Lecture 20:

Course Summary + Graphics at Stanford Today

Computer Graphics: Rendering, Geometry, and Image Manipulation
Stanford CS248A, Winter 2025

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?

Visual Computing classes (coming quarters) at Stanford

SPRING

CS348K: "Visual Computing Systems", creating efficient systems for photography, 3D graphics, and generative AI (Fatahalian)

CS 231N: "Deep Learning for Computer Vision" (F. Li)

EE267: "Virtual Reality" (Wetzstein)

CS348E: "Character Animation: Modeling, Simulation, and Control of Human Motion" (Liu) (not offered next quarter, but look for it again soon)

FALL

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

CS448B: "Data Visualization" (Agrawala)

CS149: "Parallel Computing" (Fatahalian, Olukotun)

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)

Modern trends/debates in graphics?

Video generated by OpenAI's Sora.

Interactive video generators

Next frame is conditioned on user input… e.g., movement controls

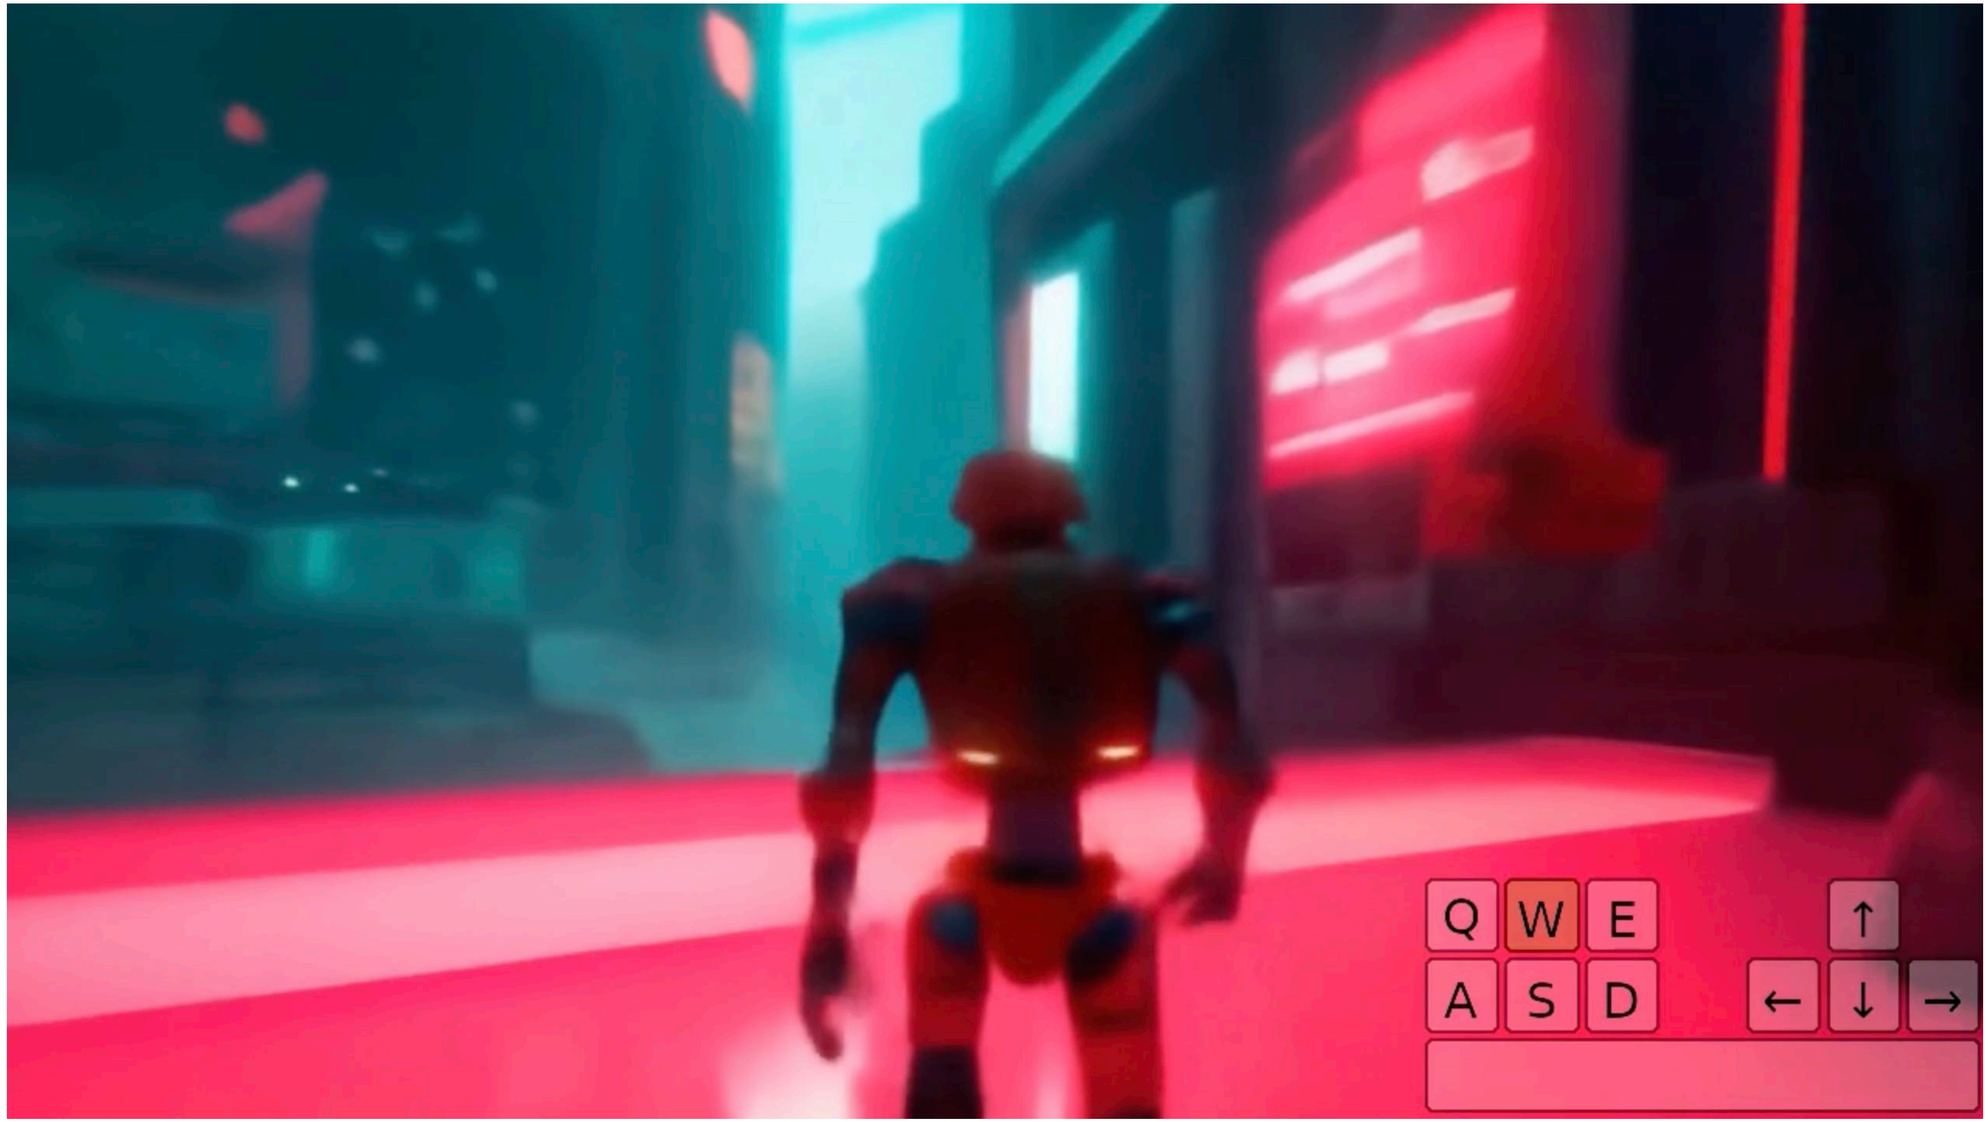

Video generated by DeepMind's Gemini 2.

