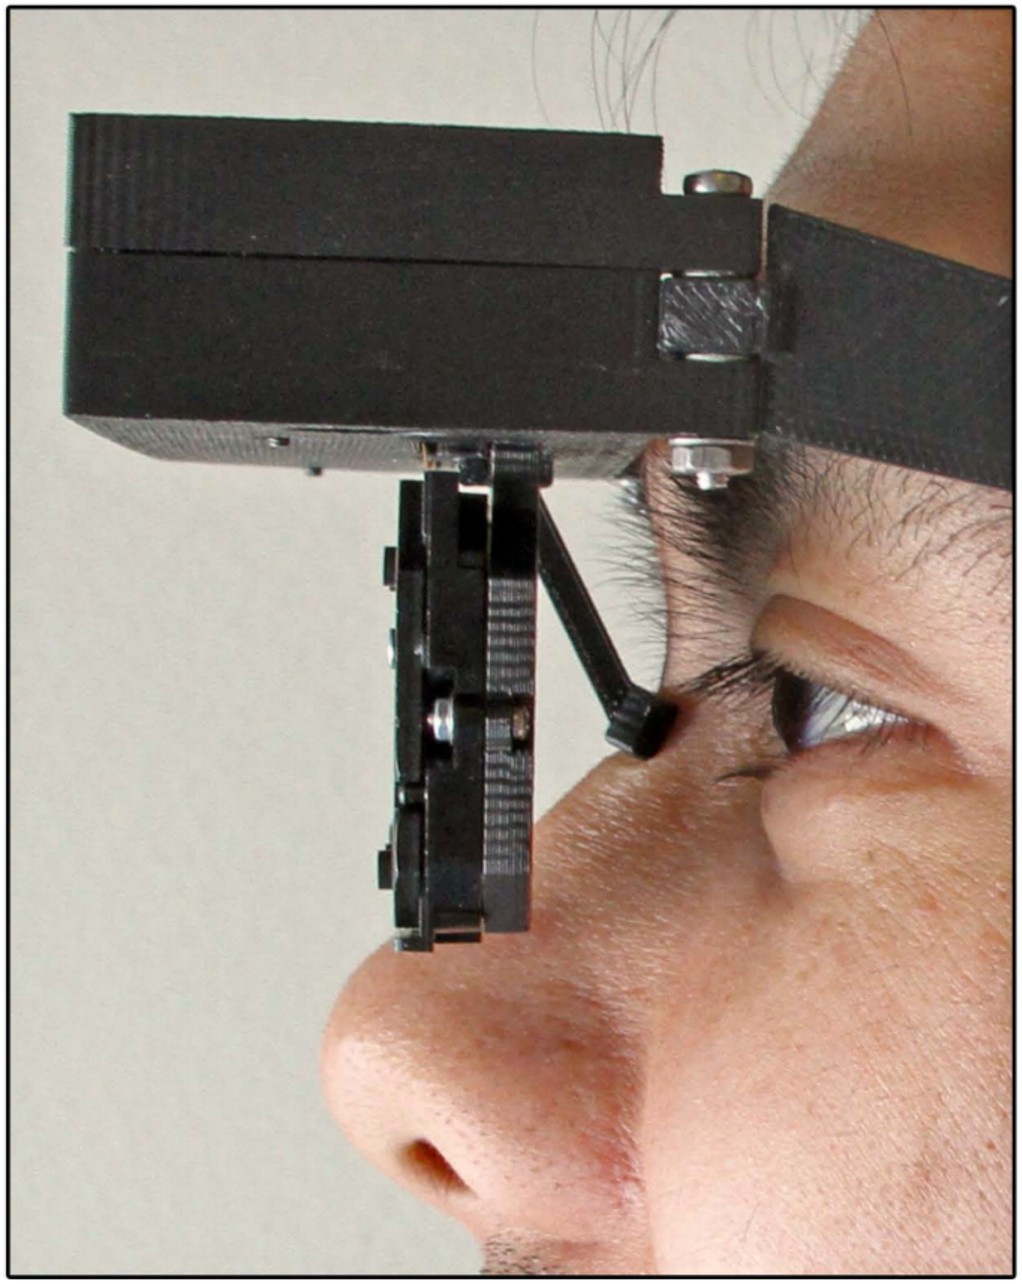

Near eye light field display

Content creation and capture

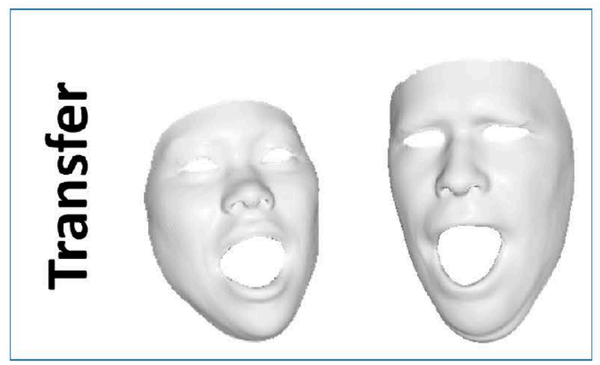

Manipulating actors by performance capture

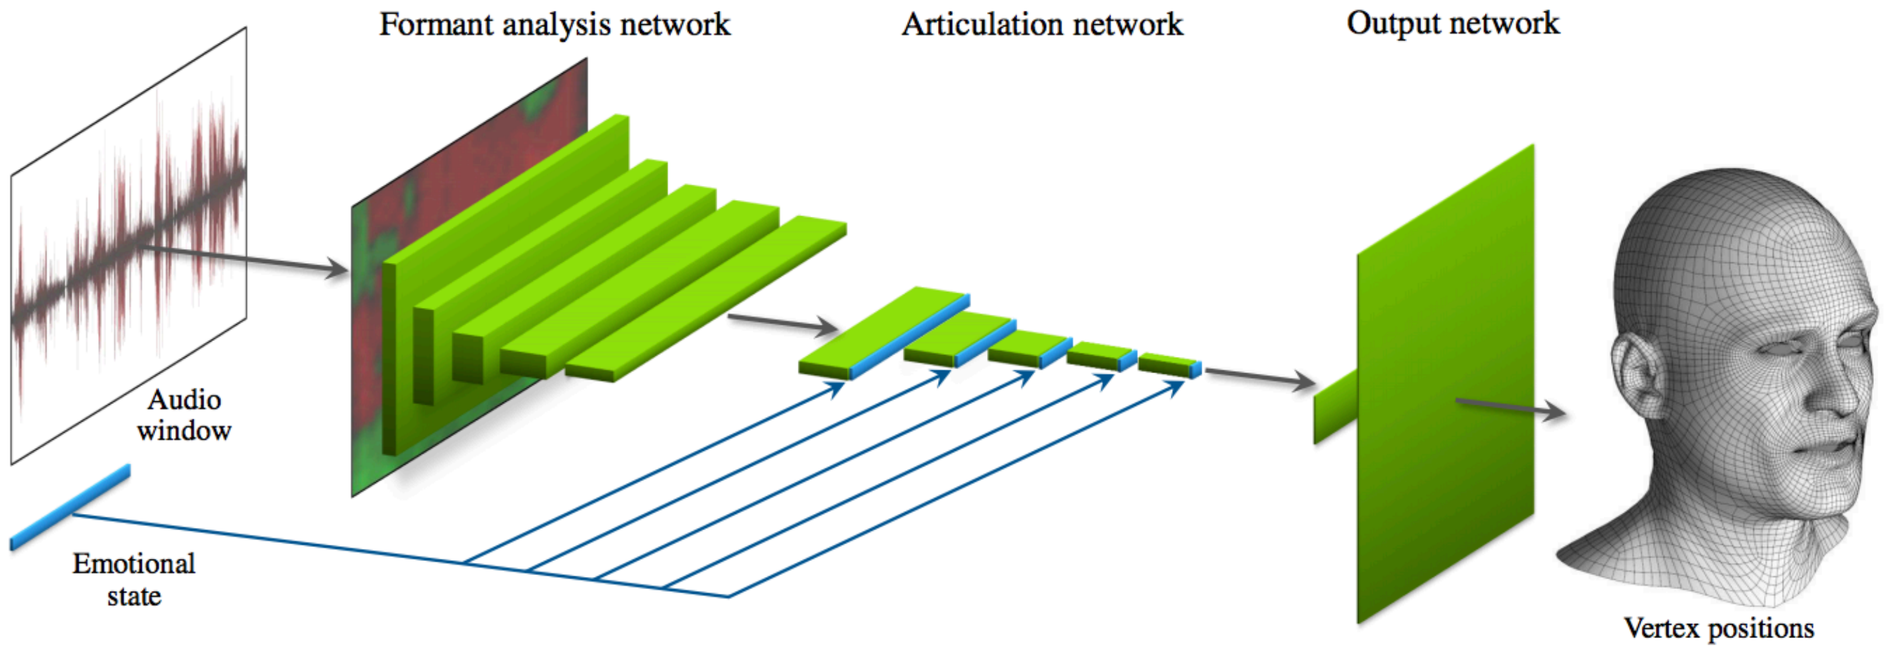

Audio input to mesh animation

The other direction: graphics helping machine learning

Input labels

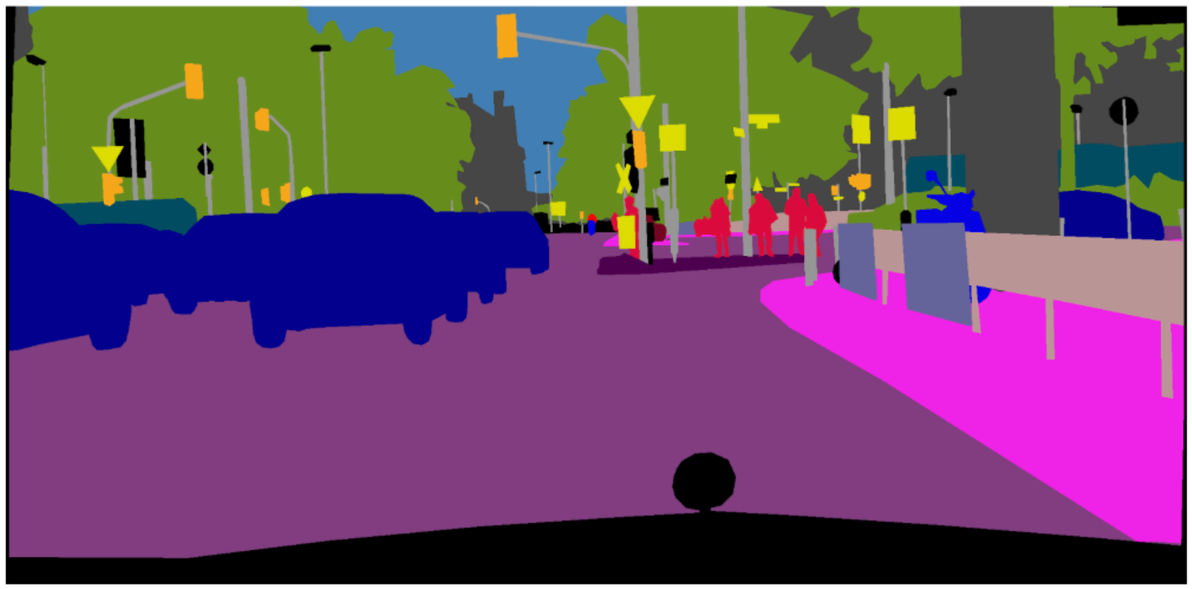

Synthesized image

Pix2pixHD

Grand Theft Auto Screenshots

Synthesized "photorealistic" image

Carla autonomous driving simulator

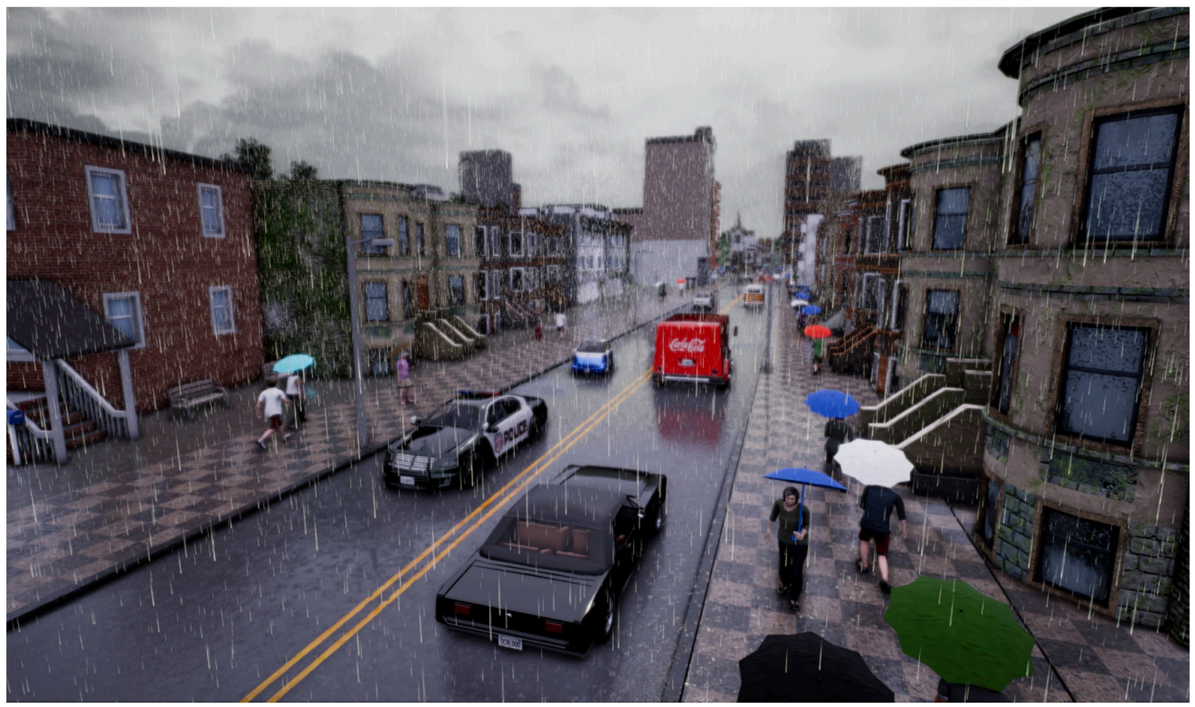

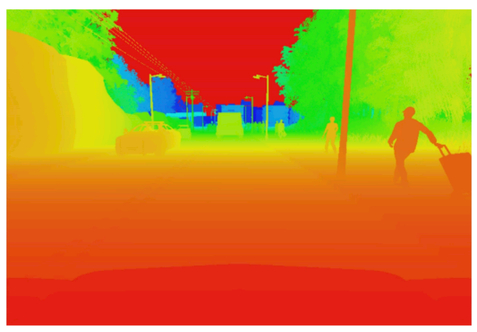

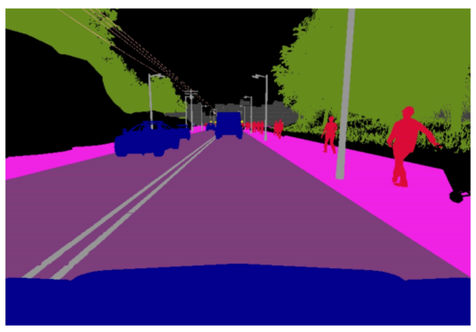

A fun resource

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

http://kesen.realtimerendering.com/

SIGGRAPH 2018 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 the page to find the actual document.

ACM Digital Library: ACM Transactions on Graphics (TOG) Volume 37, Issue 4 (July 2018) Proceedings of ACM SIGGRAPH 2018

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

(01) A Race to the Bottom (of the Geometric Energy Plot)

