Lecture 20:

Course Summary +
Graphics at Stanford Today

Computer Graphics: Rendering, Geometry, and Image Manipulation
Stanford CS248A, Winter 2023

As accomplished CS248A students you've now learned the basics of drawing shapes, representing surfaces/light/materials, manipulating images, etc…

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

What's next?

Graphics classes (coming quarters) at Stanford

SPRING

CS348K: "Visual Computing Systems", creating efficient systems for photography, 3D graphics, and modern AI (Fatahalian) - TTh 10:30am

CS348E: "Character Animation: Modeling, Simulation, and Control of Human Motion" (Liu) - MW 1:30am

CS 348N: "Neural Models for 3D Geometry" (Guibas) - MW 3:00pm

CS 231N: "Deep Learning for Computer Vision" (F. Li) - TTh 12:00pm

FALL

CS248B: "Fundamentals of Computer Graphics: Animation and Simulation " (Liu, James)

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)

Graphics Research at Stanford Today

Maneesh Agrawala

ControlNet: more precise ways to control diffusion-based generative AI

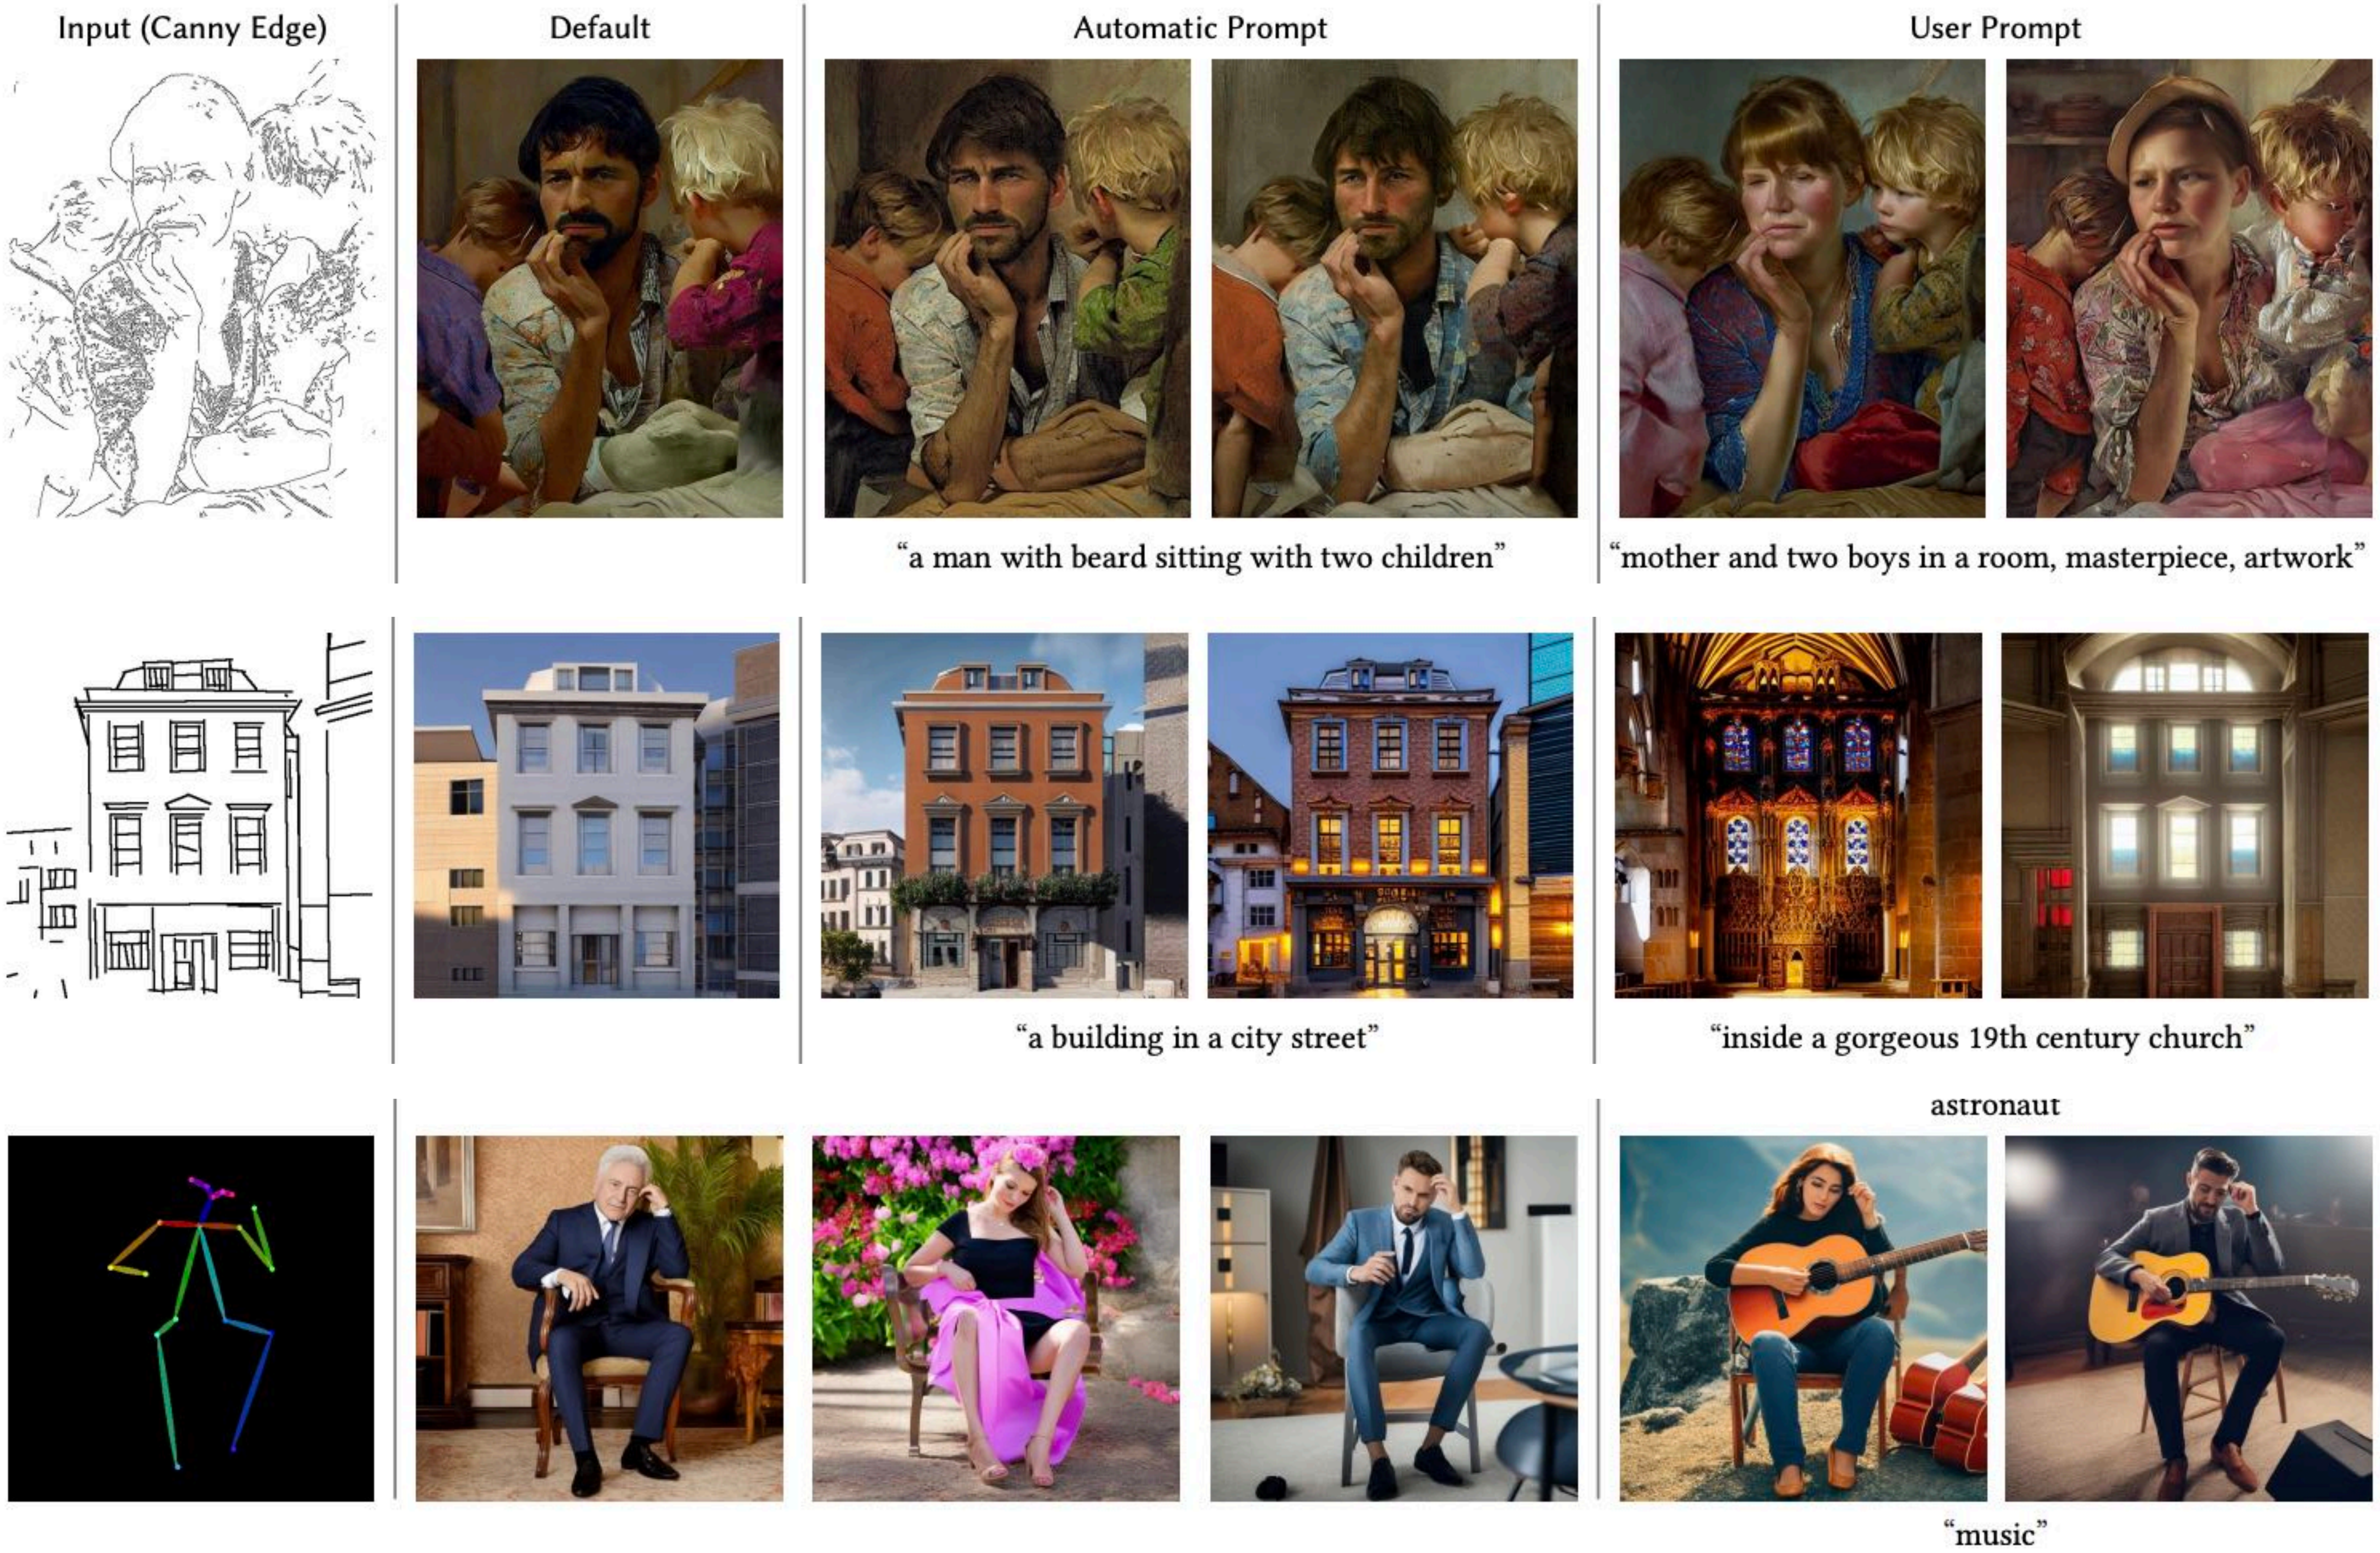

General theme: intuitive controls for content creation

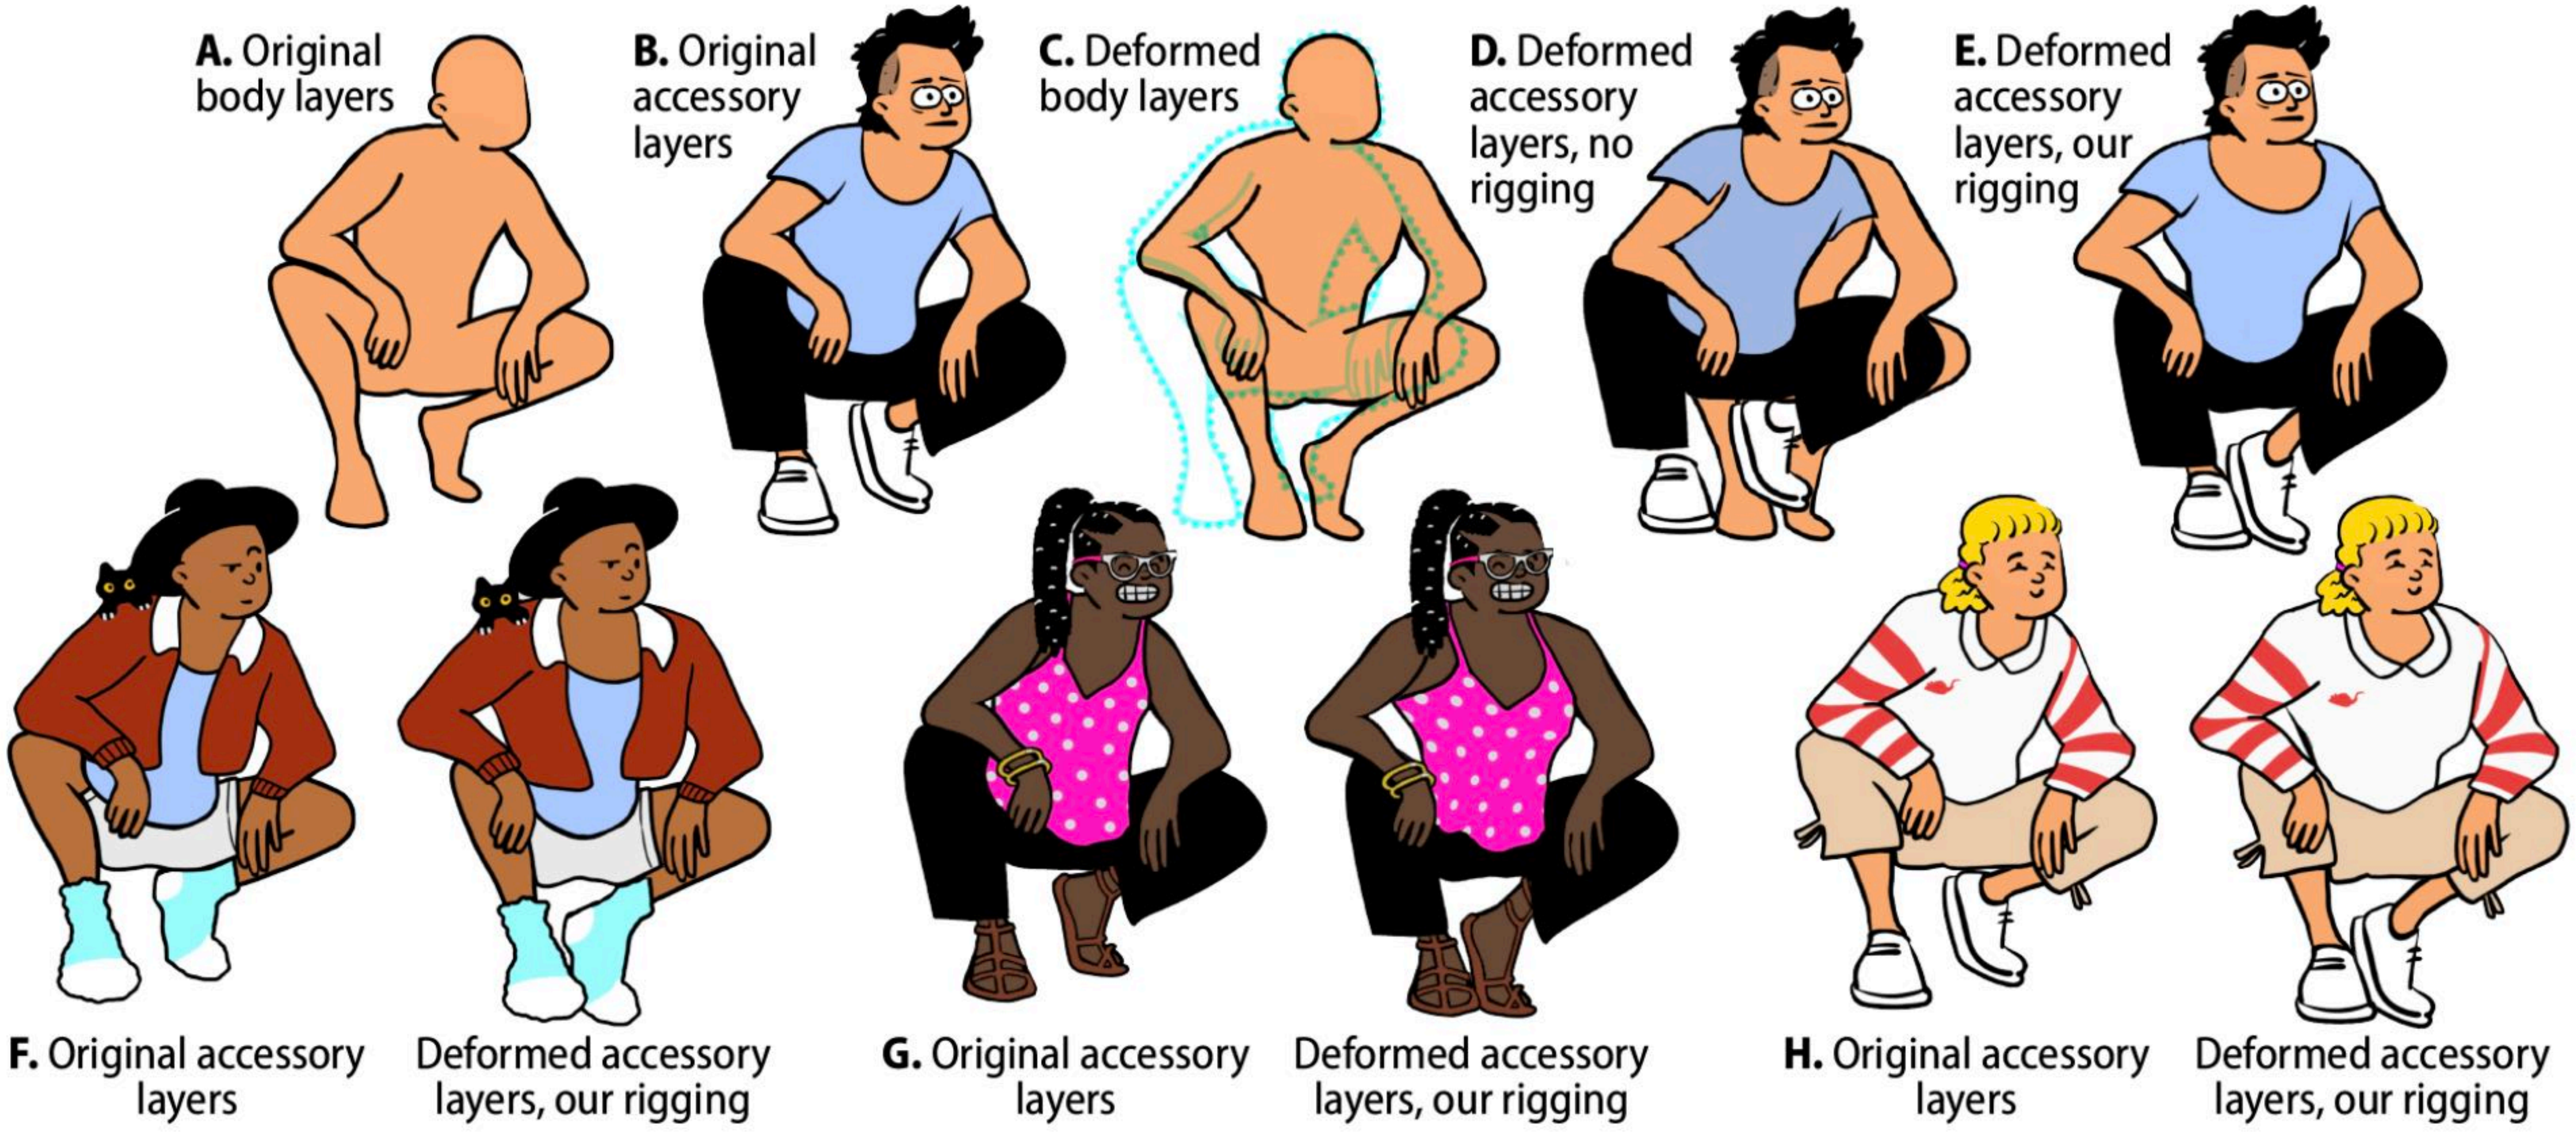

[Le et al. UIST 21]