Frustrating aspects of modern generative AI

- Exciting new capabilities in terms of what output can be created.

- But… text prompts are terrible user interfaces for many tasks
  - See Agrawala's "Unpredictable Black Boxes are Terrible User Interfaces"
  - https://magrawala.substack.com/p/unpredictable-black-boxes-are-terrible

- Photoshop demo…

Graphics Research at Stanford Today

Increasing our ability to control generative AI methods

(Agrawala) ControlNet: more precise ways to control generative AI

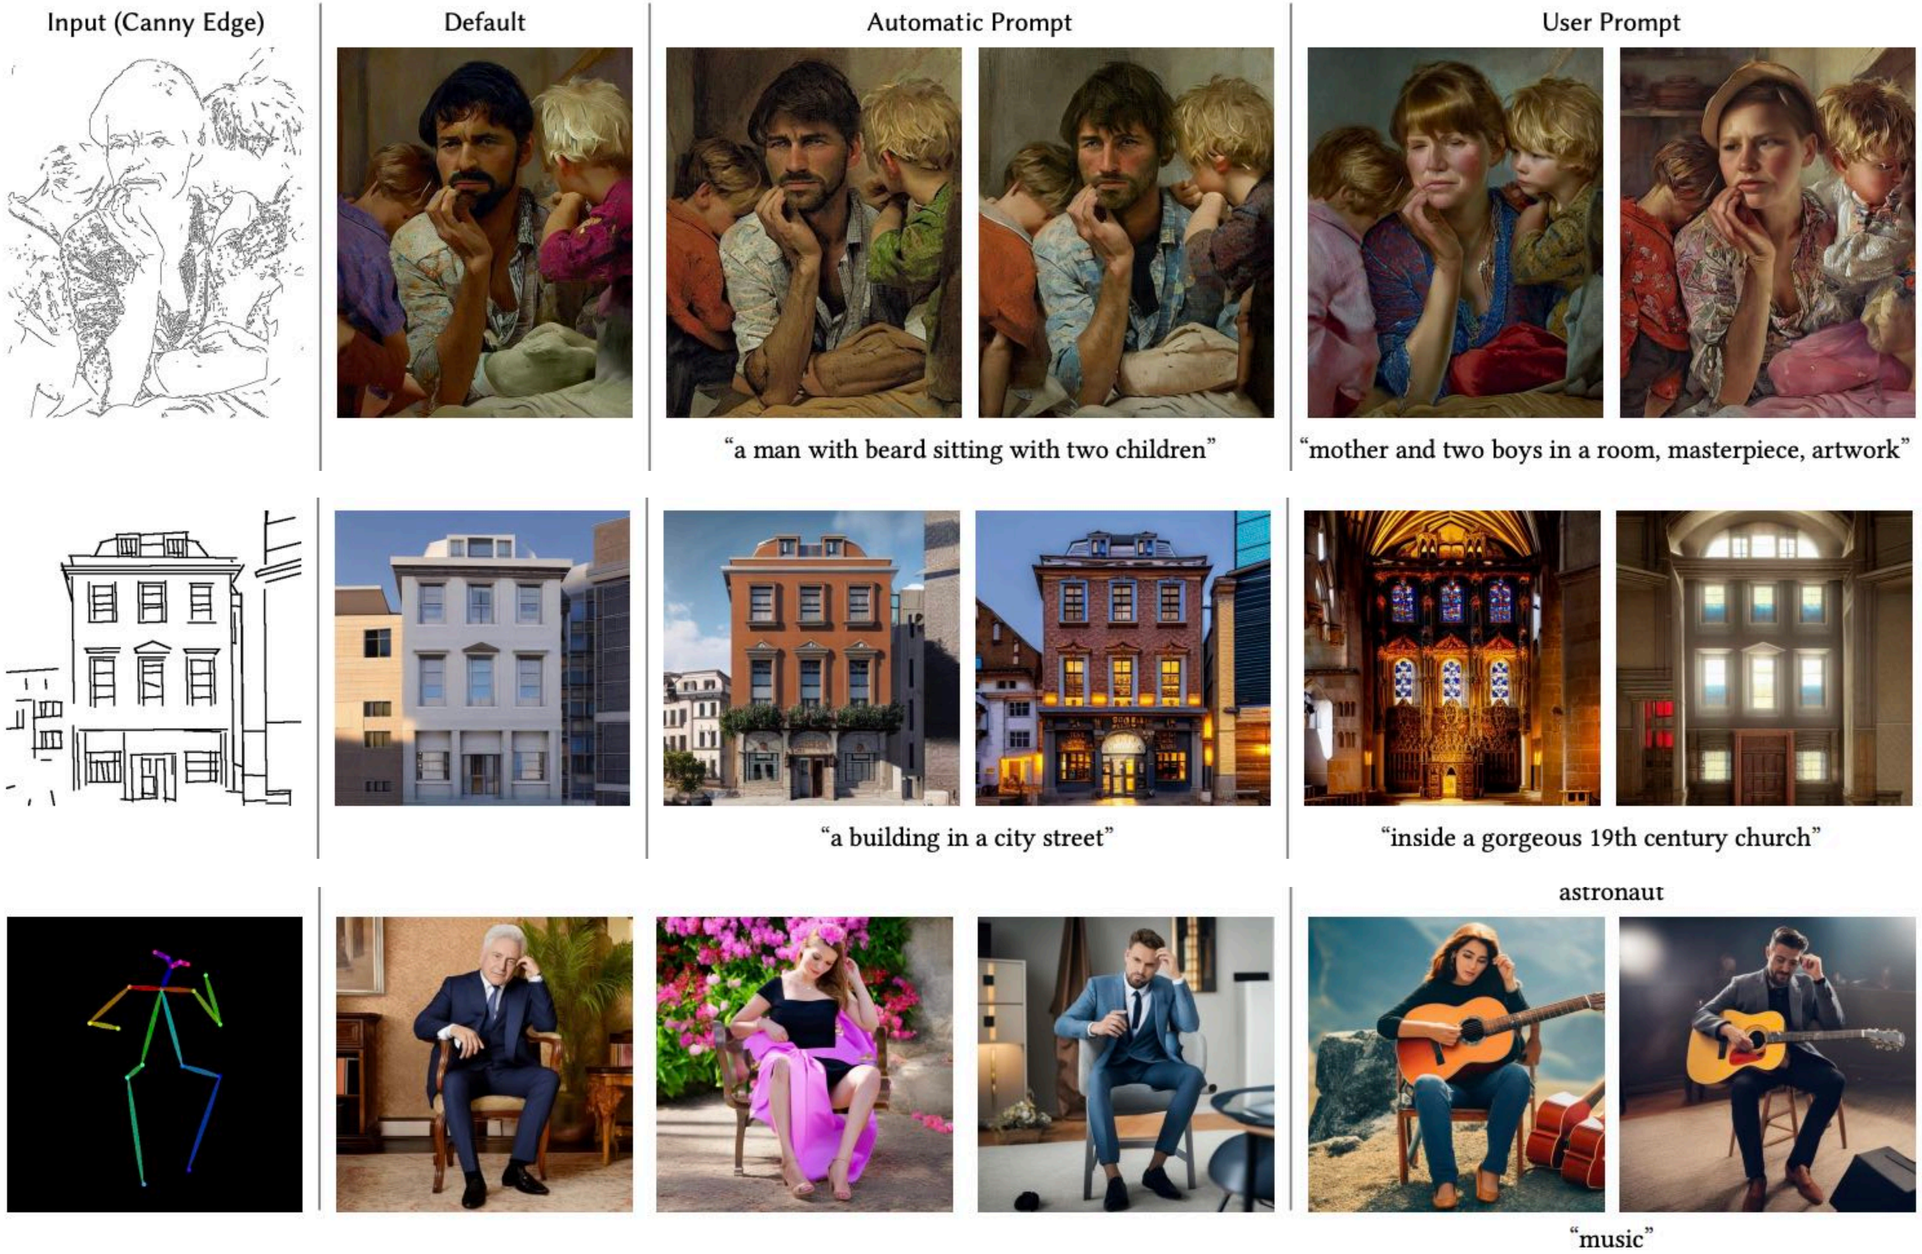

(Agrawala+Fatahalian) Integrating knowledge of the artist workflow into generative AI controls