Blended Cured Quasi-Newton for Distortion Optimization
Yufeng Zhu (The University of British Columbia and Adobe Research), Robert Bridson (The University of British Columbia and Autodesk), Danny Kaufman (Adobe Research)

Progressive Parameterizations
Ligang Liu, Chunyang Ye, Ruiqi Ni, Xiao-Ming Fu (University of Science and Technology of China)

Anderson Acceleration for Geometry Optimization and Physics Simulation
Yue Peng (University of Science and Technology of China), Bailin Deng (Cardiff University), Juyong Zhang, Fanyu Geng, Wenjie Qin, Ligang Liu (University of Science and Technology of China)

Opt: A Domain Specific Language for Non-linear Least Squares Optimization in Graphics and Imaging (TOG Paper)
Zachary DeVito Facebook Research), Michael Mara (Stanford University), Michael Zollhofer (MPI Informatics), Gilbert Bernstein (Stanford University), Jonathan Ragan-Kelley (UC Berkeley), Christian Theobalt (MPI Informatics), Pat Hanrahan (Stanford University), Matthew Fisher (Adobe Research), Matthias Niessner (Technical University of Munich)

Active Animations of Reduced Deformable Models with Environment Interactions
Zherong Pan, Dinesh Manocha (University of North Carolina at Chapel Hill)

How to get involved

- 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 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 way more 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.

Example: what my own Ph.D. students are working on these days…

- Generating efficient code from image processing or deep learning DSLs (Halide Autoscheduler), and compiling these applications directly to FGPAs
- Designing a new shading language for future real-time 3D graphics pipelines (collaboration with NVIDIA)
- Parallel computing platforms that make it simpler and more efficient to analyzing large video collections (Scanner project: "Spark for video")
- Designing programming models for querying video collections (e.g, find frames with "three people around a table" or where DNN1 disagrees with DNN2)
- Designing more efficient DNNs to accelerate image processing on video

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

Research is just one option…

(Despite what many of us biased faculty tell you,
there are many other good ones as well)

Why not start your own project?

Interested in applying computer science to a problem that excites you? Give it a shot!

Like a topic enough to be your own boss? Consider starting your own company.

Why go work for Google or Facebook when you can start a company that beats them?

(yes, those are great jobs too!)

Your professors encourage you to be brave and take risks.

You are lucky because you are extremely talented. The cost of "messing up" for you is actually much less than for other students because your backup plan is very good.

Be ambitious while at Stanford with opportunities beyond just classes. If it doesn't work out, you'll try something else and you'll probably succeed... or end up getting the good job you would have gotten anyway.

Thanks for being a great class!

Good luck on projects! Make sure you have fun, that's the point!

See you on the 21st!