Ron Fedkiw

Ron Fedkiw

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

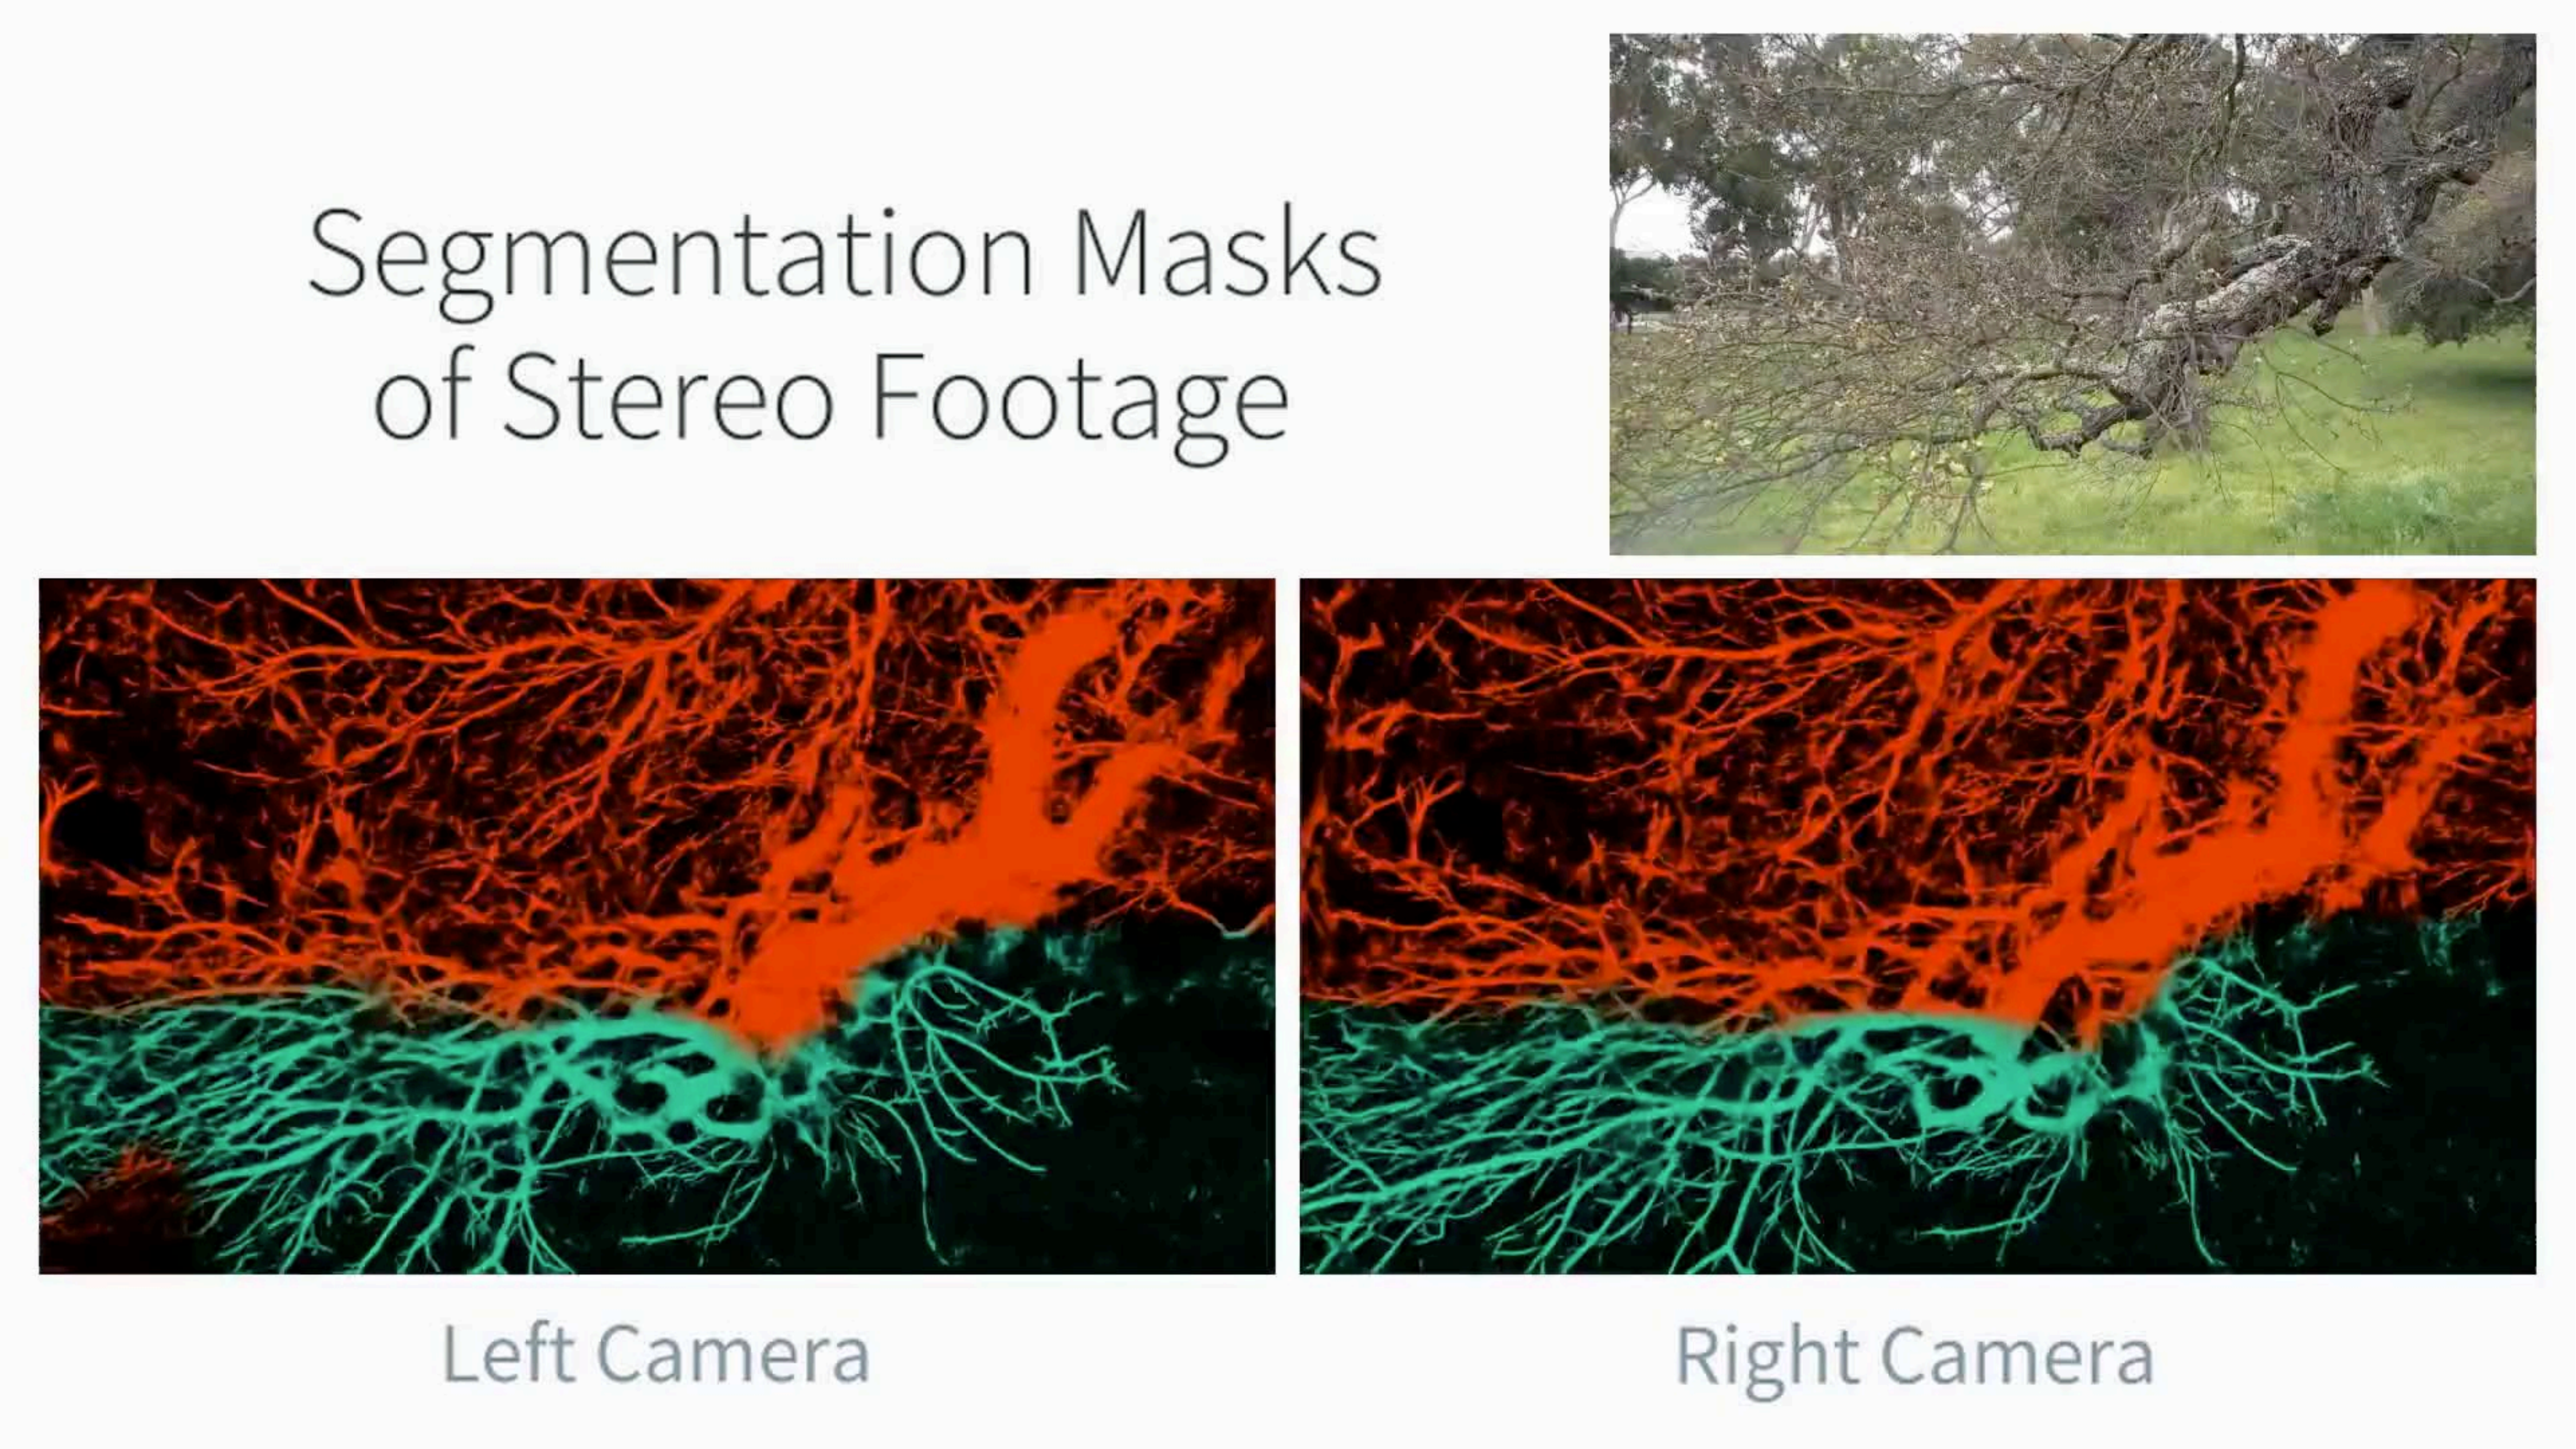

Ron Fedkiw

Karen Liu

Human motion synthesis and estimation

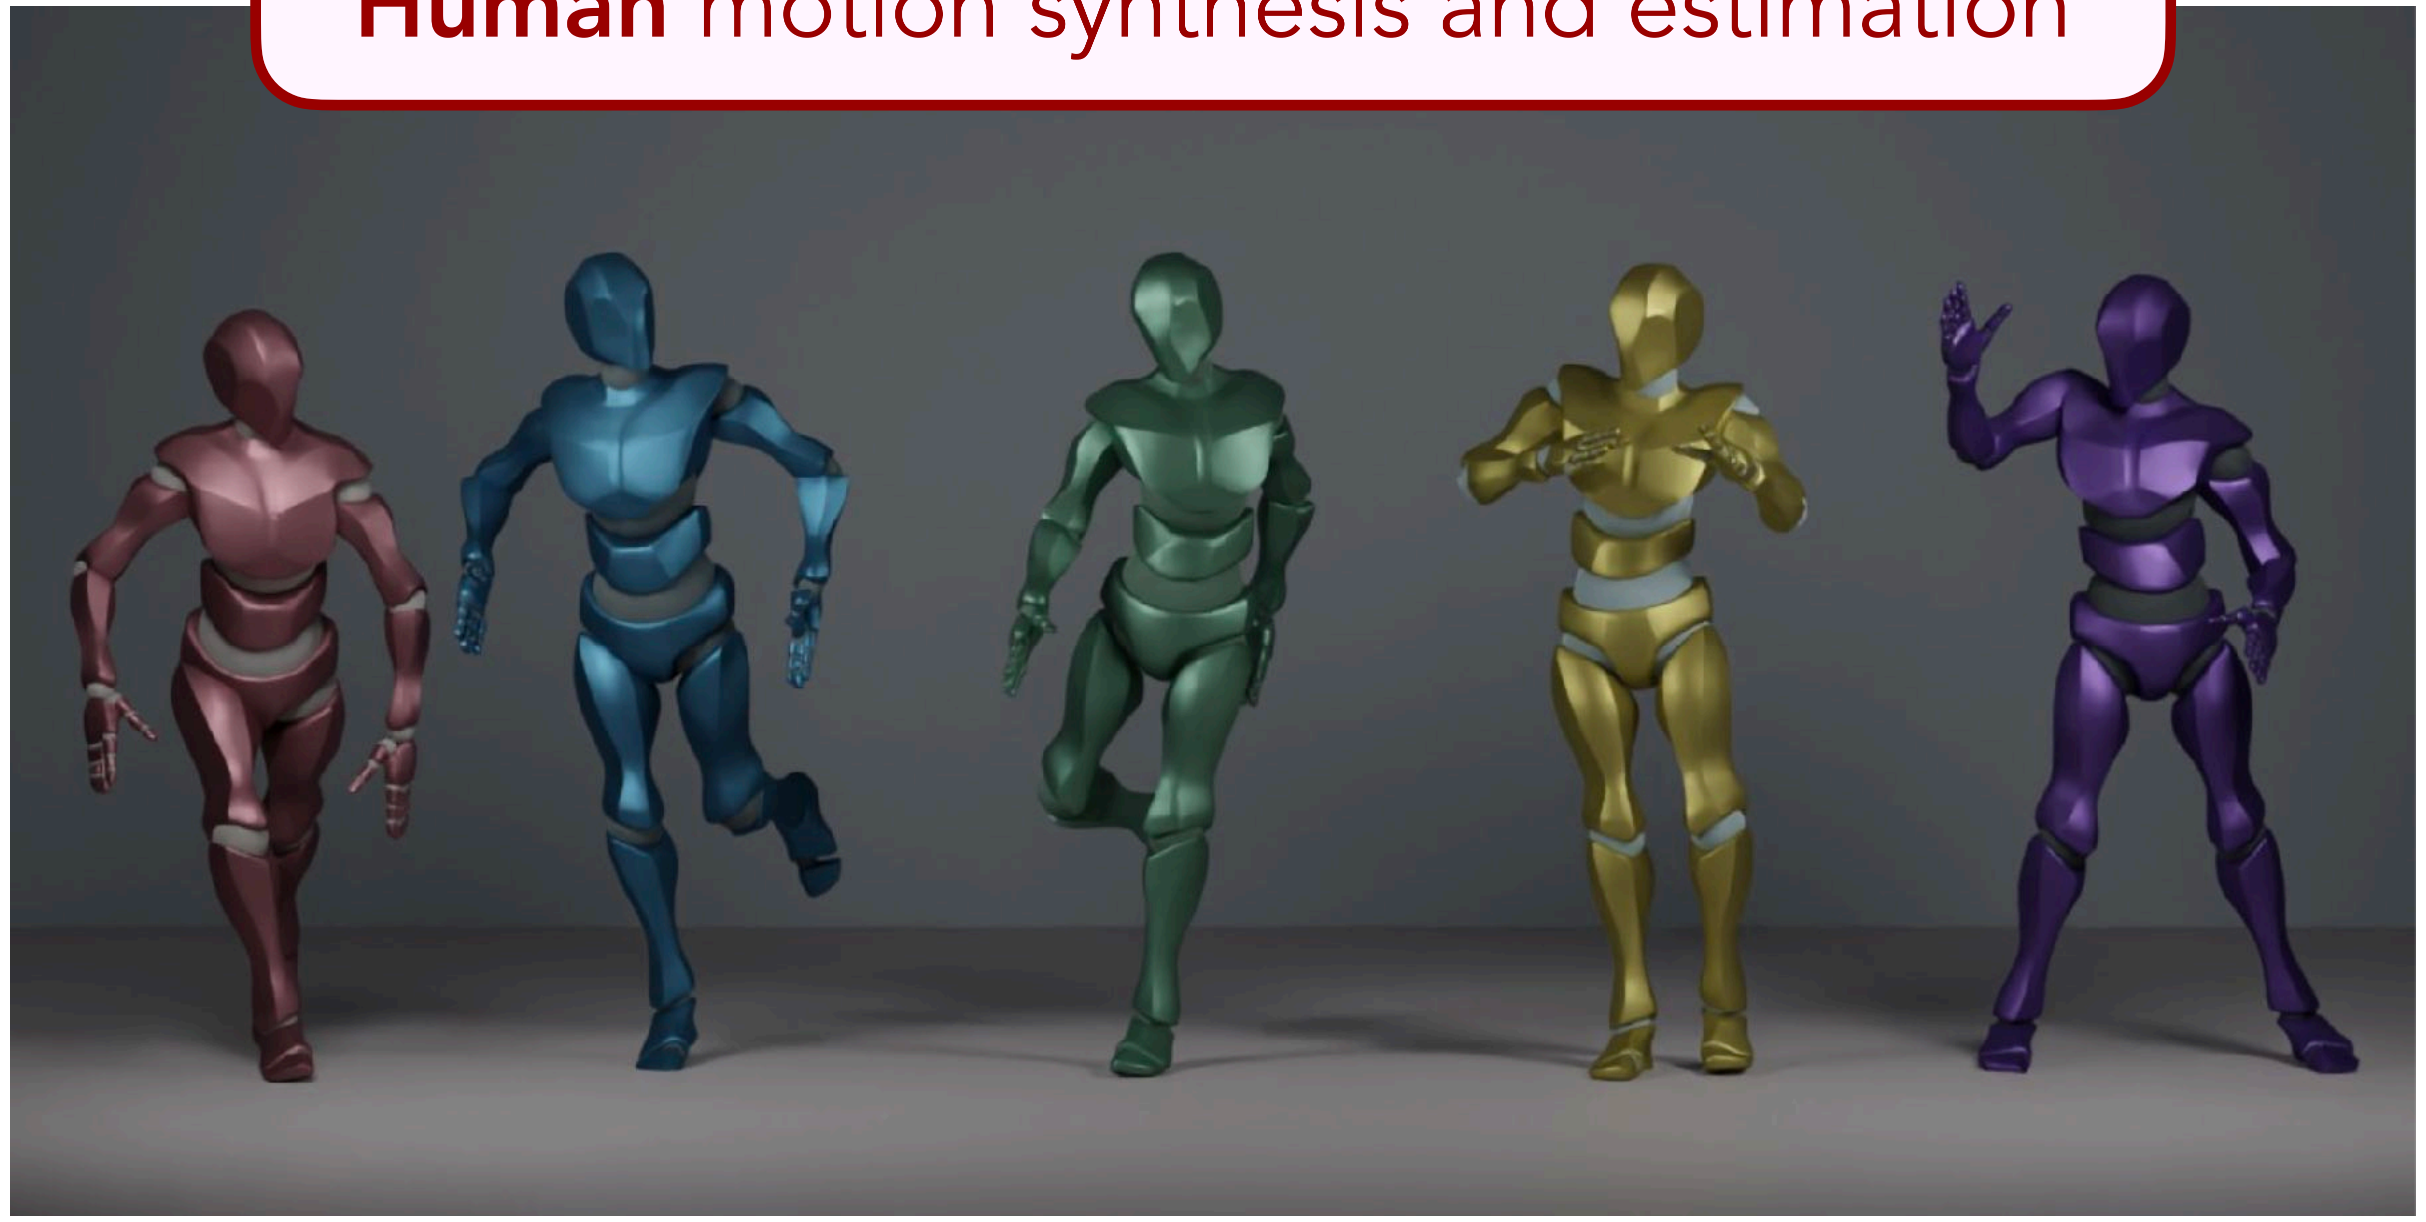

Scene-aware **human** behaviors

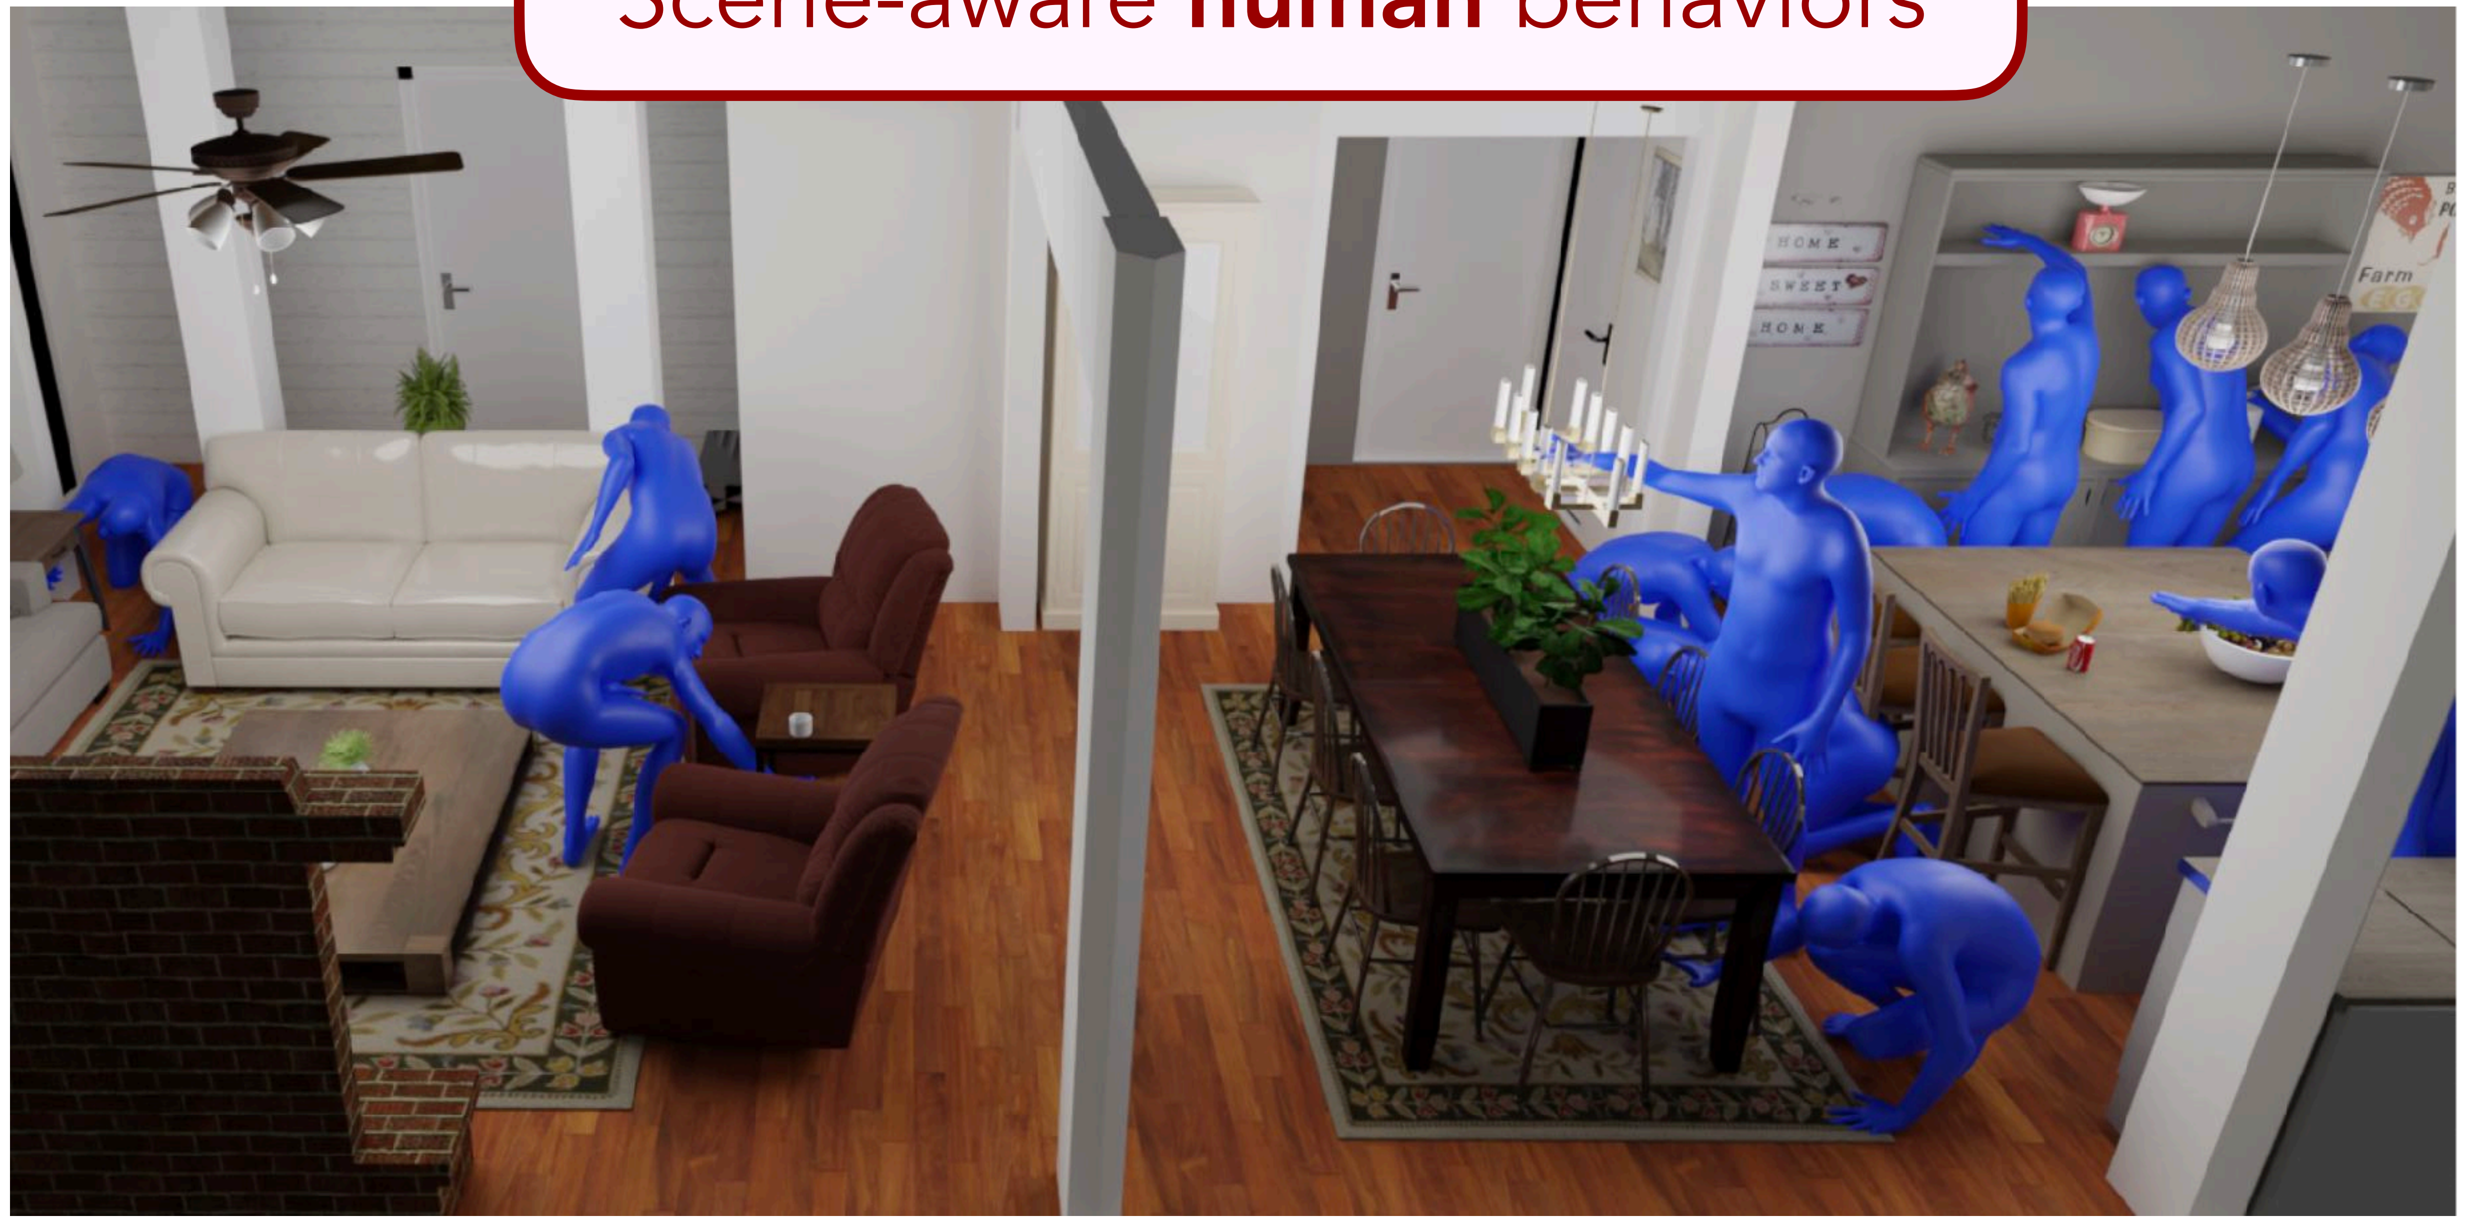

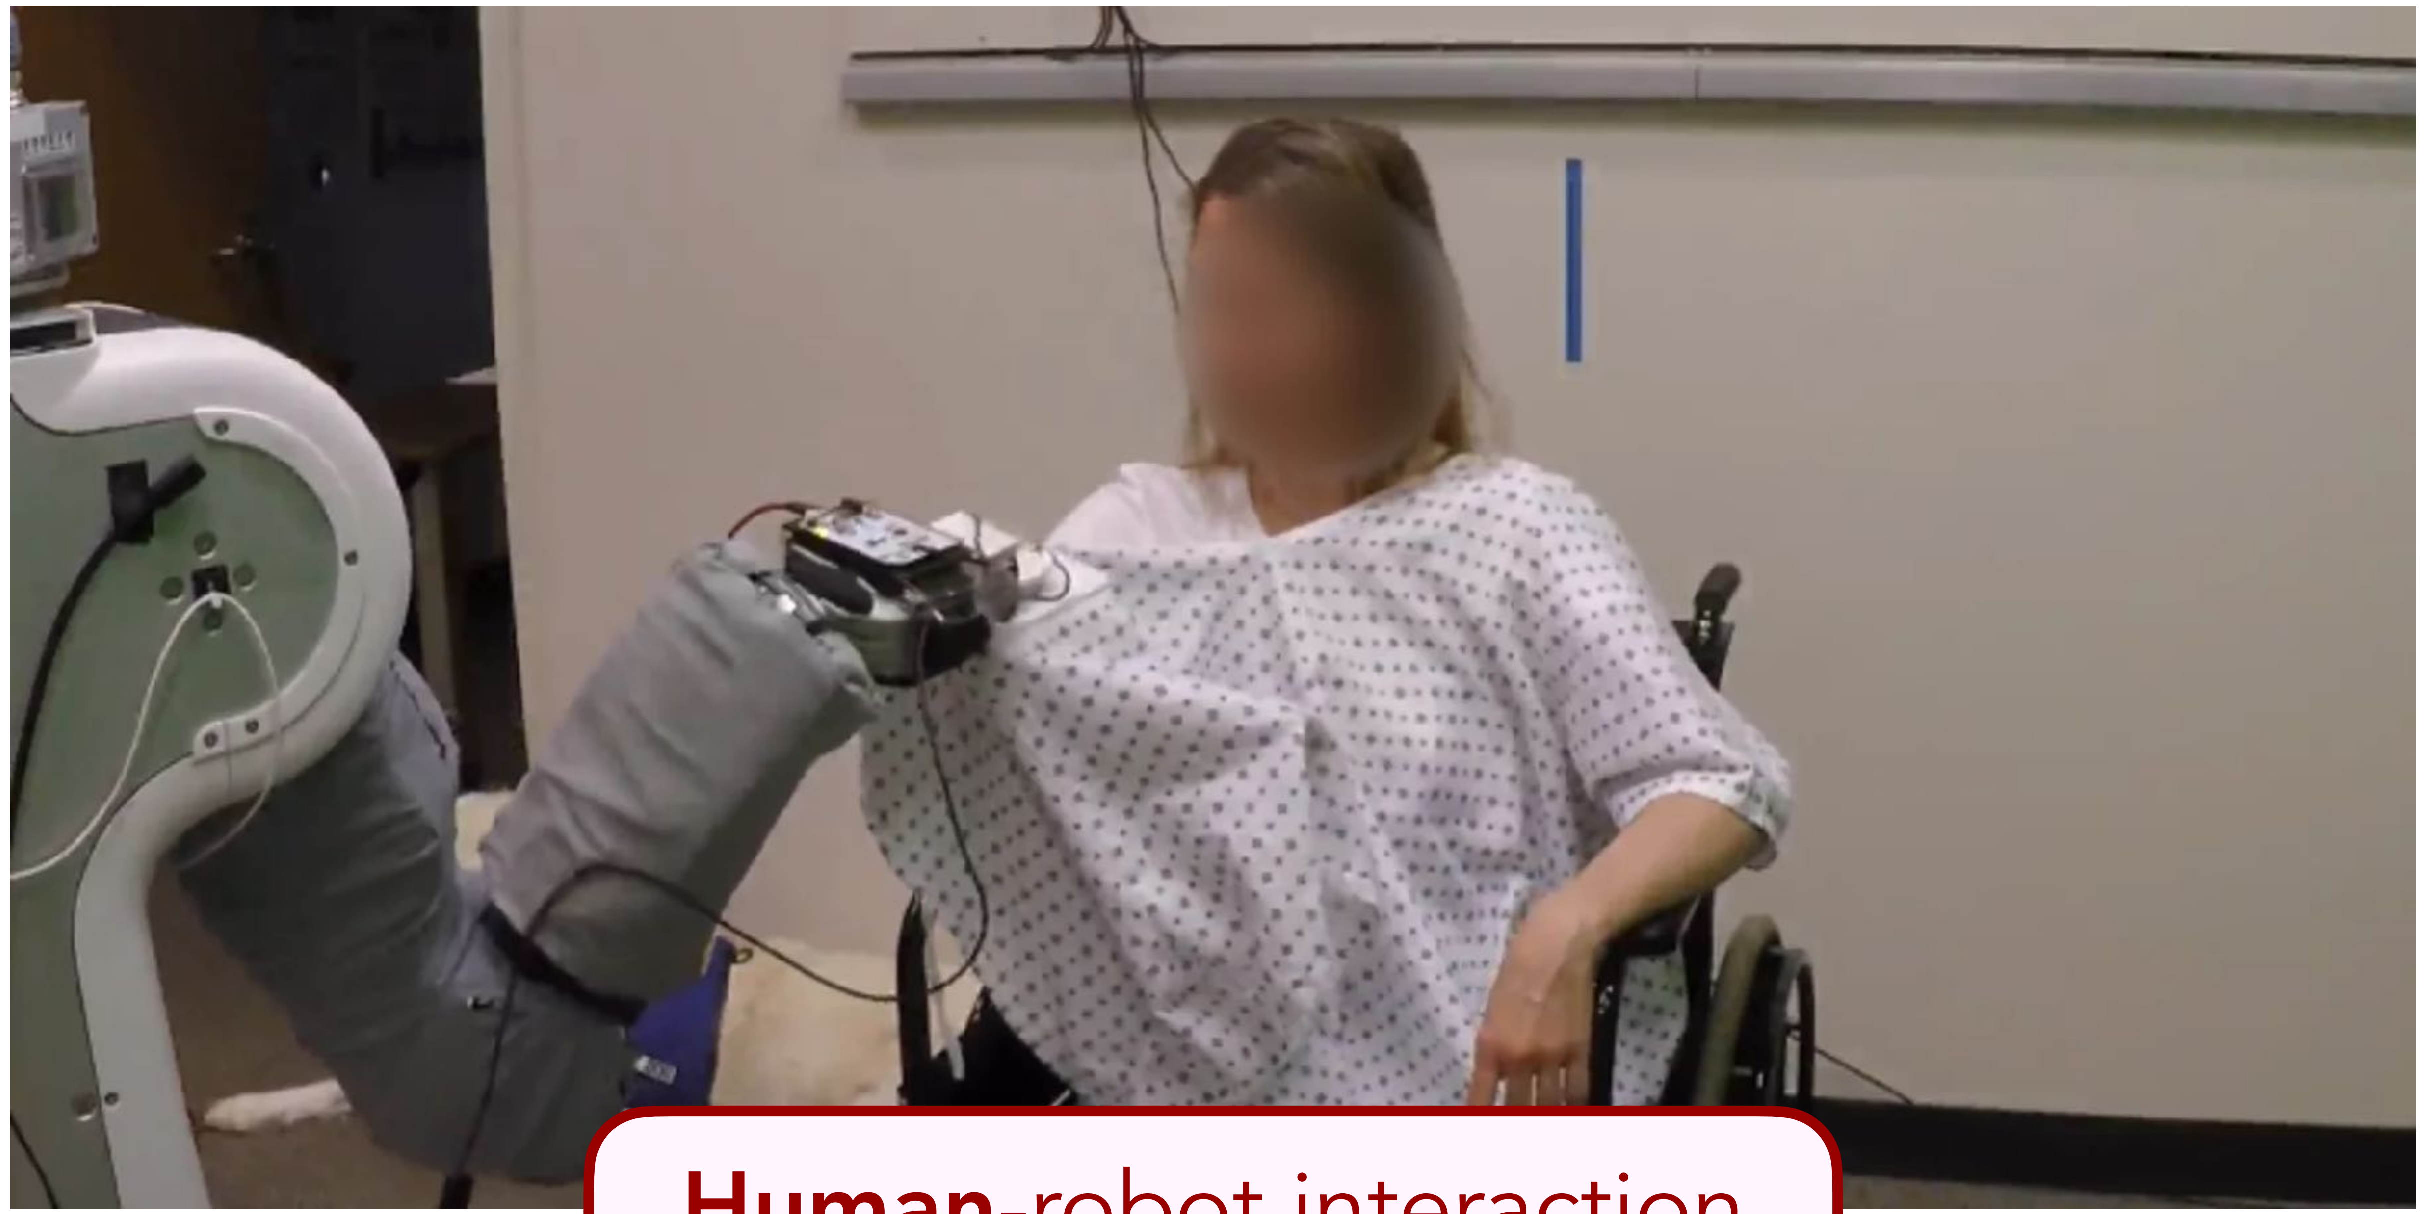

Human-robot interaction

Human dexterity

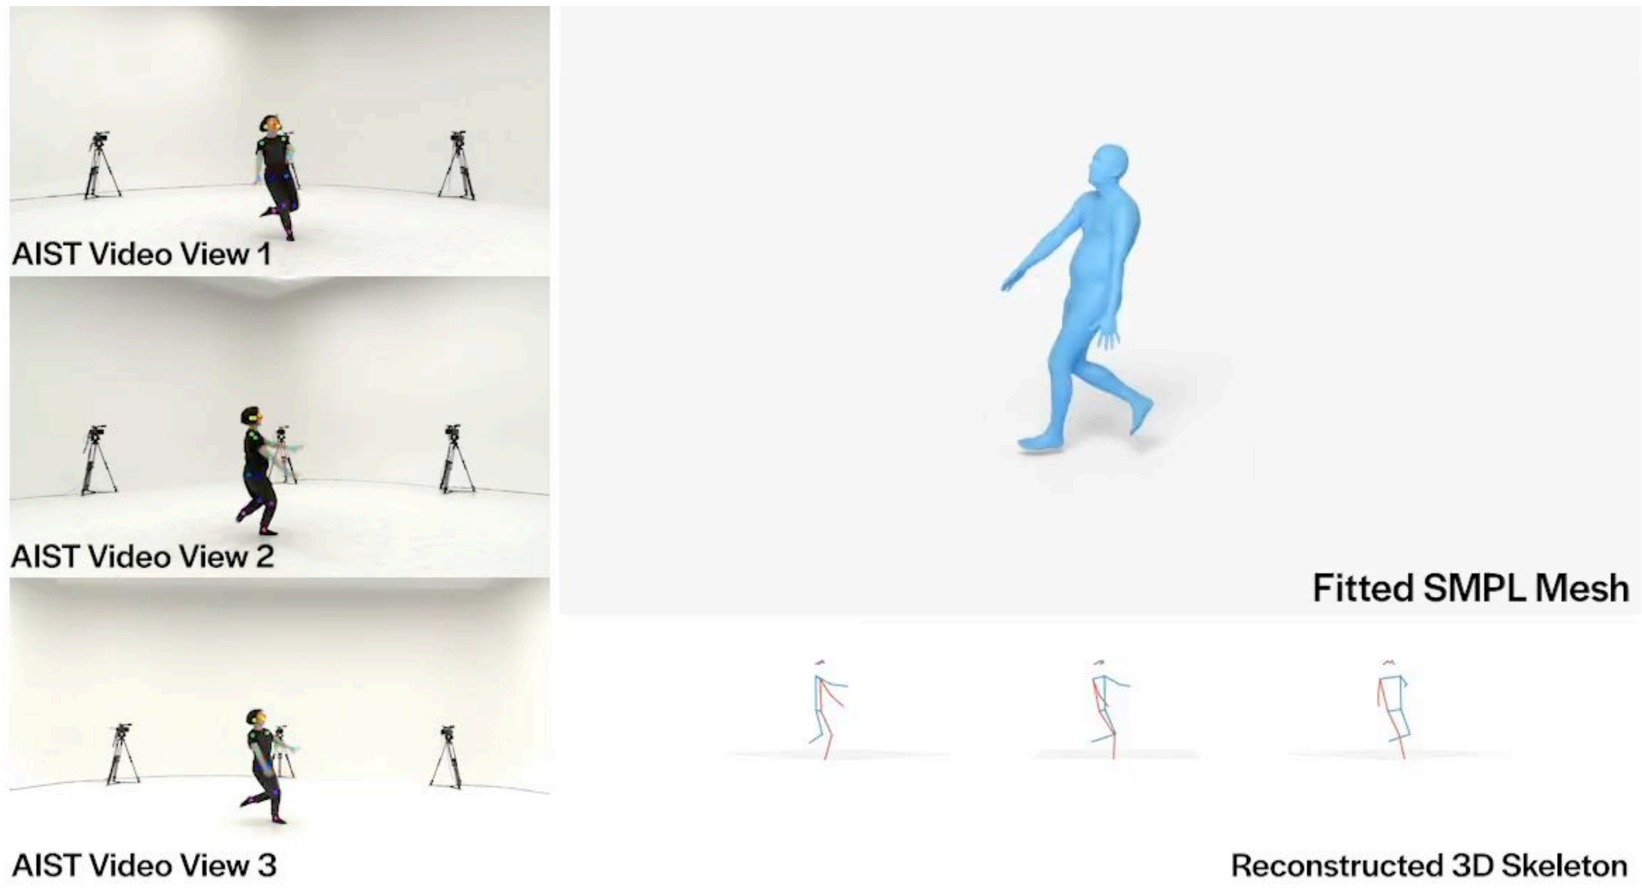

Human dance dataset

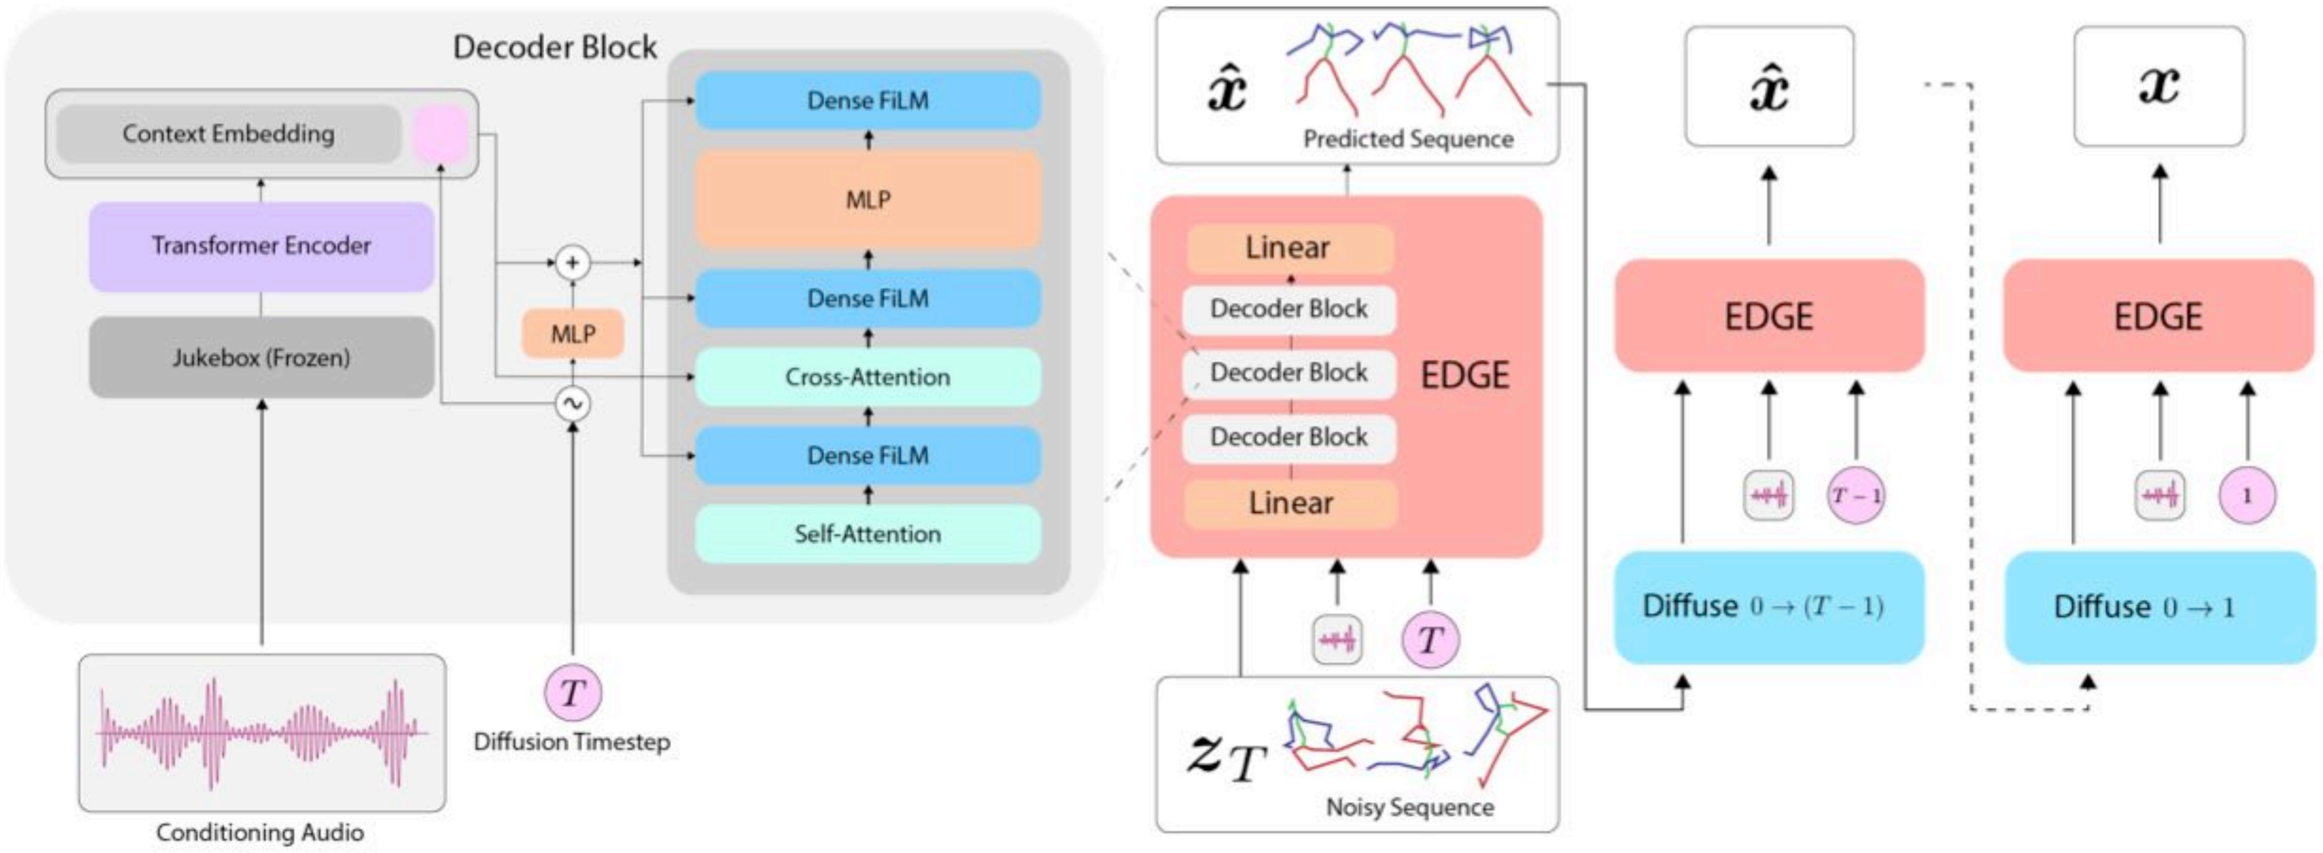

Train an editable generative model to imitate human motor skills and musicality in dance using Diffusion Models

Tseng et al CVPR 2023

Karen Liu

Interests in animation, simulation, and control

Leo Guibas

Recent Guibas Lab projects

- Computer vision and sensor networks
- Geometric and topological data analysis
- 3D machine learning and 3D representations
- 3D shape/scene analysis and synthesis
- Neural methods for navigation and manipulation
- Affective computing

Algorithmic problems in modeling physical objects and phenomena in vision/graphics/robotics

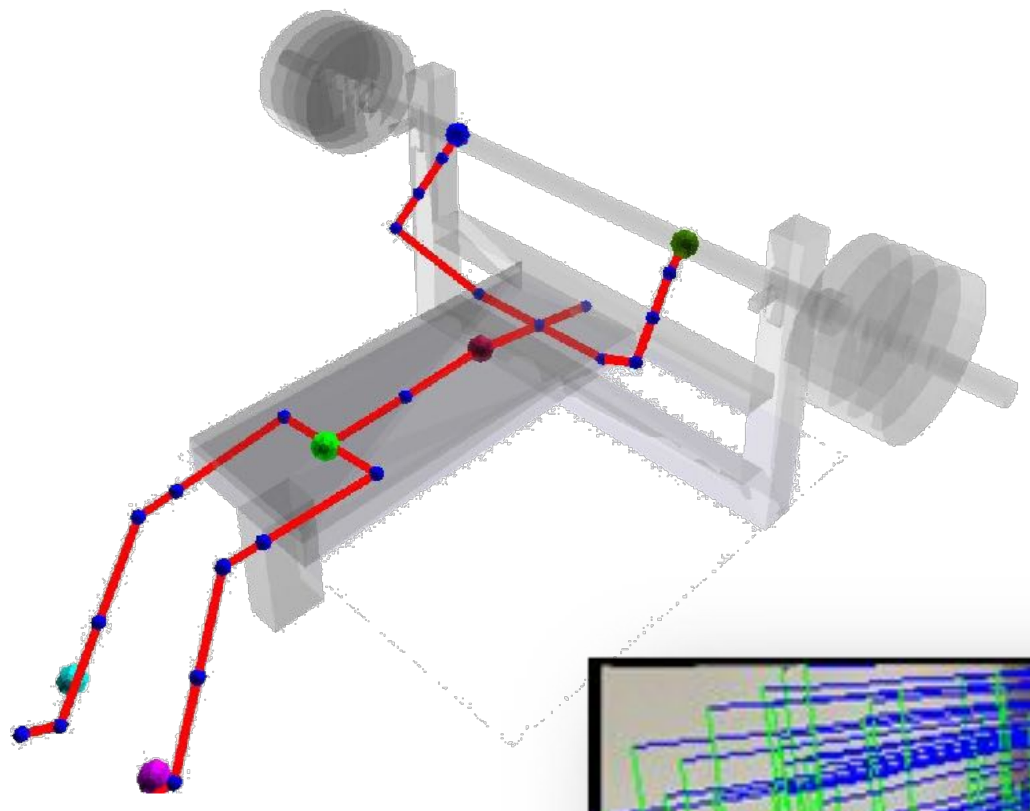

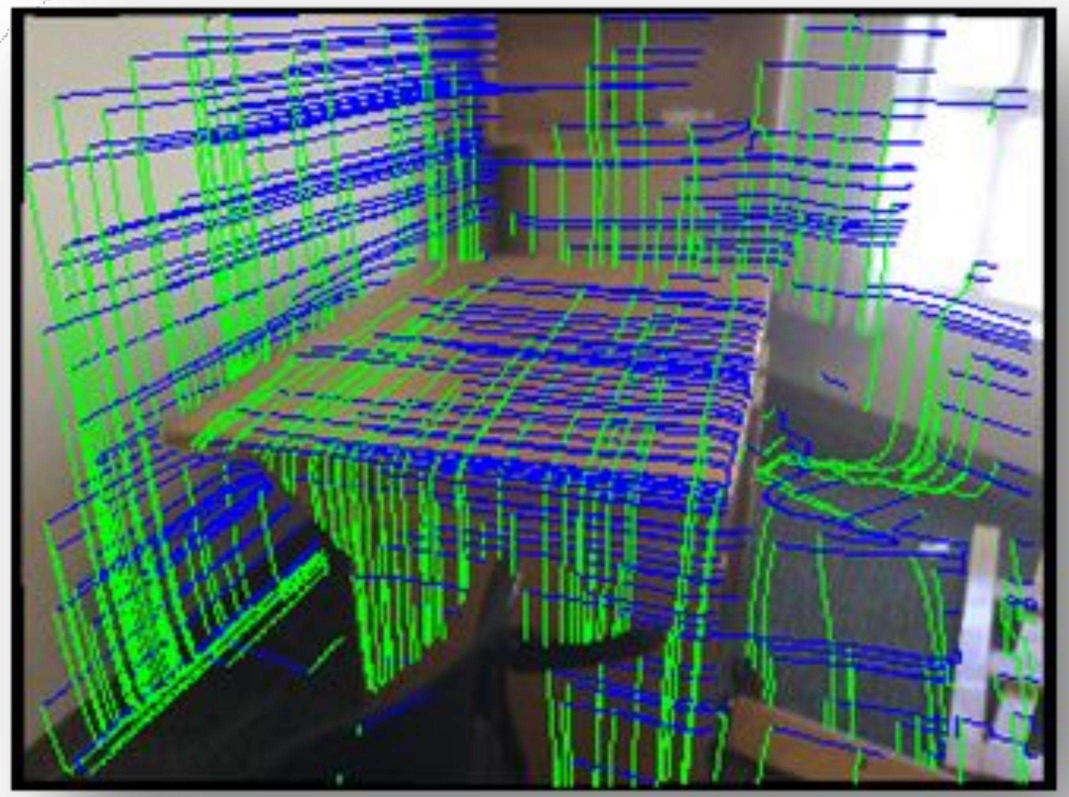

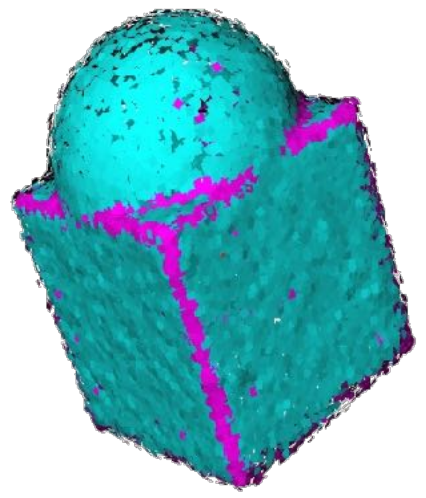

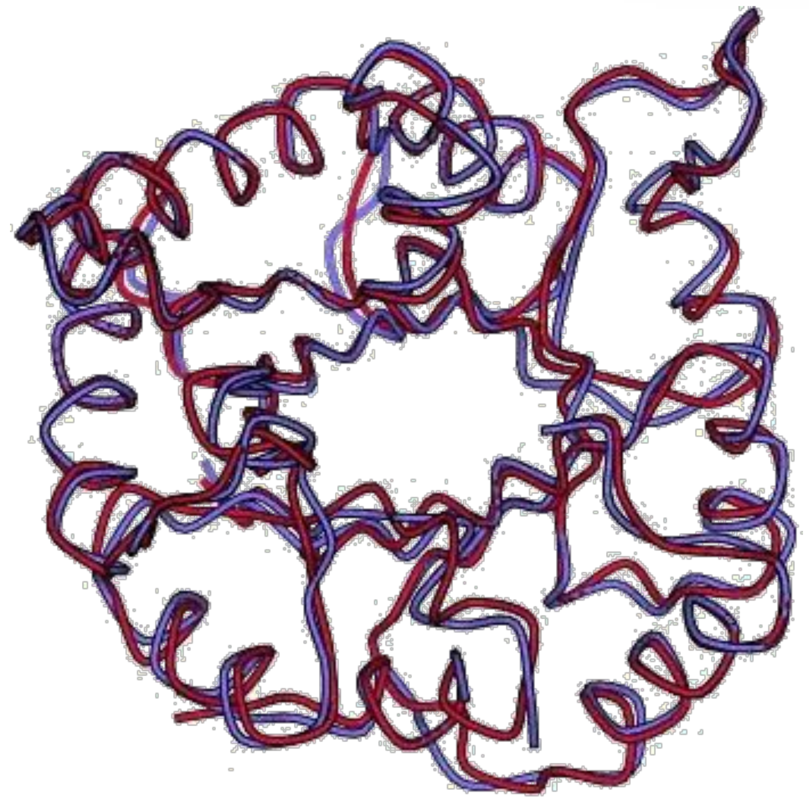

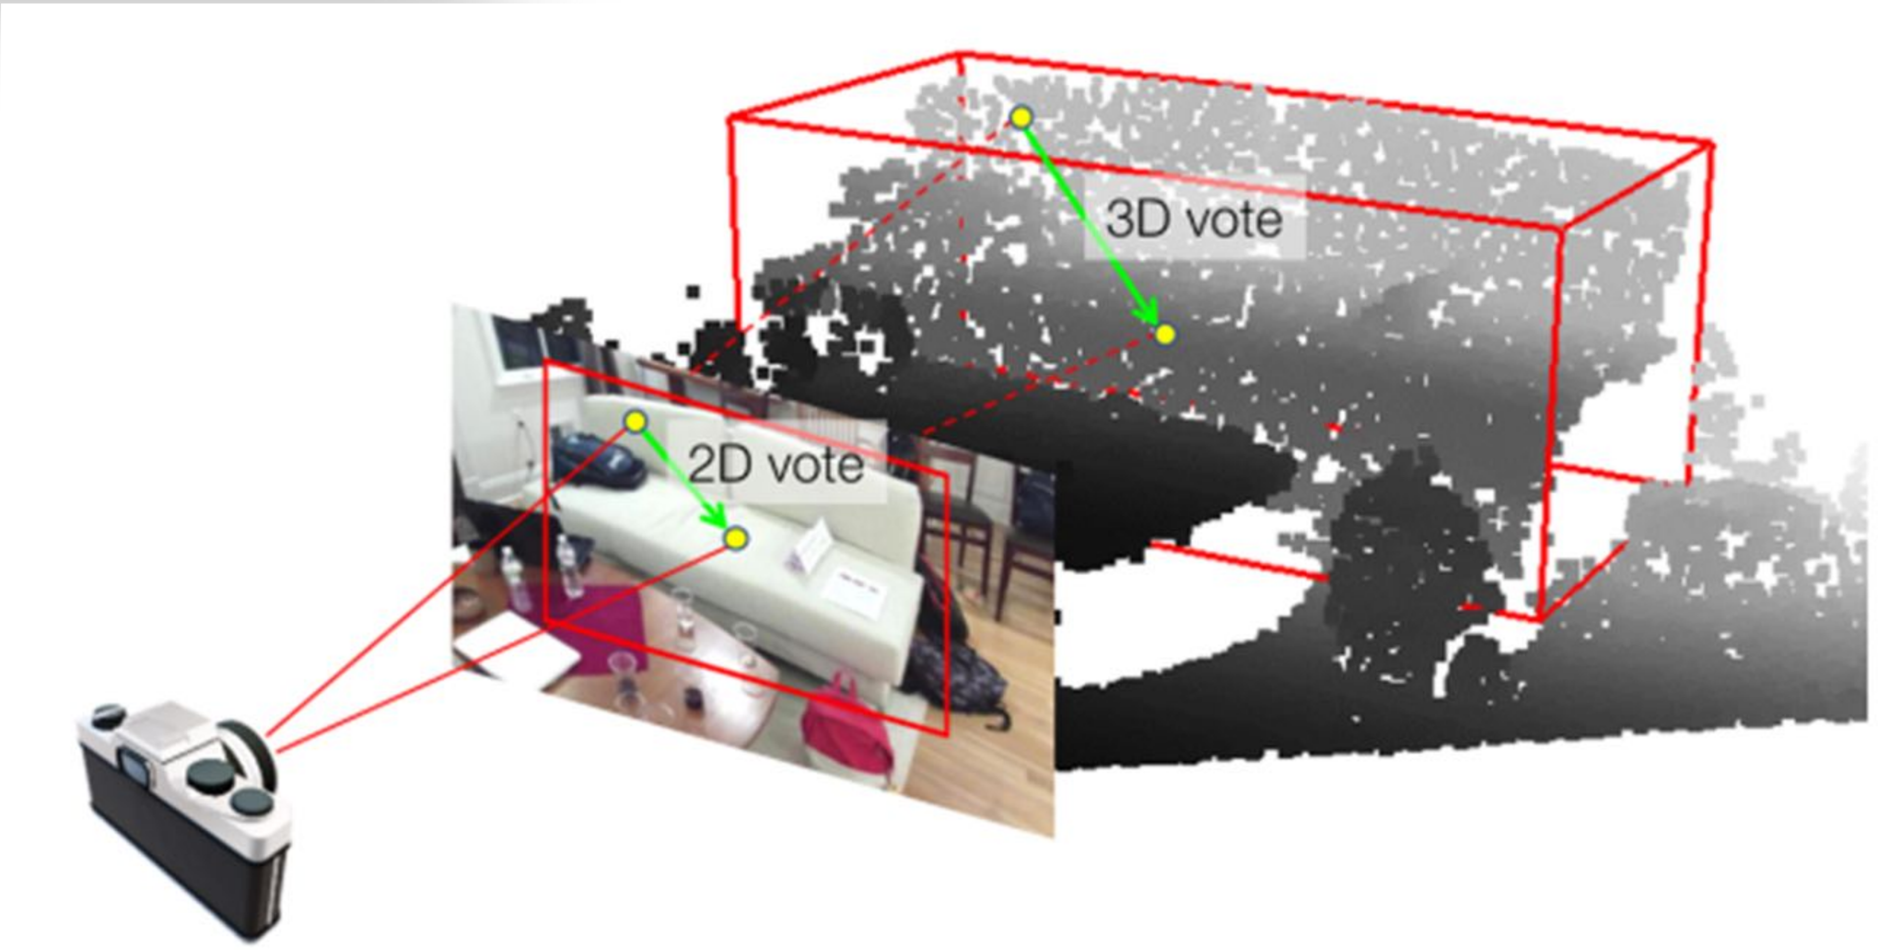

Leo Guibas

Geometry processing and analysis

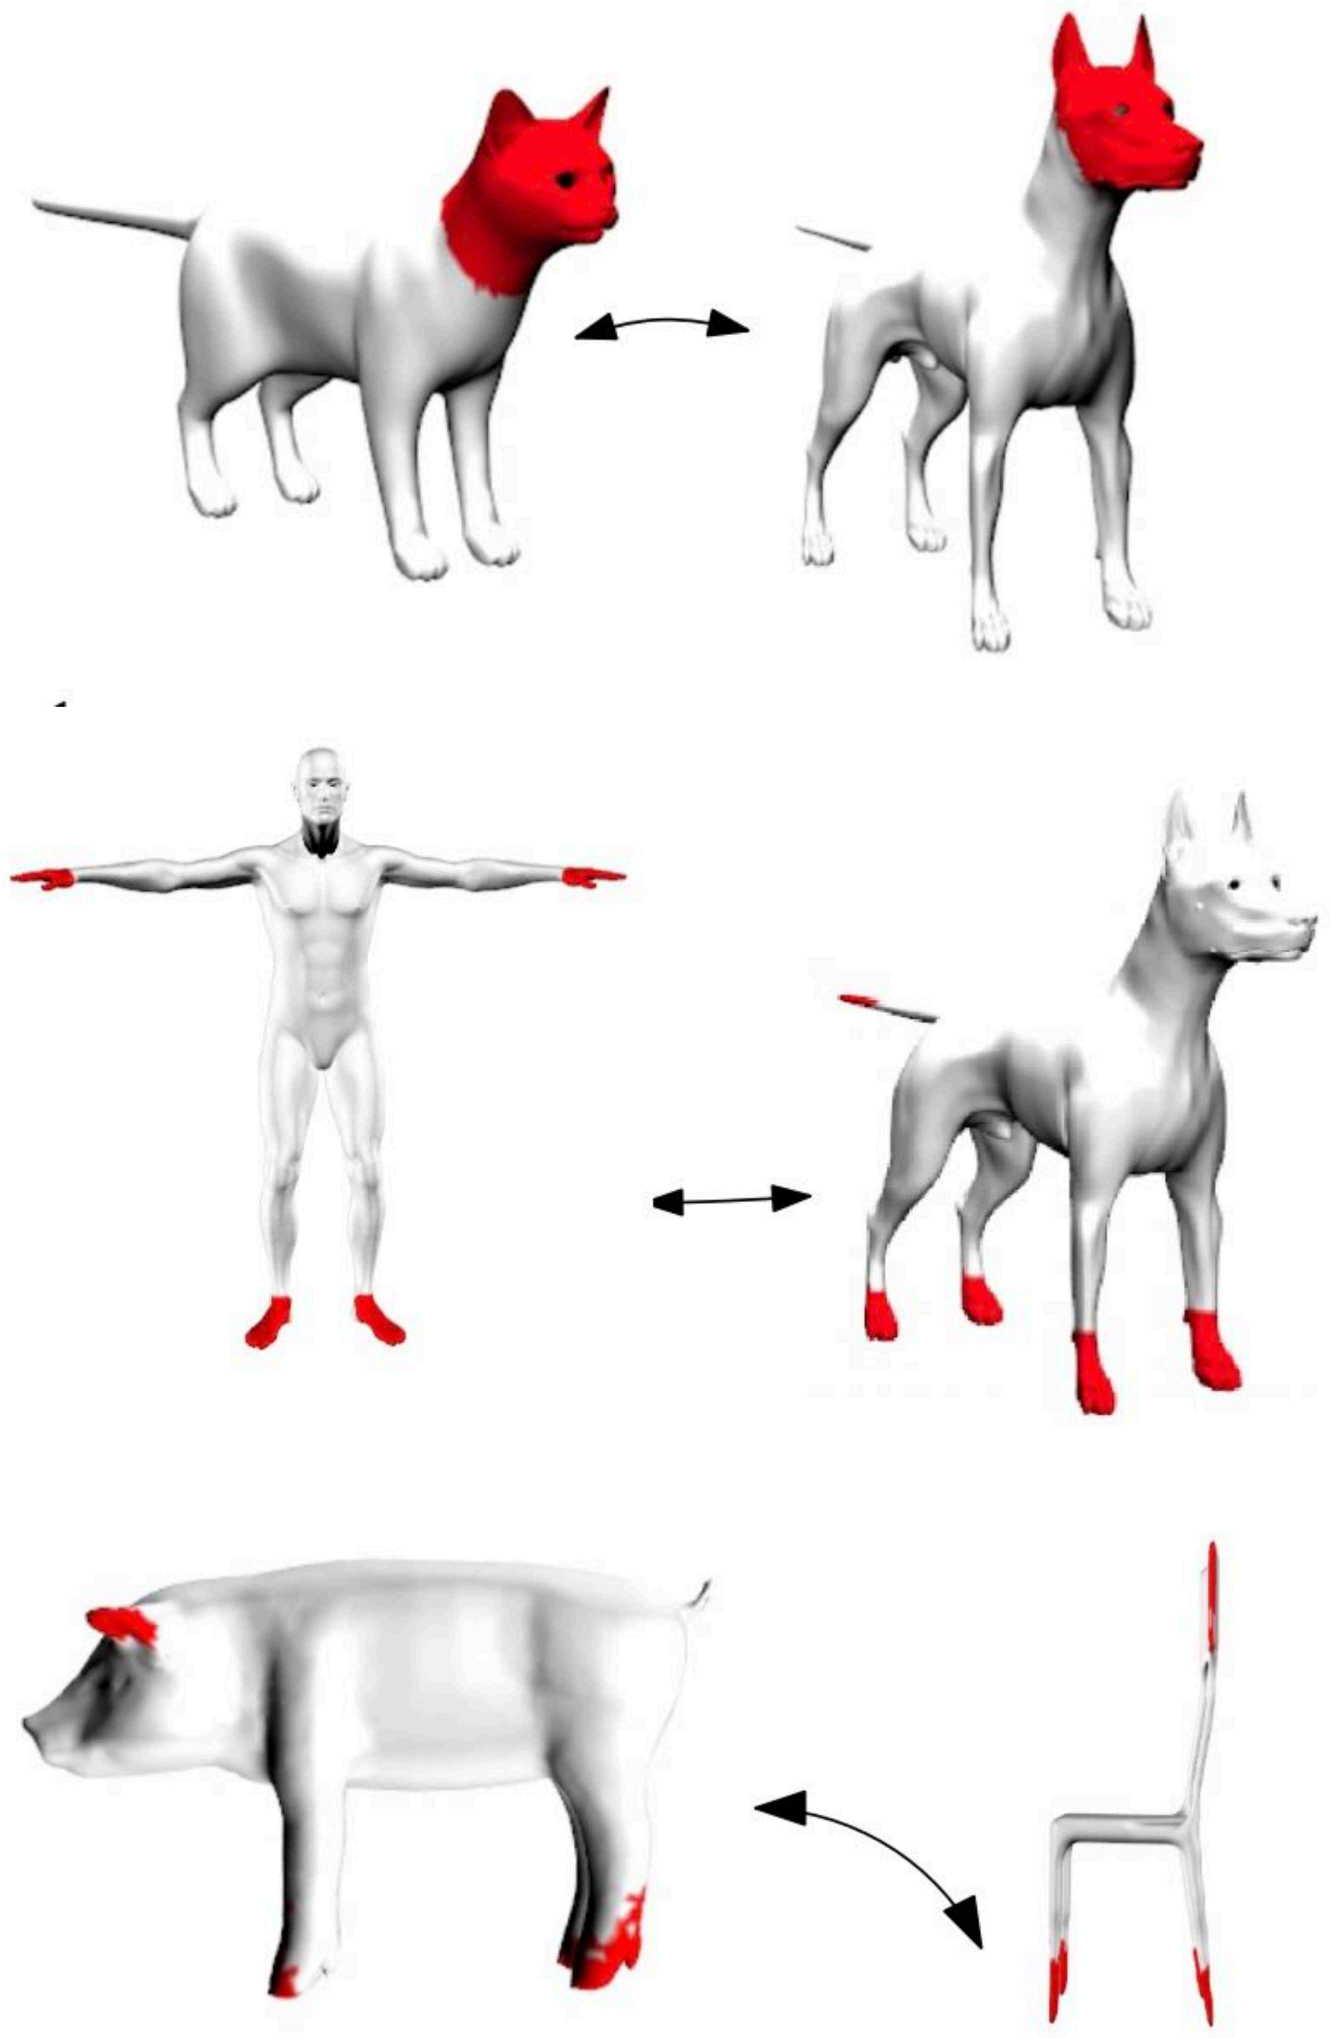

Shape Similarity and Correspondence

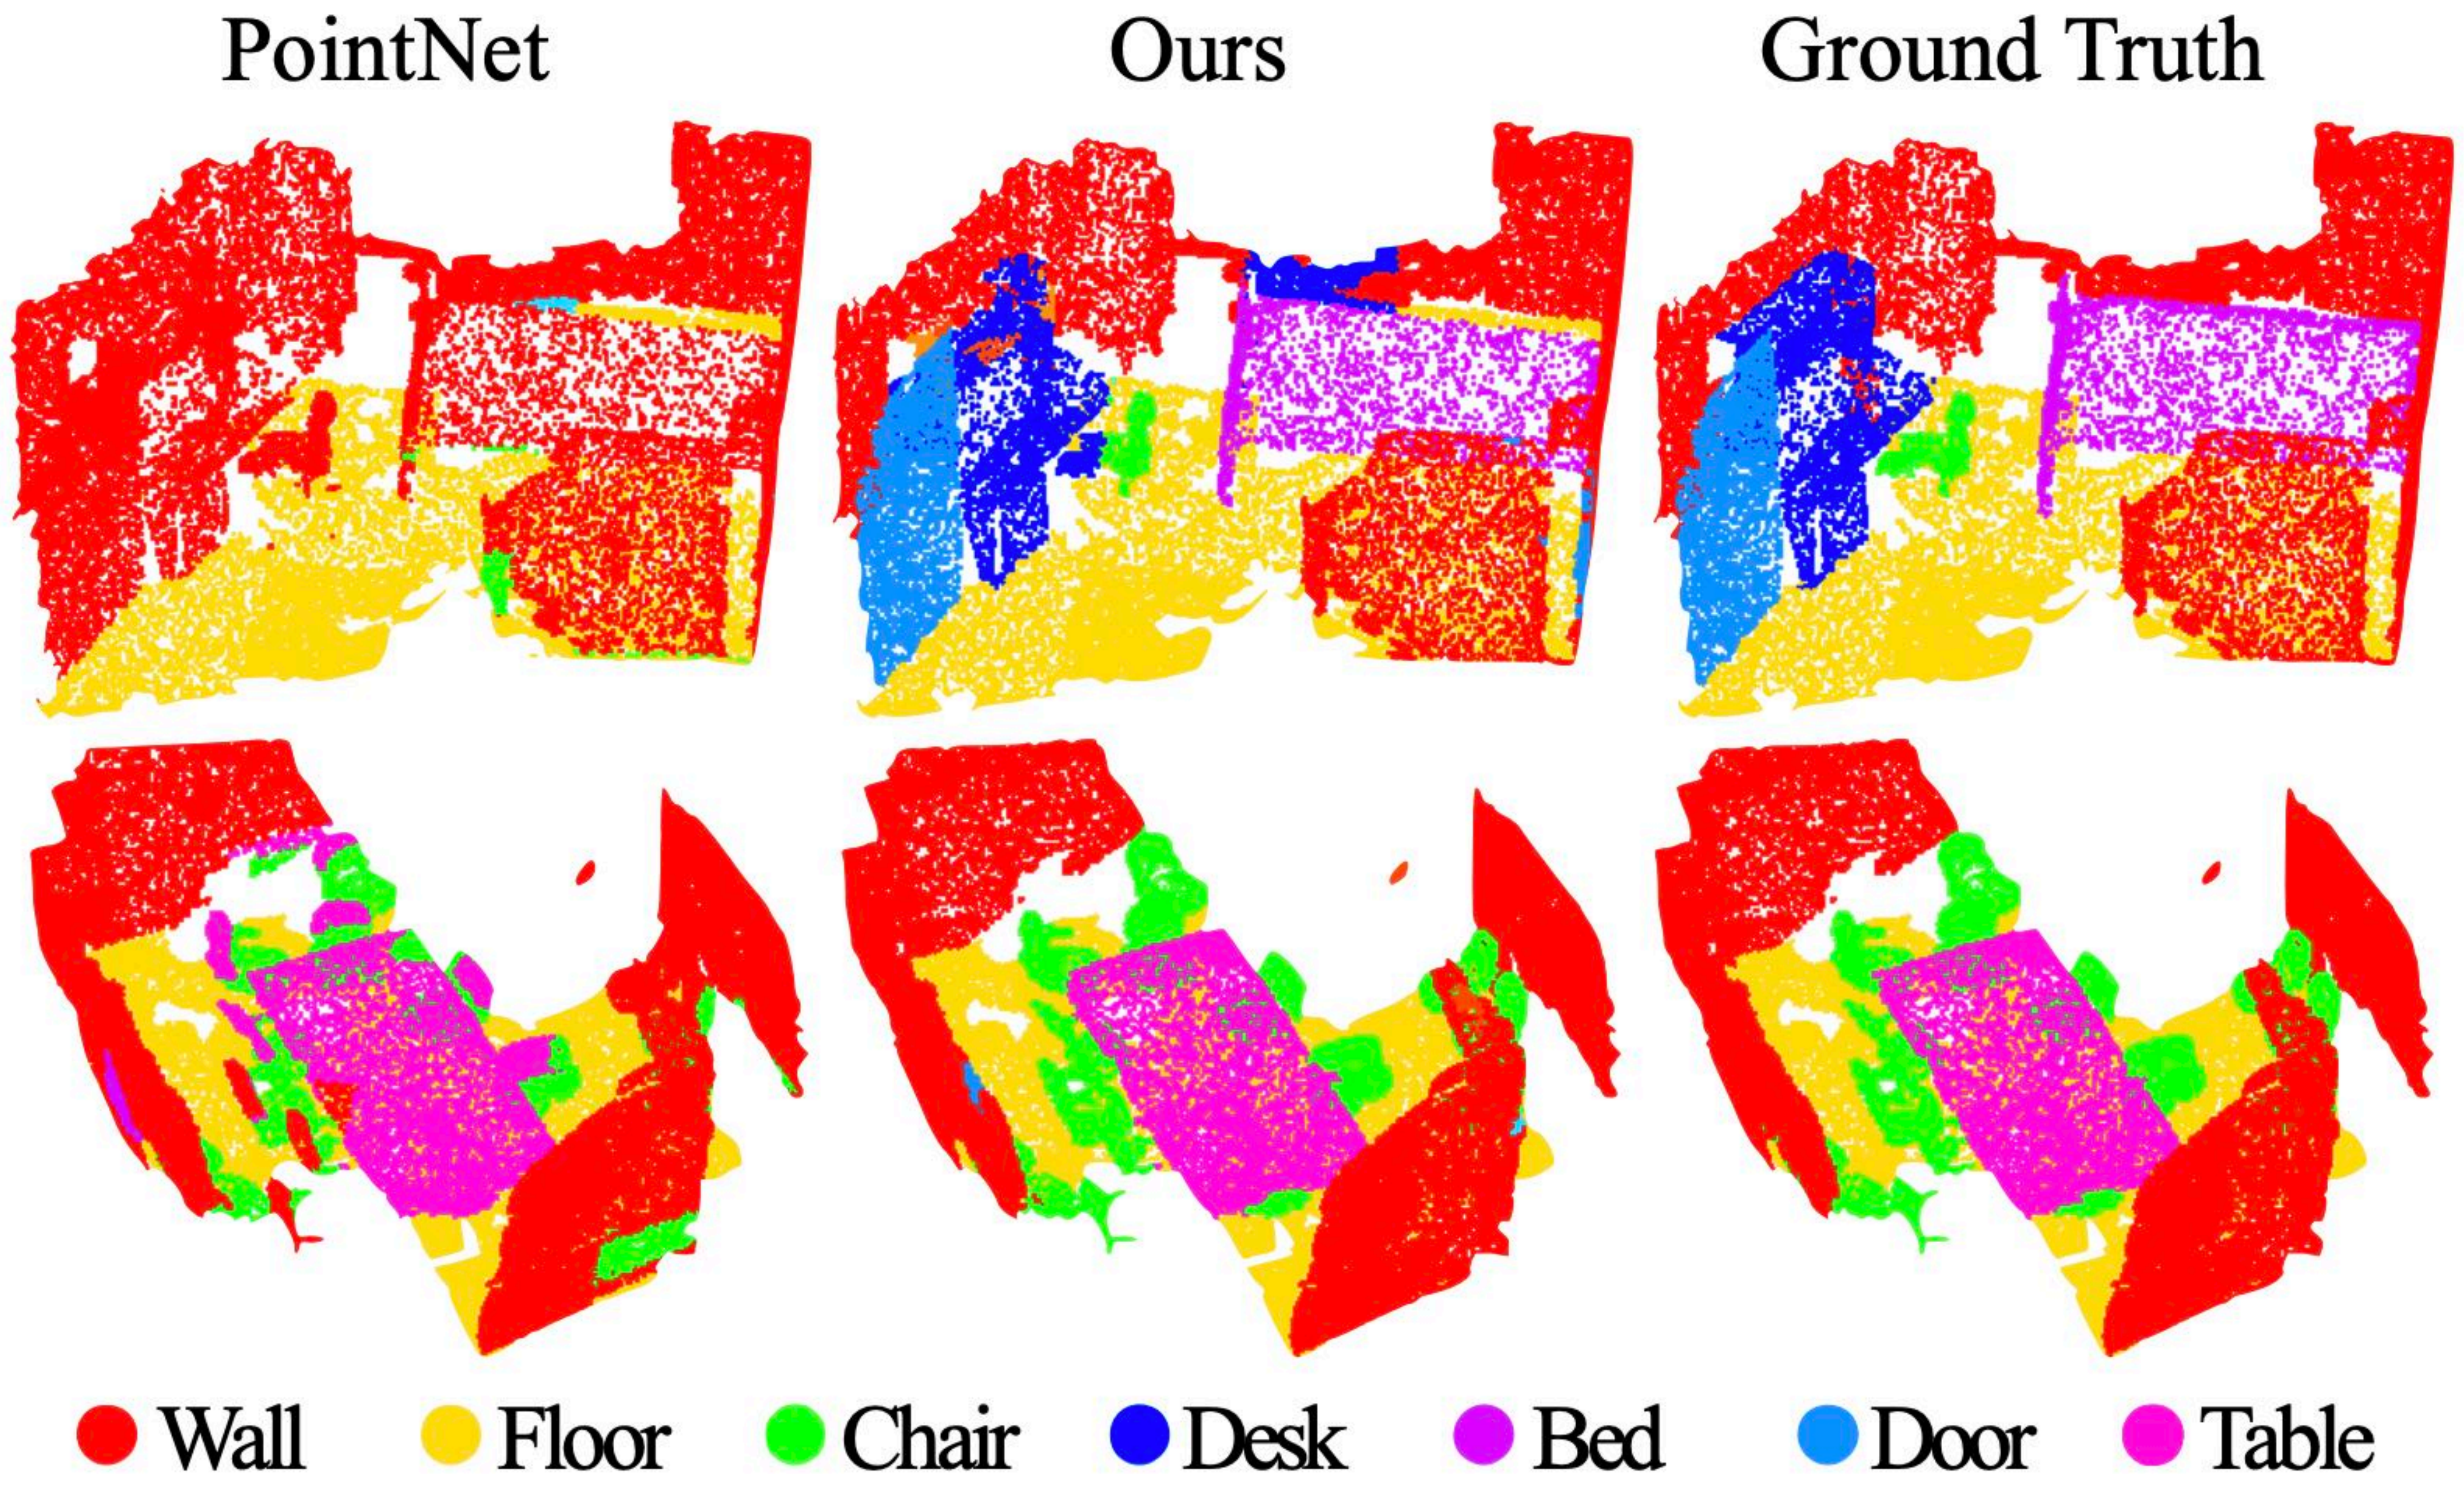

PointNet: Deep Learning on Point Clouds

Gordon Wetzstein

Stanford Computational Imaging Lab

Neural Rendering

XR & Wearable Computing

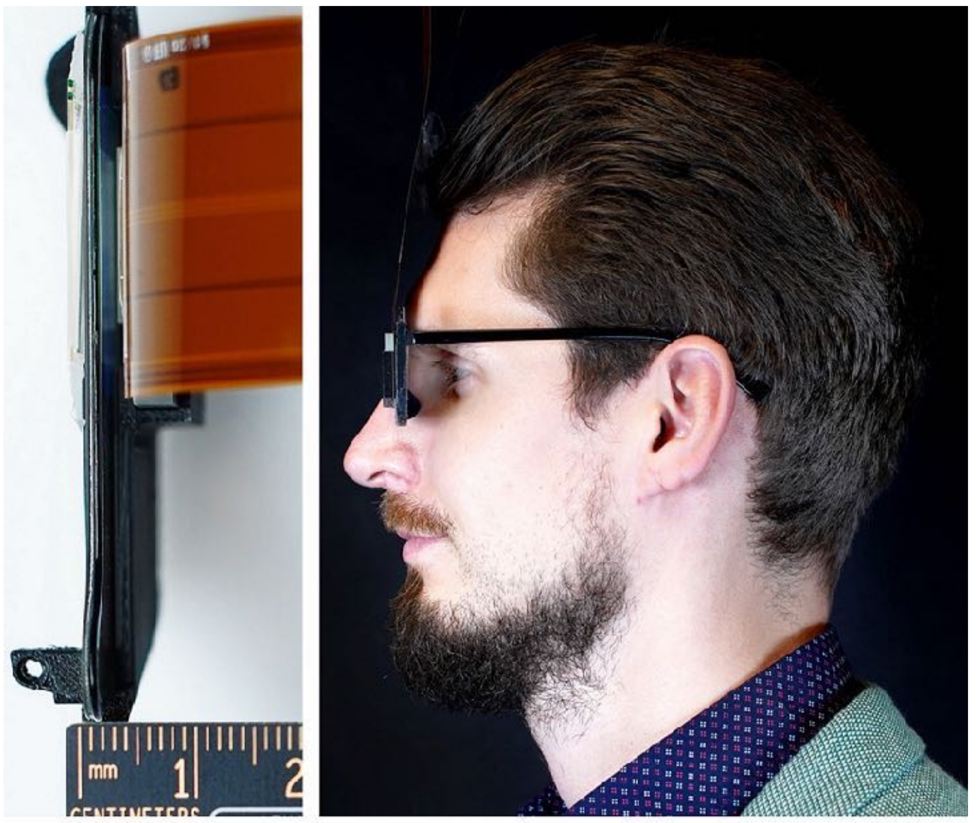

Deep Optics

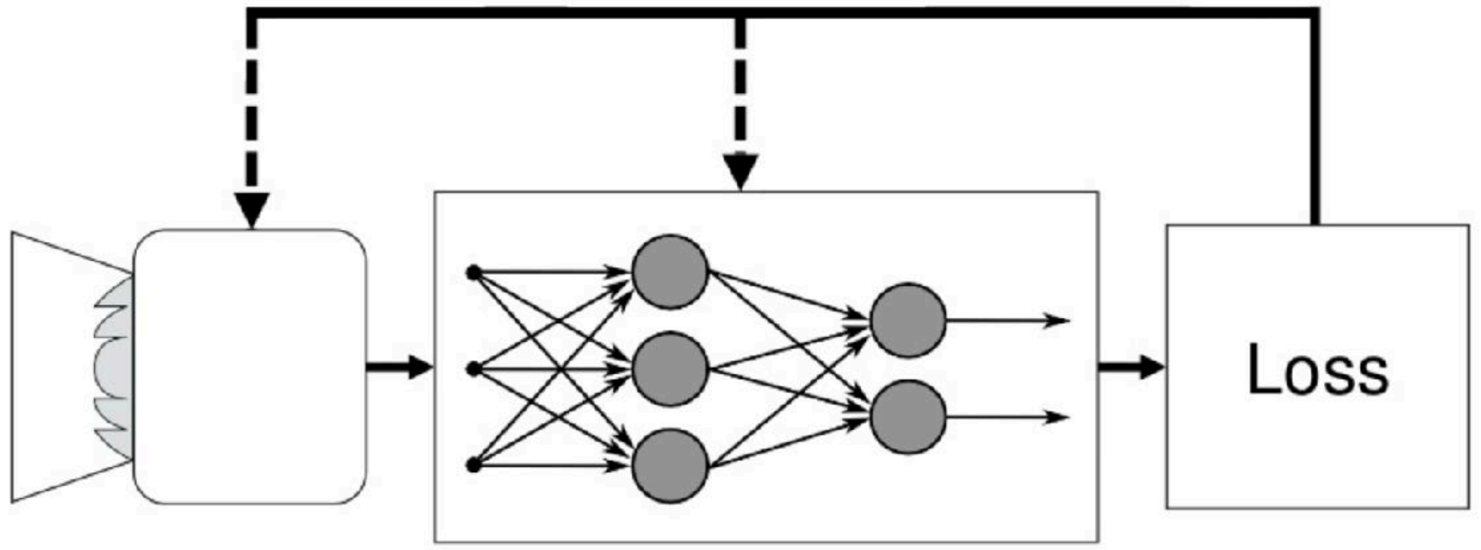

Single-photon Imaging

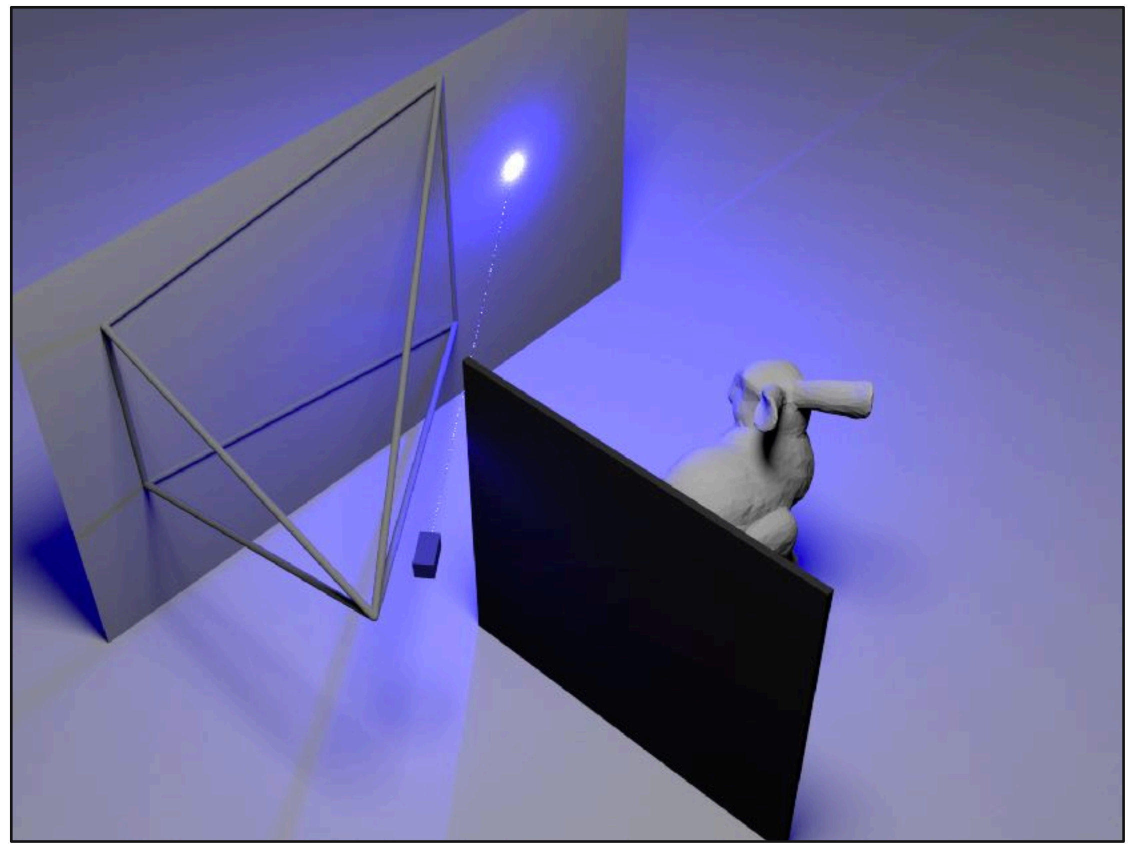

Computational Microscopy

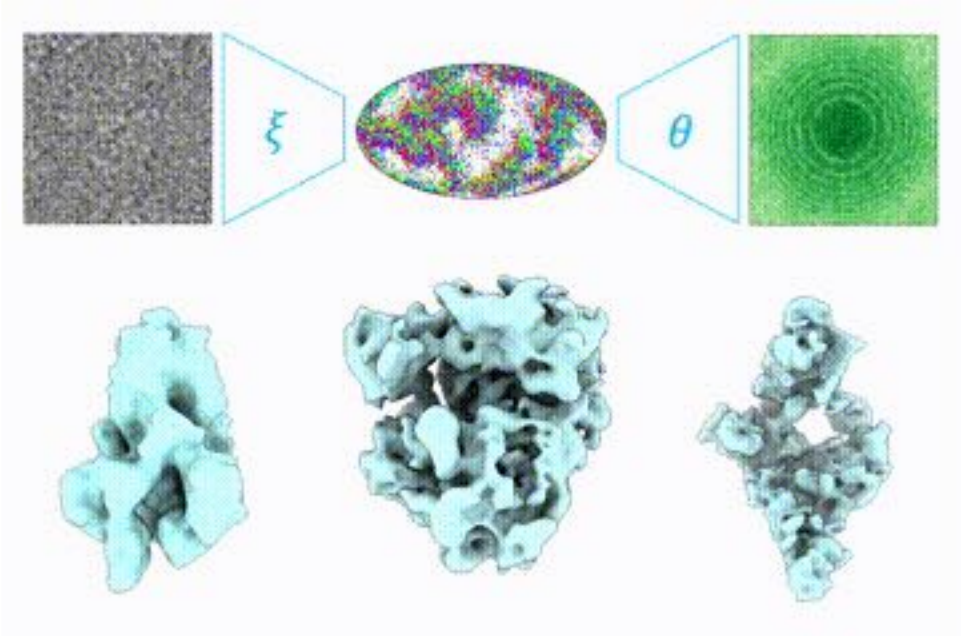

Computational Cameras

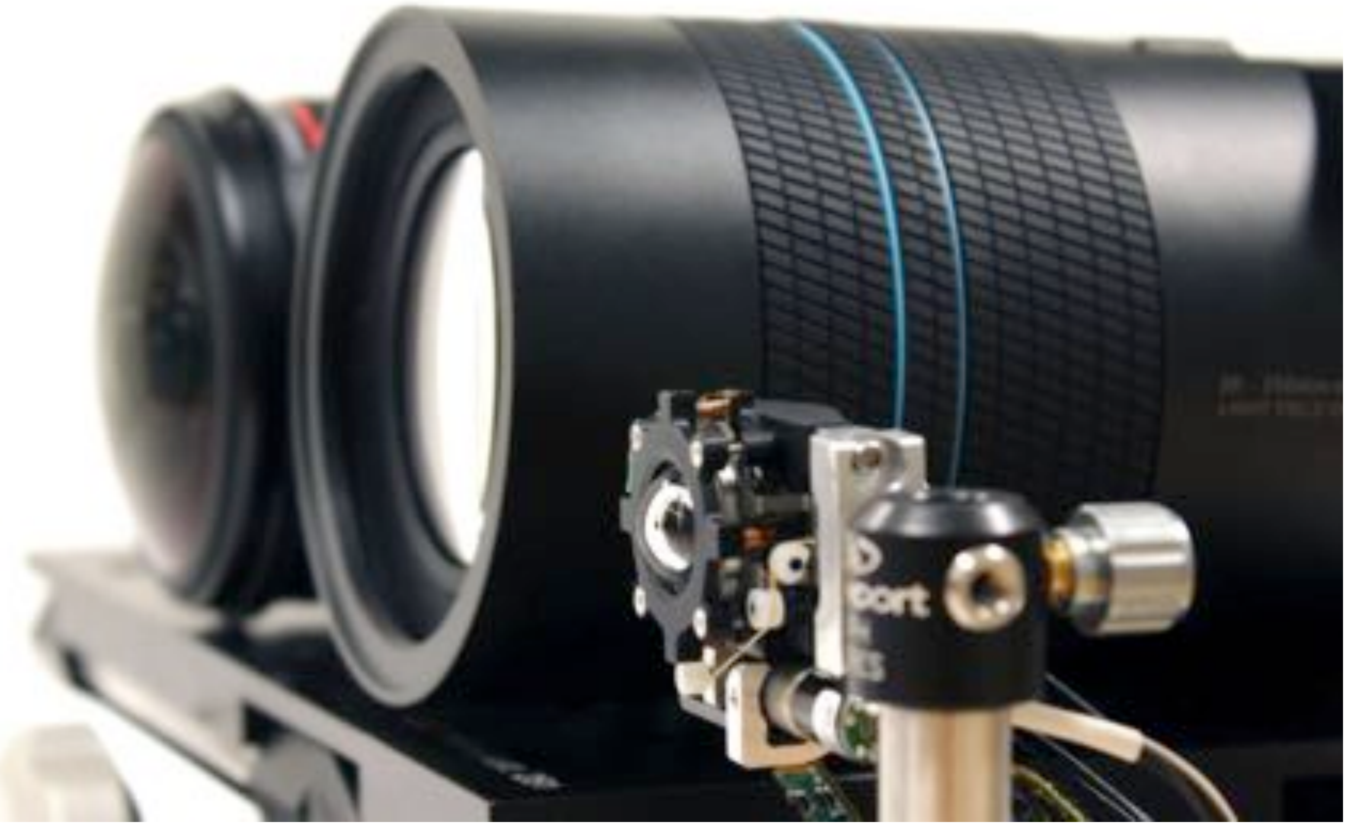

Efficient 3D GANs – Latent Code Interpolation

Chan et al. CVPR 2022

Doug James


Cymbal
SURGE

Jiajun Wu

Kayvon Fatahalian (me)


ROLEX
0.06
1
1
oppo
106 mph
THE CHAMPIONSHIPS WIMBLEDON


ROLEX
0.06
1
1
OPPO
106 mph
THE CHAMPIONSHIPS WIMBLEDON
Federer
Djokovic

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 2023

Controlling generative AI by making collages

User creates a collage with standard graphics tools

Collage defines what should be in the image, where it goes, and what it should look like

Generative AI turns the collage into a "plausible" realistic image

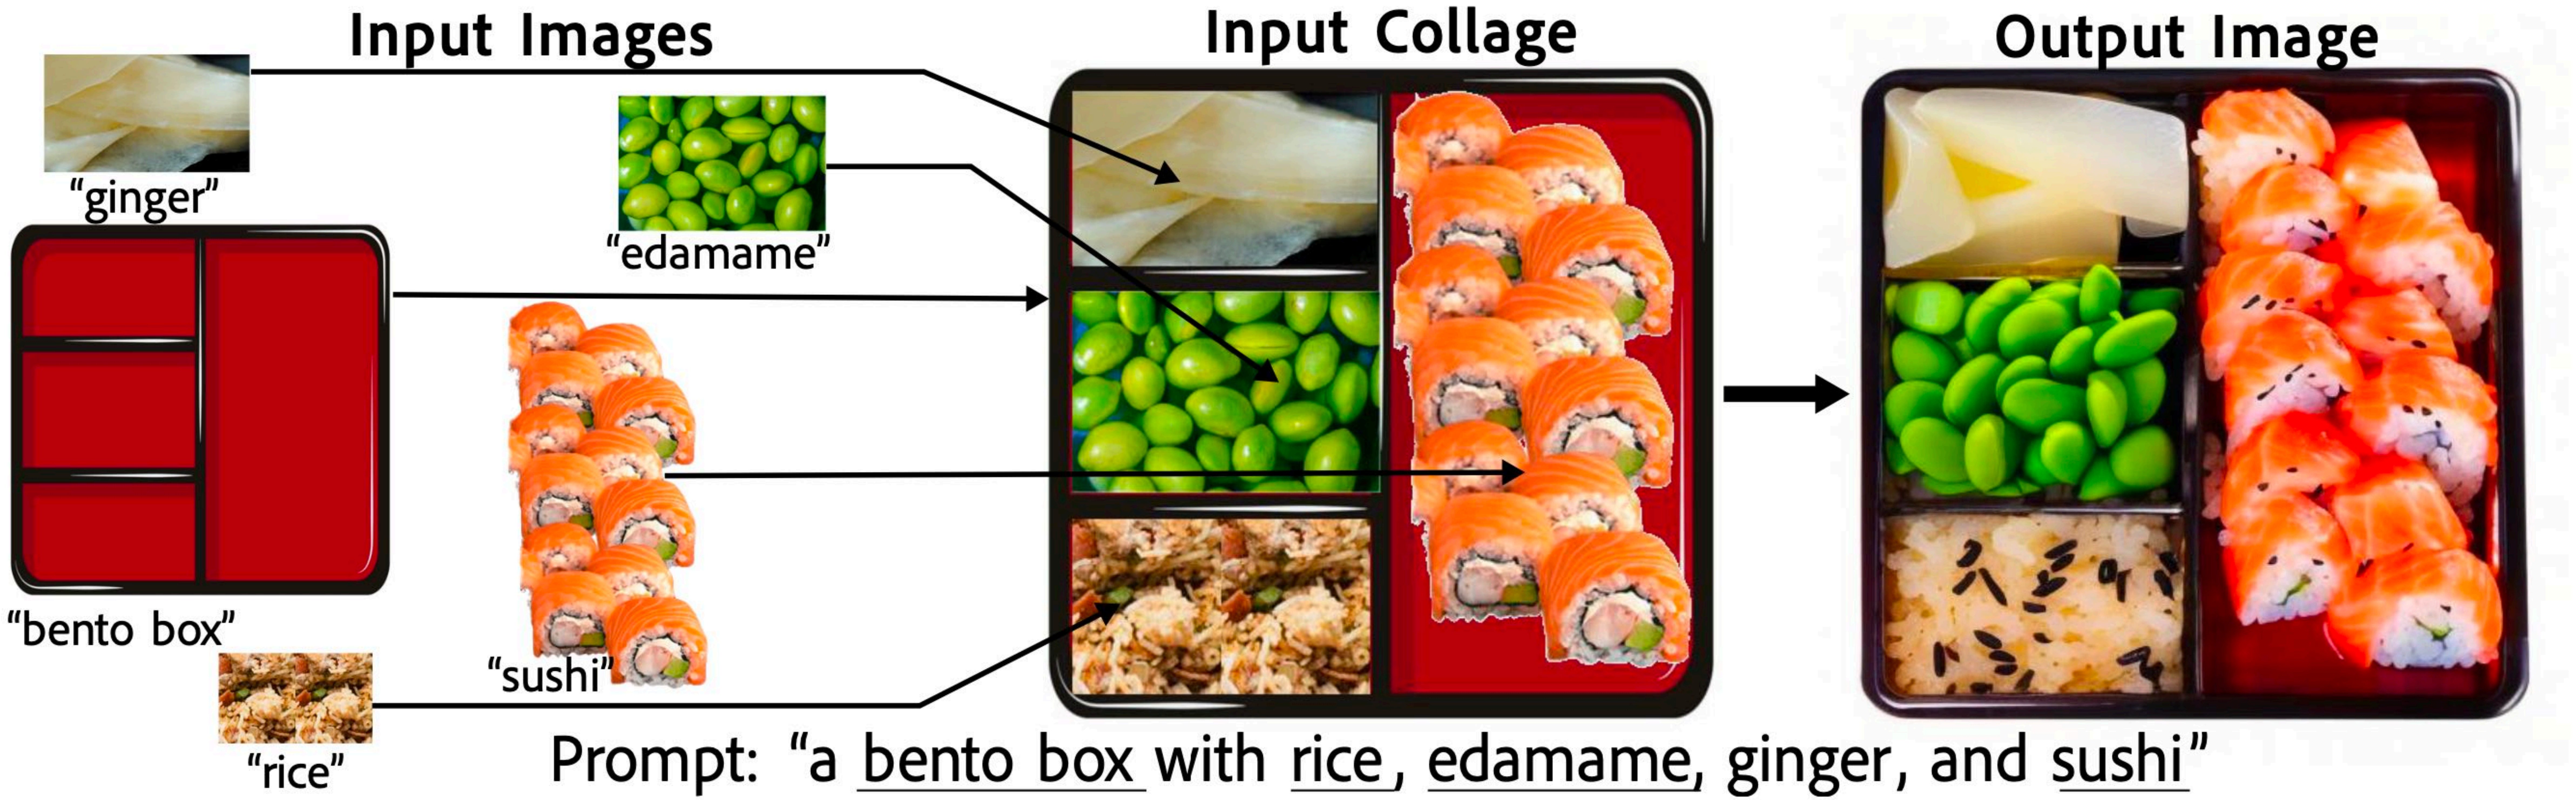

High Interest in AI Agents Across Disciplines

Game-Playing Agents

Dota 2

Google Research Football

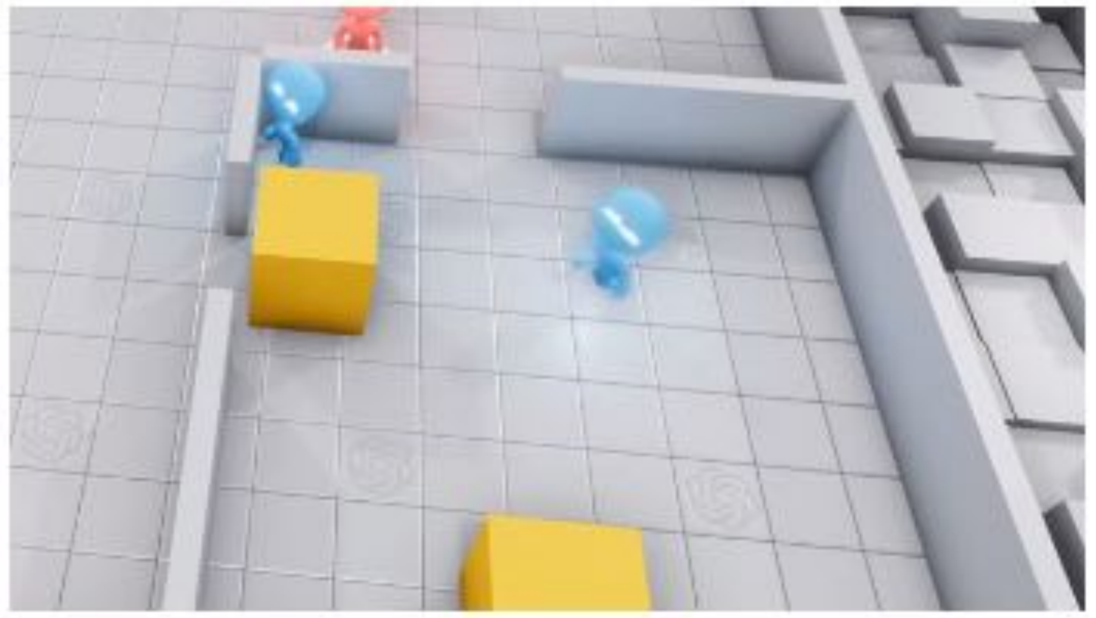

OpenAI Hide & Seek

Robotics

Habitat Rearrangement Challenge

Autonomous Vehicles (Carla, NV Drive Sim)

ProcTHOR

Game Development & Debugging

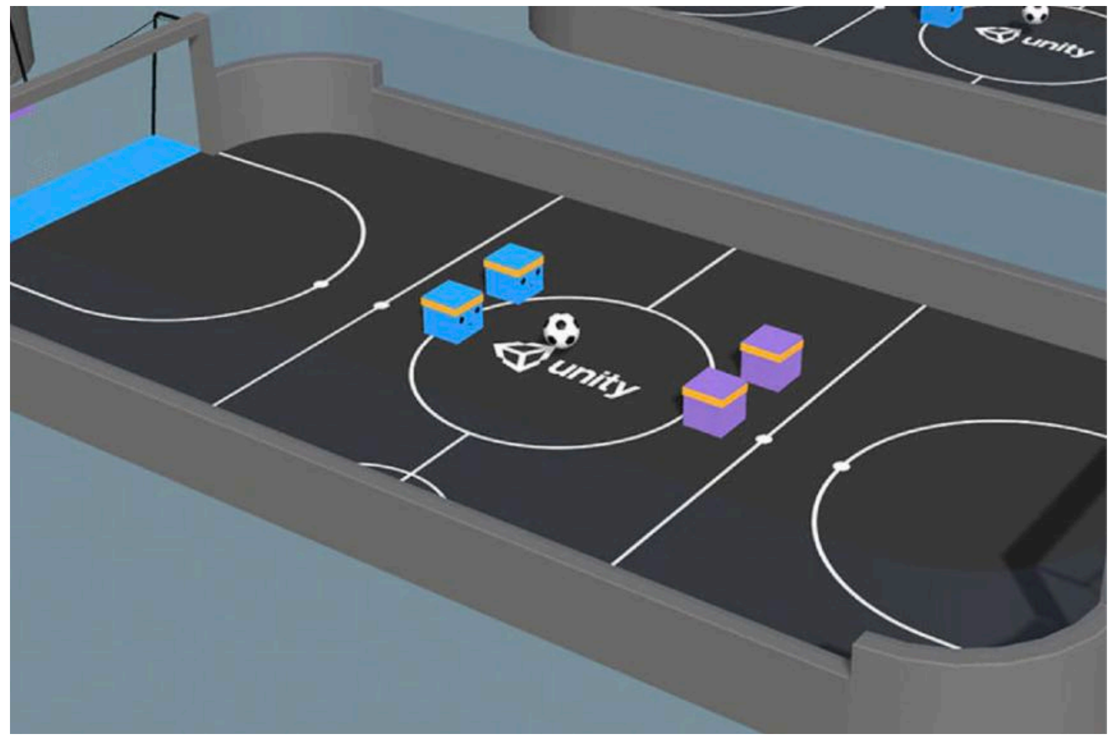

Unity ML-Agents

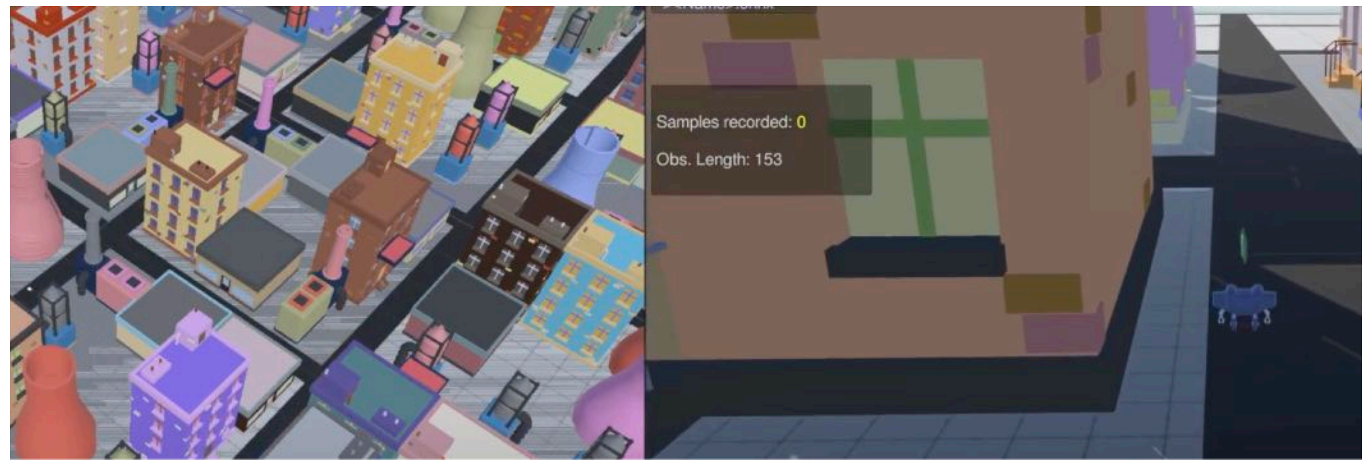

Automated QA & Design (EA SEED)

Generating Simulated Experience is Computationally Demanding: Slow & Expensive Training!

OpenAI Five

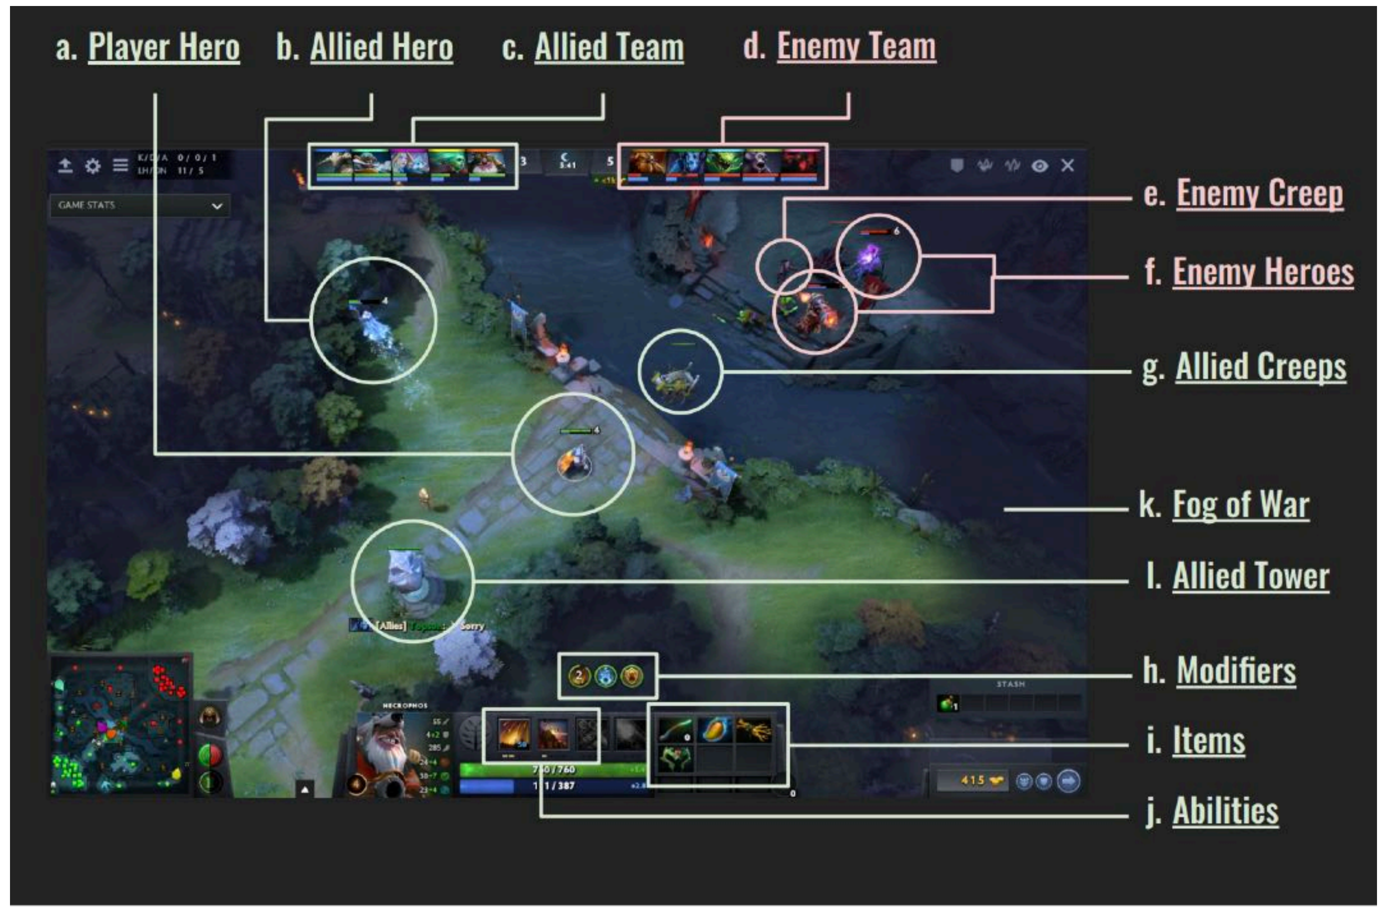

Rapid: 100k+ CPUs,
months of training

Habitat 2.0

25000 FPS on 8 GPUs,
Months to Learn Rearrangement

Unity ML-Agents

Thoughput Limited by
Scaling Strategy

Example: OpenAI Hide & Seek:

What if we redesigned a game engine for running billions of independent training runs?
How fast could you go?

Simulating 16,000 worlds at 700,000 fps

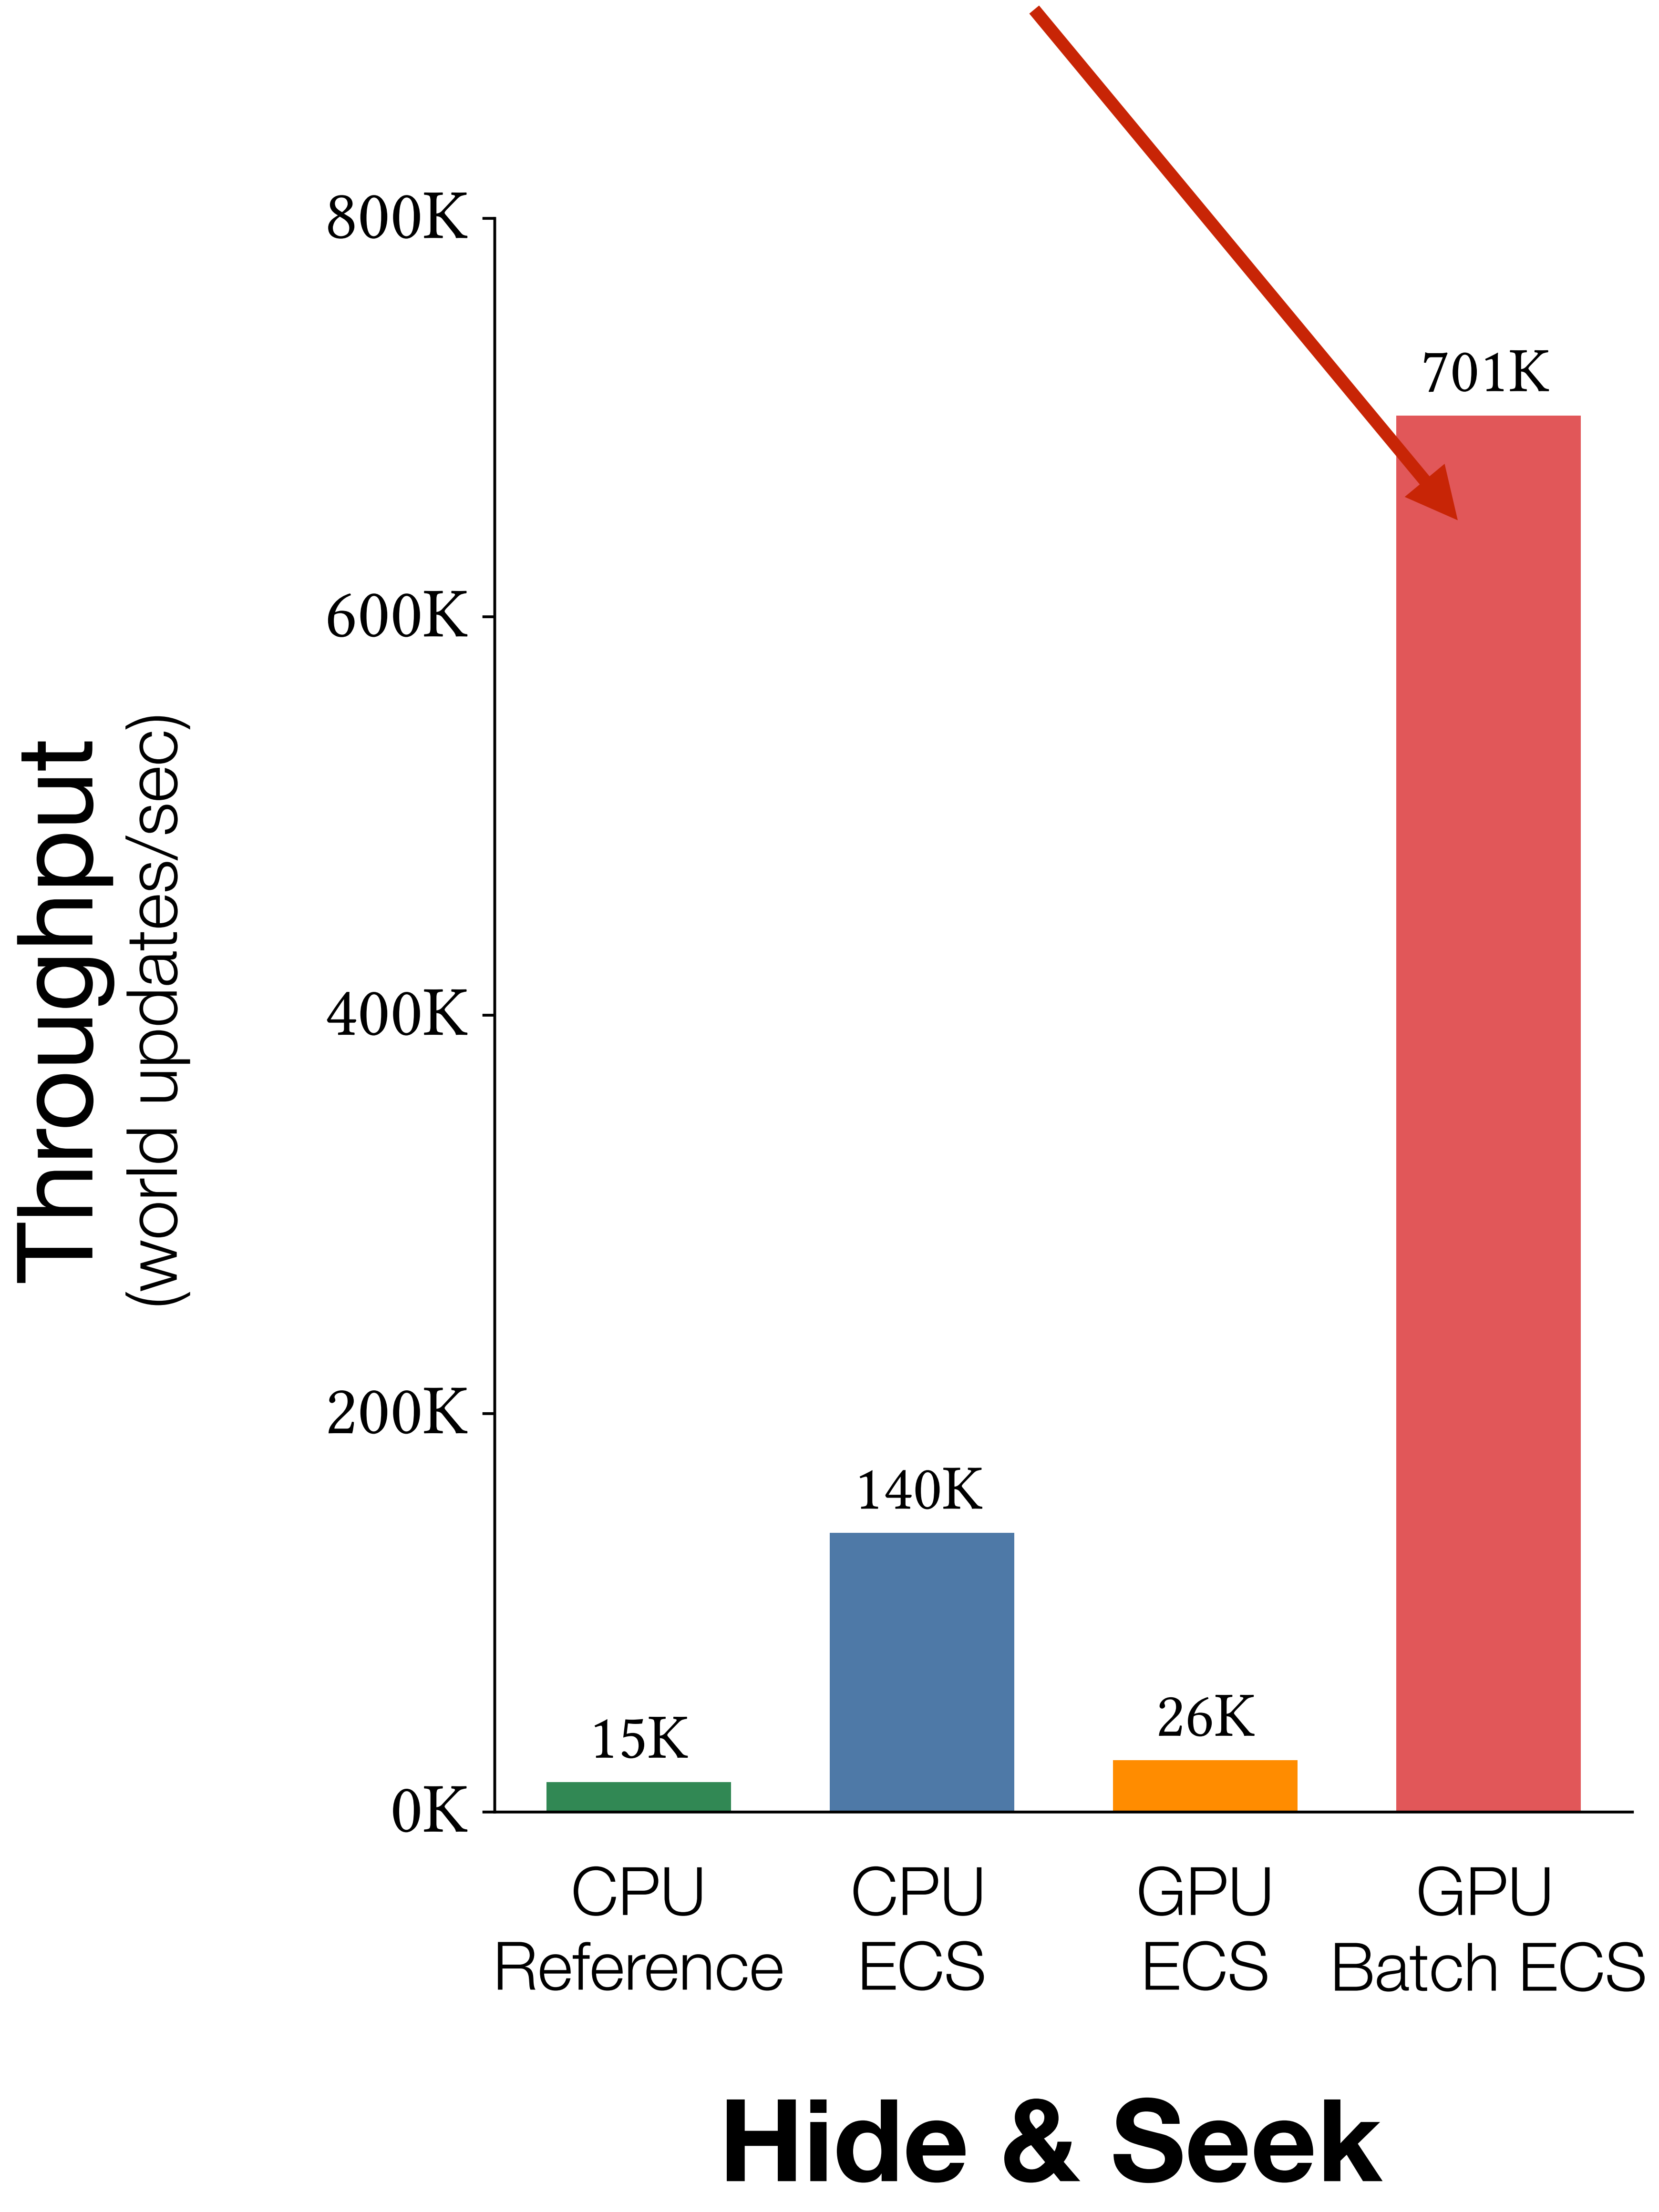

7 Will the slides be available?
+1 Asked by Jerrick Hoang

1 how can I remove the box that says: "some audio/video is being blocked....." It's blocking the title of the slides, thanks
+1 Asked by Zen

0 Where is the =?
+1 Asked by Daniel Wexler

0 What are the threads?
+1 Asked by Sunny Manchanda

Ask a question... Add

Ask anonymously

Radoslaw Jurga
the multiplications are independent of each other

Prateek G
pipelining

praveen dhanuka
yes

Octave Crespel
yeah, x y and z are independent

Kada Situ
pipeline

Philipp Munkes
Superscalar

Type message to auditorium...

What if up to two instructions can be performed at once?

```
a = x*x + y*y + z*z
```

Assume register
R0 = x, R1 = y, R2 = z

```
mul R0, R0, R0
mul R1, R1, R1
mul R2, R2, R2
add R0, R0, R1
add R3, R0, R2
```

R3 now stores value of program variable 'a'

time	Volunteer 1	Volunteer 2
1		
2		
3		
4		
5		

Step to mic


Colin
Emmanuel
NYC
Frank
Ann
Devon
Jean
Jason
Clean up
Red and black antennae waving
They all do it the saaaame
Phil
live from Brooklyn
Nate+Pam
Anne
Seth
Back to Karaoke Shower

Other popular research topics in computer graphics…

Creating physically plausible models

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

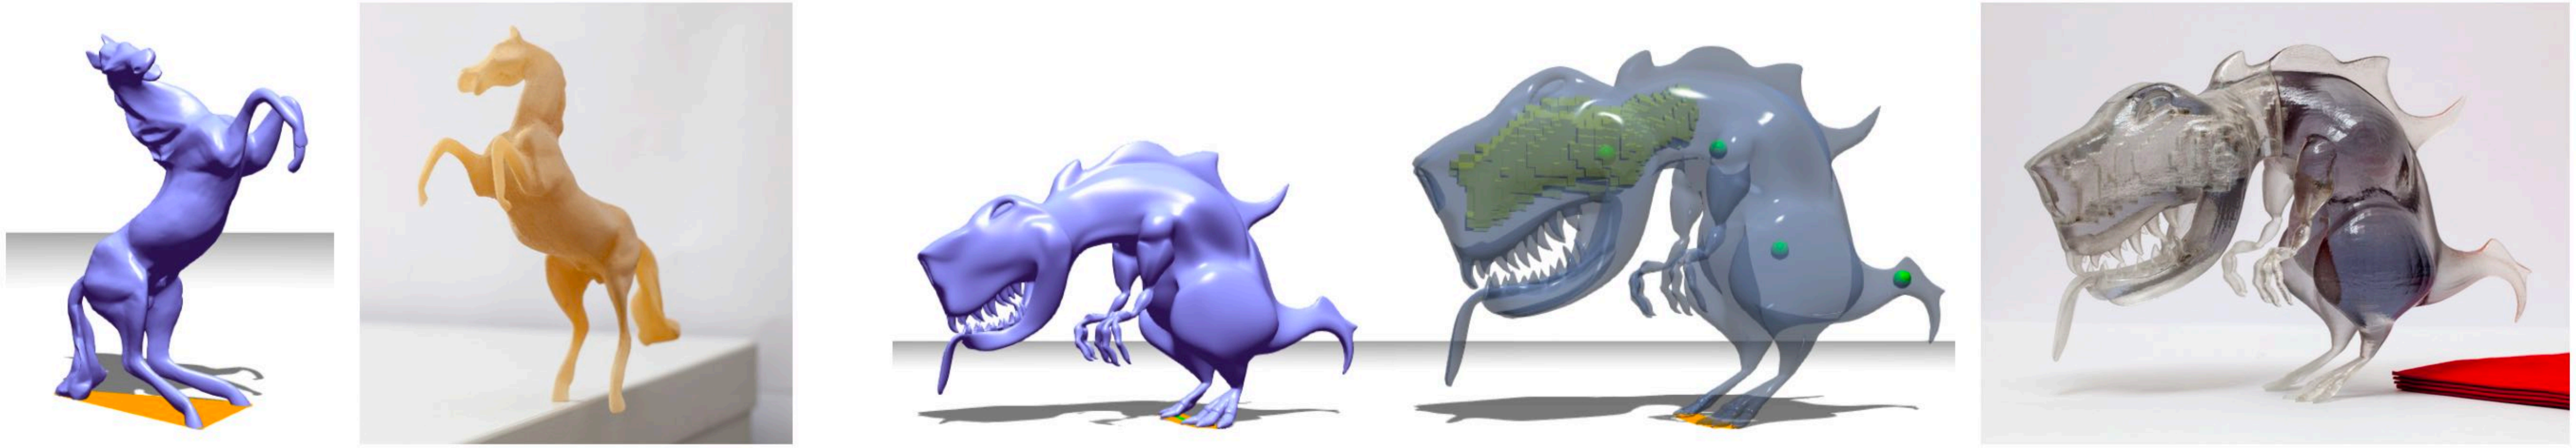

Fabricate models that are balanced to stand

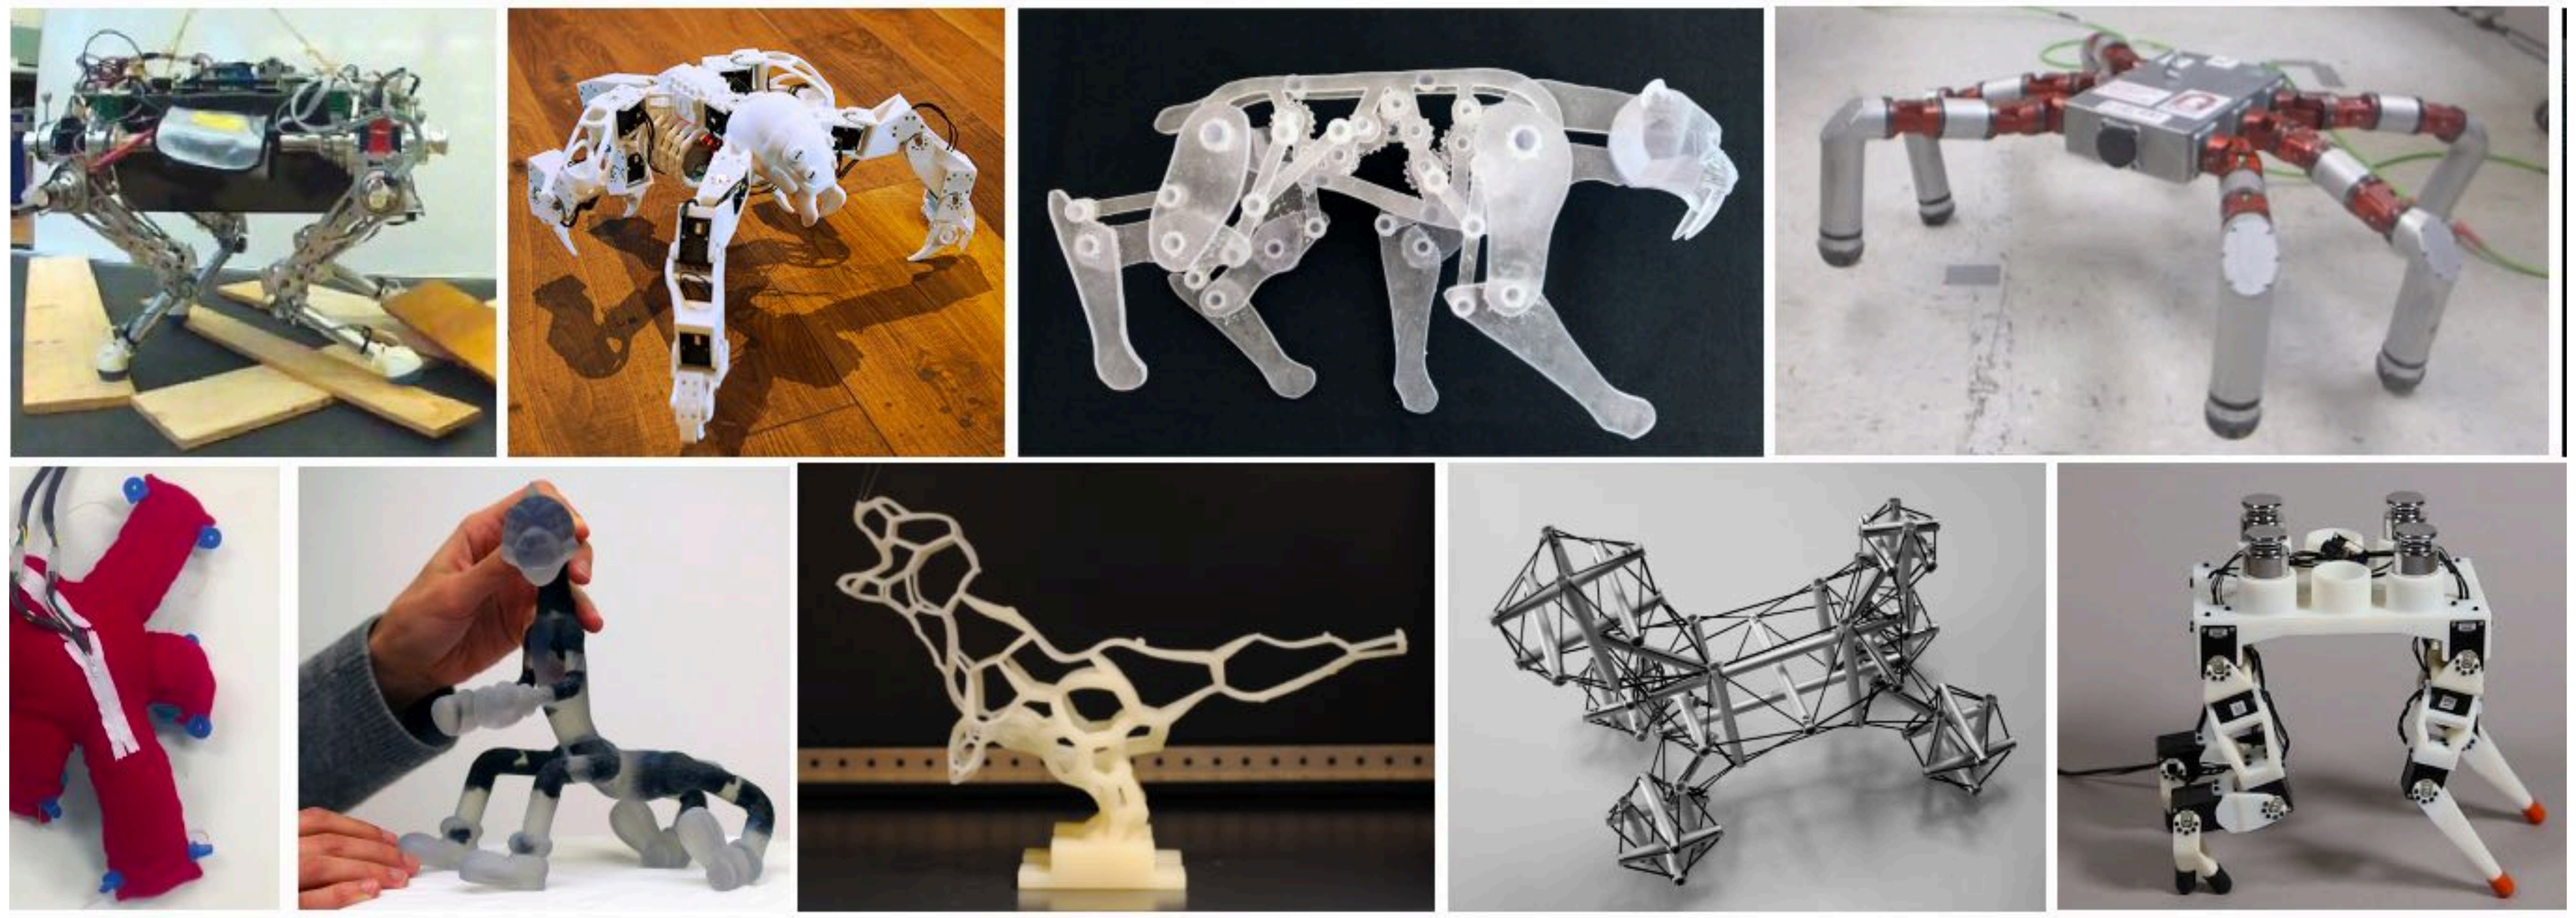

Fabricate robots that can balance and move

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 displays/rendering for VR/AR

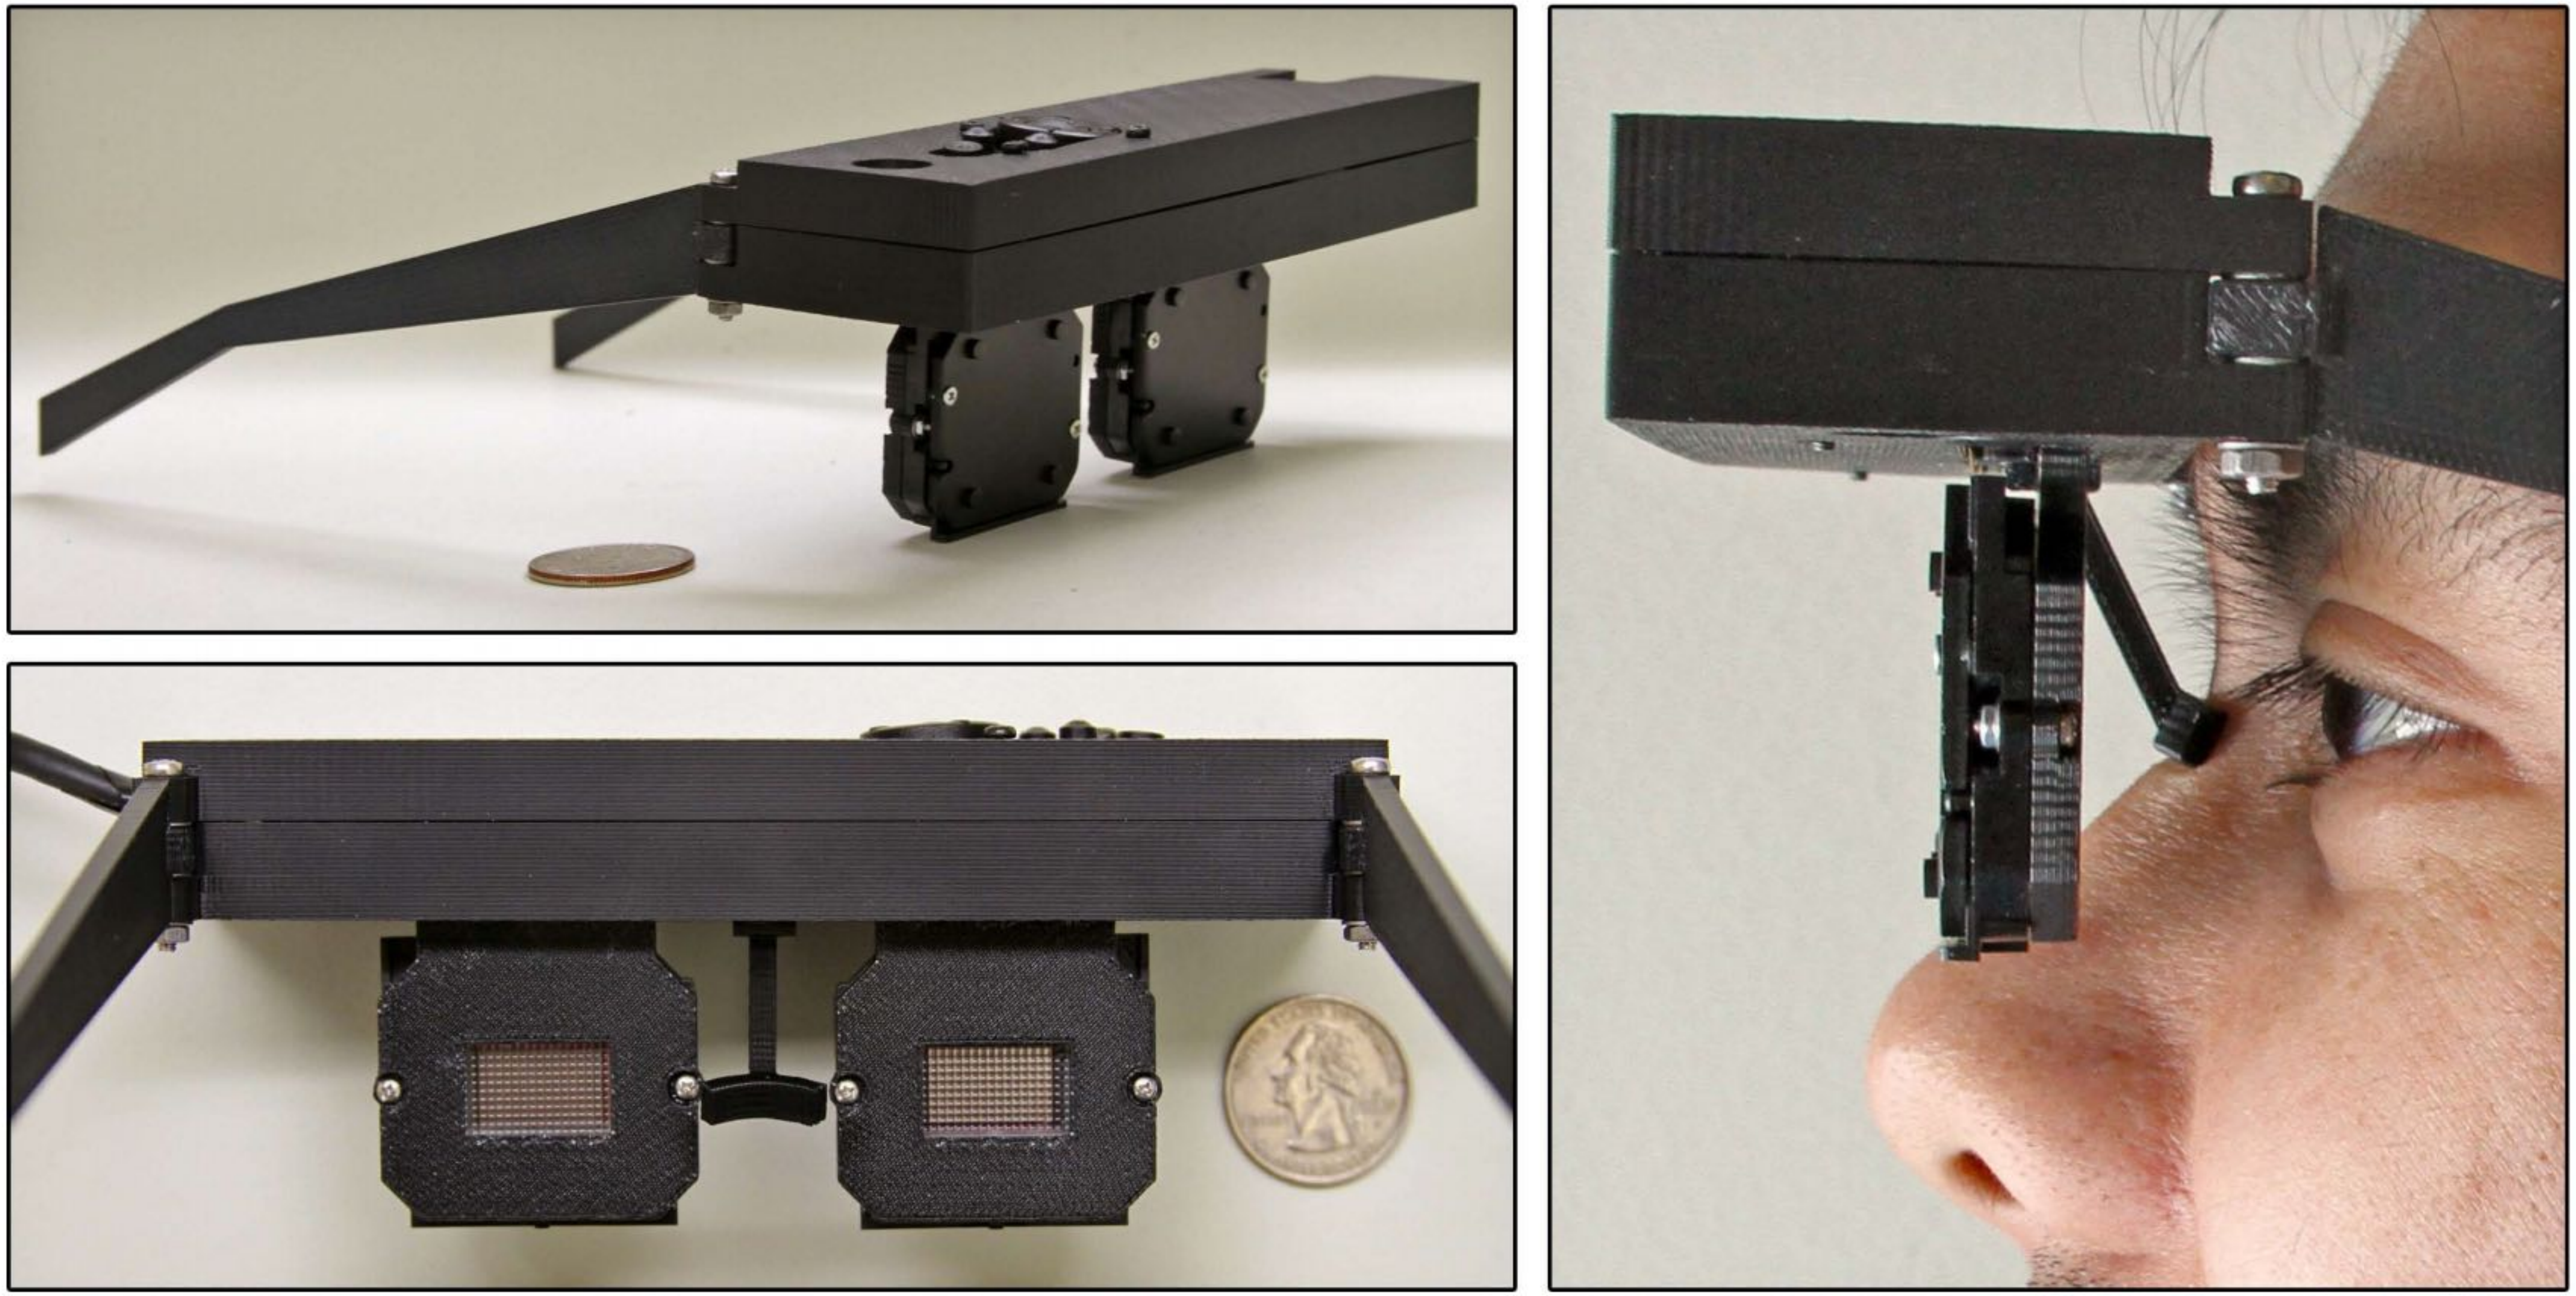

Near eye light field display

A fun resource

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

http://kesen.realtimerendering.com/

SIGGRAPH 2022 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 41, Issue 4 (July 2022) Proceedings of ACM SIGGRAPH 2022

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

Conditionally Accepted Papers / TOG Papers

Ecoclimates: Climate-Response Modeling of Vegetation
Wojtek Palubicki, Milosz Makowski, Weronika Gajda (Adam Mickiewicz University), Torsten Hadrich, Dominik L. Michels (KAUST), Soren Pirk (Adobe Research)

DSG-Net: Learning Disentangled Structure and Geometry for 3D Shape Generation
Jie Yang* (Institute of Computing Technology, Chinese Academy of Sciences and University of Chinese Academy of Sciences), Kaichun Mo* (Stanford University), Yu-Kun Lai (Cardiff University), Leonidas J. Guibas (Stanford University), Lin Gao (Institute of Computing Technology, Chinese Academy of Sciences and University of Chinese Academy of Sciences) * Authors contributed equally

A Fast Unsmoothed Aggregation Algebraic Multigrid Framework for the Large-Scale Simulation of Incompressible Flow
Han Shao*, Libo Huang*, Dominik L. Michels (KAUST) * Authors contributed equally

A Practical Model for Realistic Butterfly Flight Simulation
Qiang Chen (Jiangxi University of Finance and Economics), Tingsong Lu, Yang Tong, Guoliang Luo (East China Jiaotong University), Xiaogang Jin (Zhejiang University), Zhigang Deng (University of Houston)

Dev2PQ: Planar Quadrilateral Strip Remeshing of Developable Surfaces
Floor Verhoeven (ETH Zurich), Amir Vaxman (Utrecht University), Tim Hoffmann (TU Munich), Olga Sorkine-Hornung (ETH Zurich)

DiffCloth: Differentiable Cloth Simulation with Dry Frictional Contact
Yifei Li, Tao Du (MIT CSAIL), Kui Wu (Tencent Lightspeed & Quantum Studios), Jie Xu, Wojciech Matusik (MIT CSAIL)

DiffusionNet: Discretization Agnostic Learning on Surfaces
Nicholas Sharp (Carnegie Mellon University and University of Toronto), Souhaib Attaiki (LIX, Ecole Polytechnique), Keenan Crane (Carnegie Mellon University), Maks Ovsjanikov (LIX, Ecole Polytechnique)

HRBF-Fusion: Accurate 3D reconstruction from RGB-D data using on-the-fly implicits
Yabin Xu (Nanjing University of Aeronautics and Astronautics and Delft University of Technology), Liangliang Nan (Delft University of Technology), Laishui Zhou, Jun Wang (Nanjing University of Aeronautics and Astronautics), Charlie C.L. Wang (The University of Manchester and Delft University of Technology)

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 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.

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!