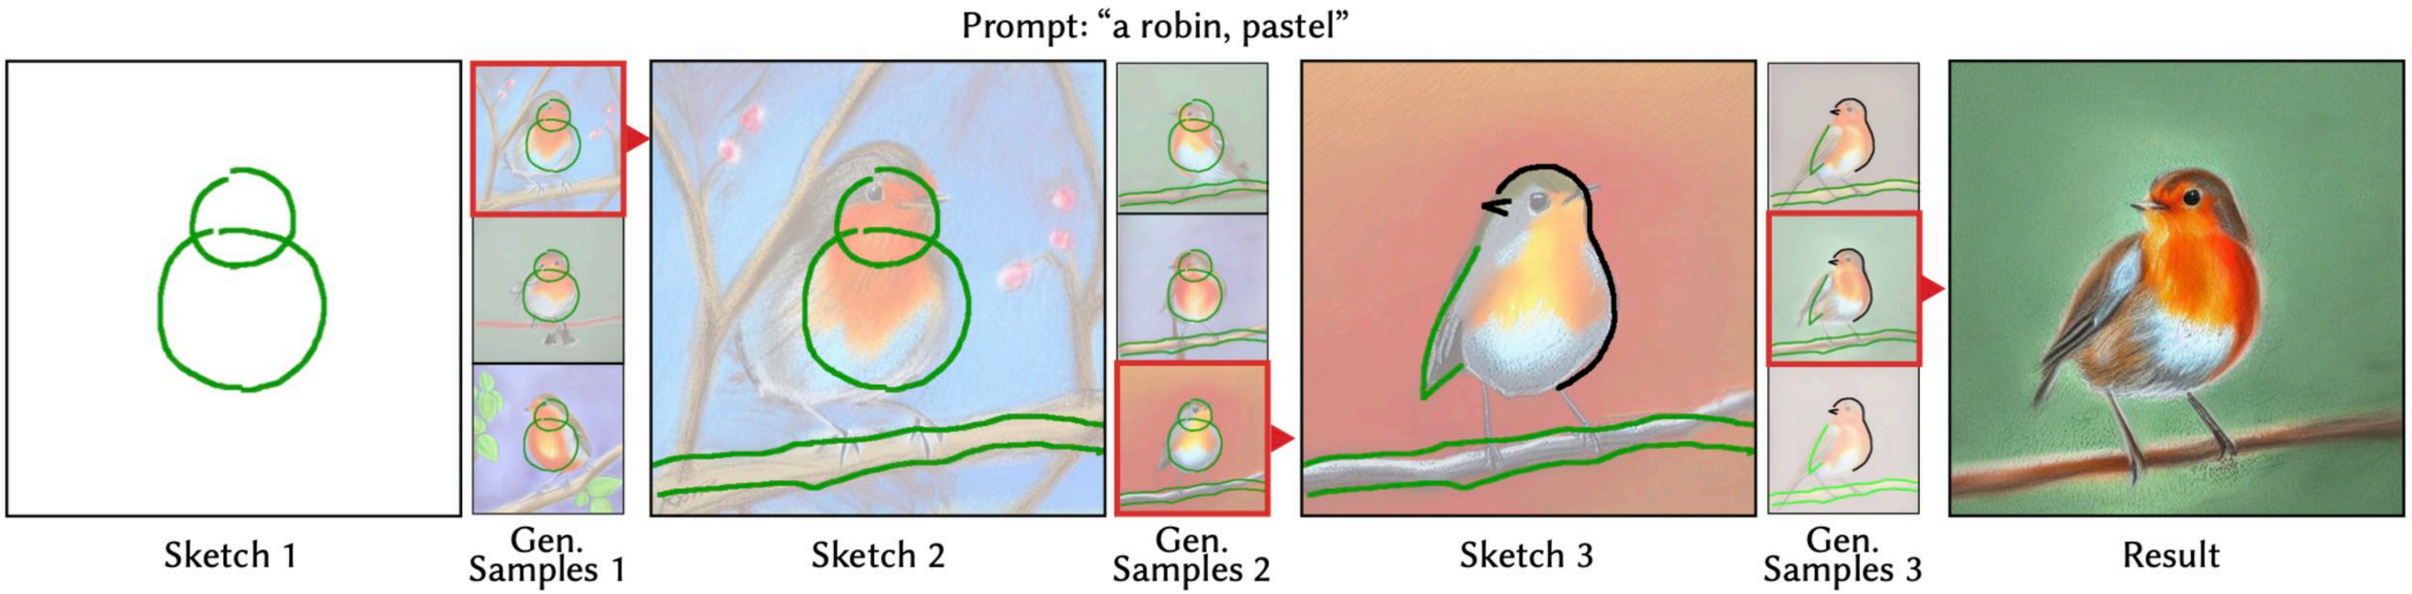

Neuro-symbolic approaches to generative AI

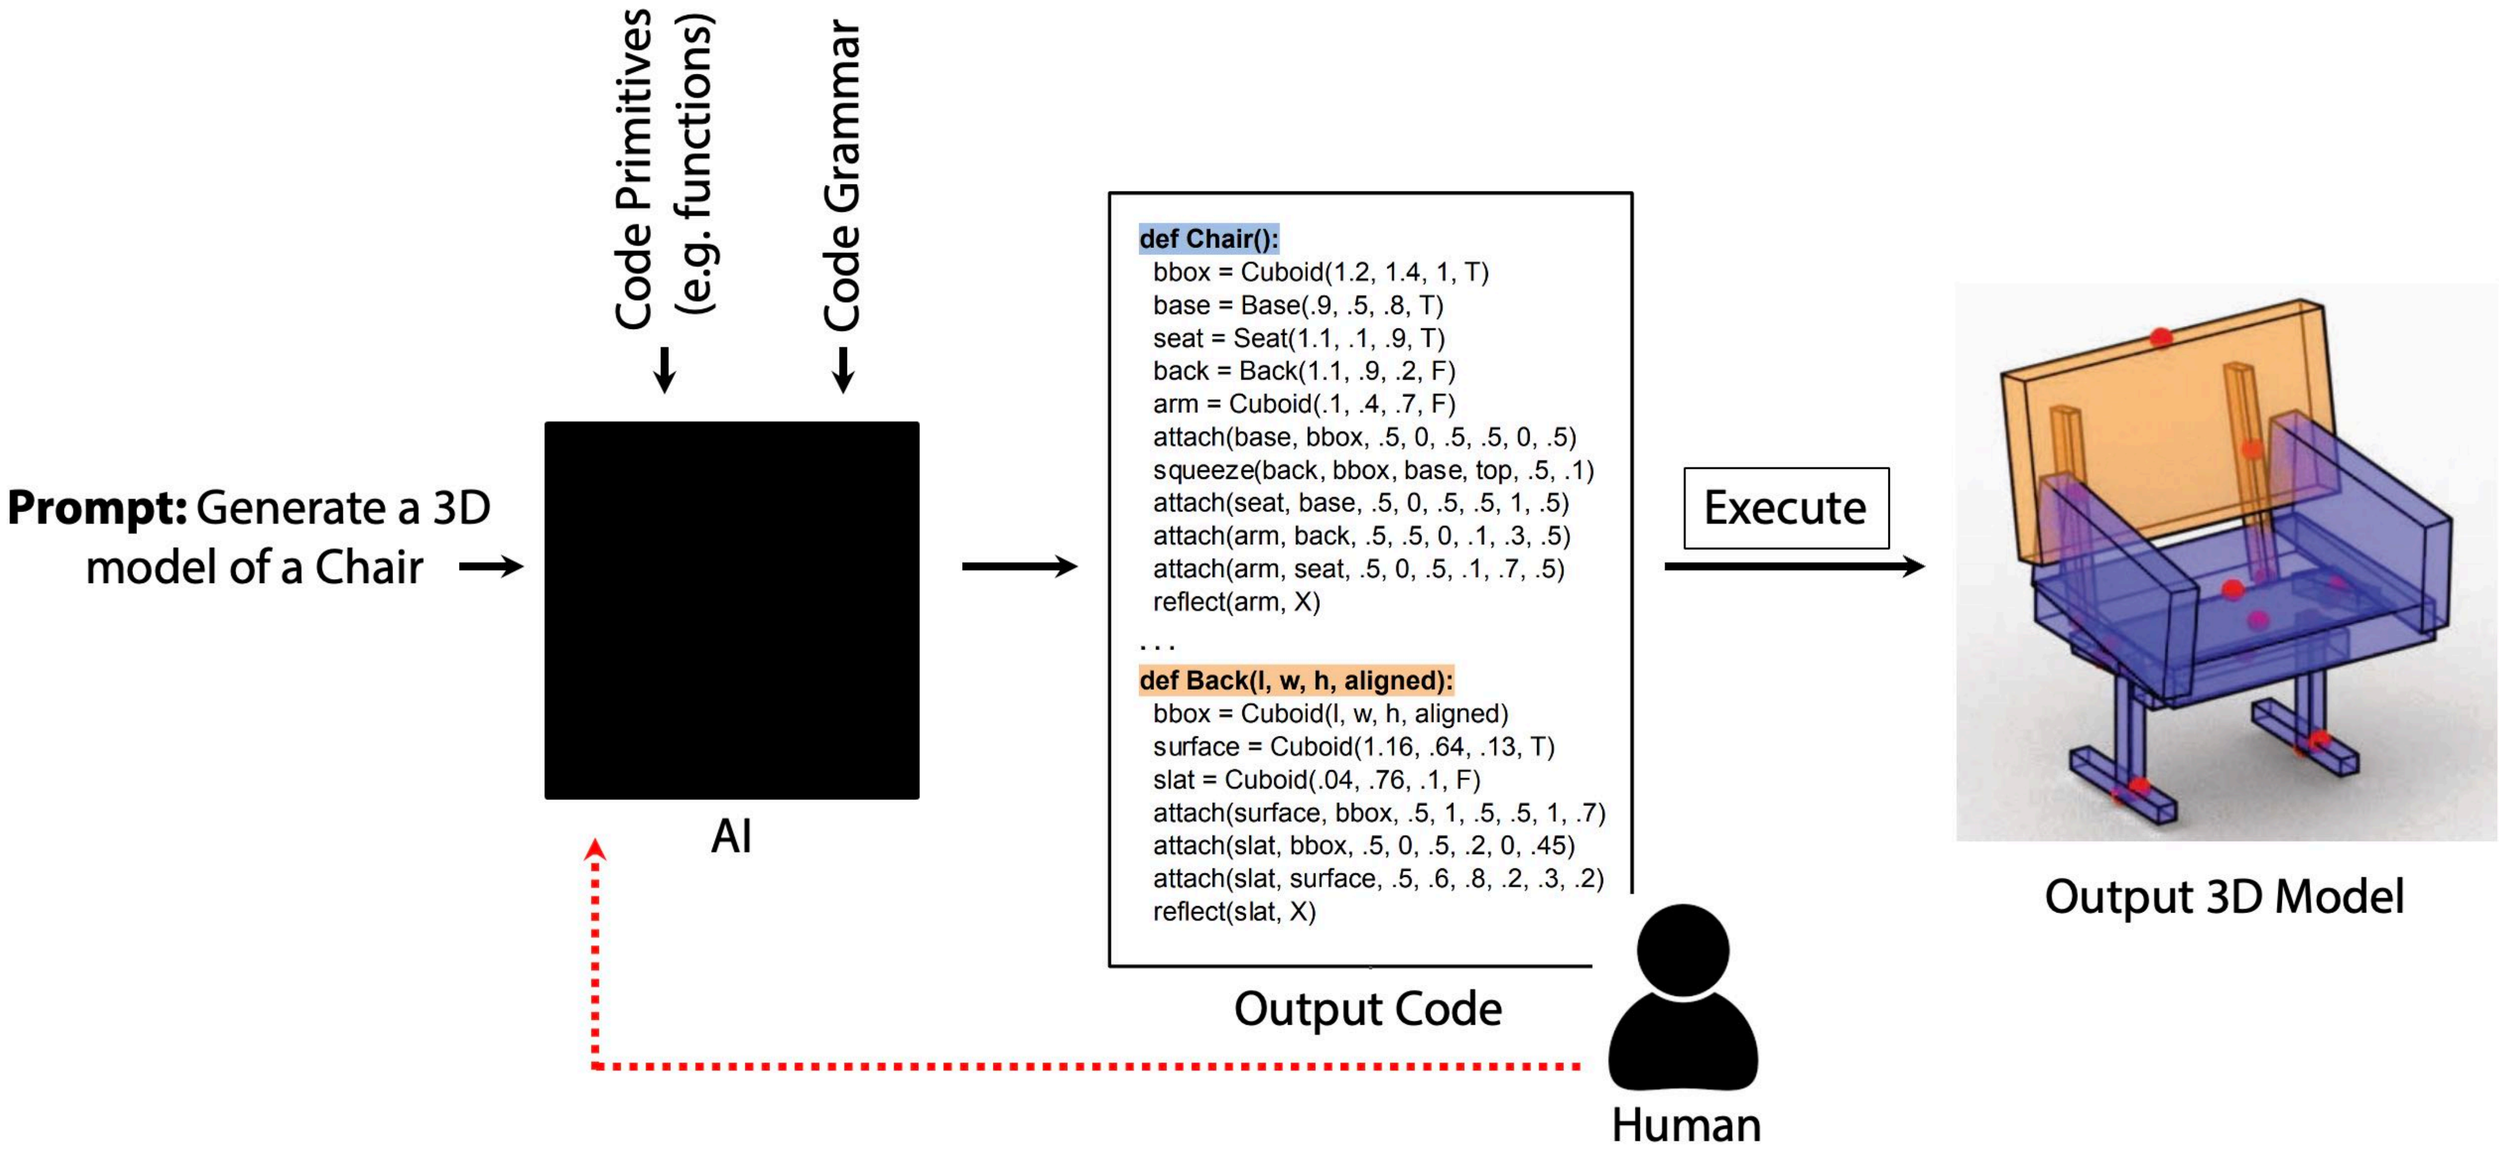

Automated verification of SVG animation programs

Text prompt

<svg> ... </svg>
Move the orange circle above the rectangular shape.
In the meantime, rotate the letter H clockwise by 90 degrees.

MoVer program

```
o1 = lo.clr(o,"orange")∧shp(o,"circle") # check for an orange circle
o2 = lo.shp(o,"rectangle")              # check for a rectangle
o3 = lo.id(o,"H")                       # check for the id "H"
# check for a translation by the circle to the top of the rectangle
m1 = lm.type(m,"trn")∧agt(m,o1)∧post(m,top(o1,o2))
# check for a clockwise rotation of 90 degrees by "H"
m2 = lm.type(m,"rot")∧agt(m,o3)∧dir(m,"cw")∧mag(m,90)
# assert that the translation and the rotation overlap in time
while(m1,m2)
```

MoVer verification report

```
m1 = lm.type(m,"trn")∧agt(m,o1)∧post(m,top(o1,o2))          false
     top(o1,o2)                                             false
     post(m,top(o1,o2))                                     false

m2 = lm.type(m,"rot")∧agt(m,o3)∧dir(m,"cw")∧mag(m,90)       true

while(m1,m2)                                                false
```

Generated motion graphics animation

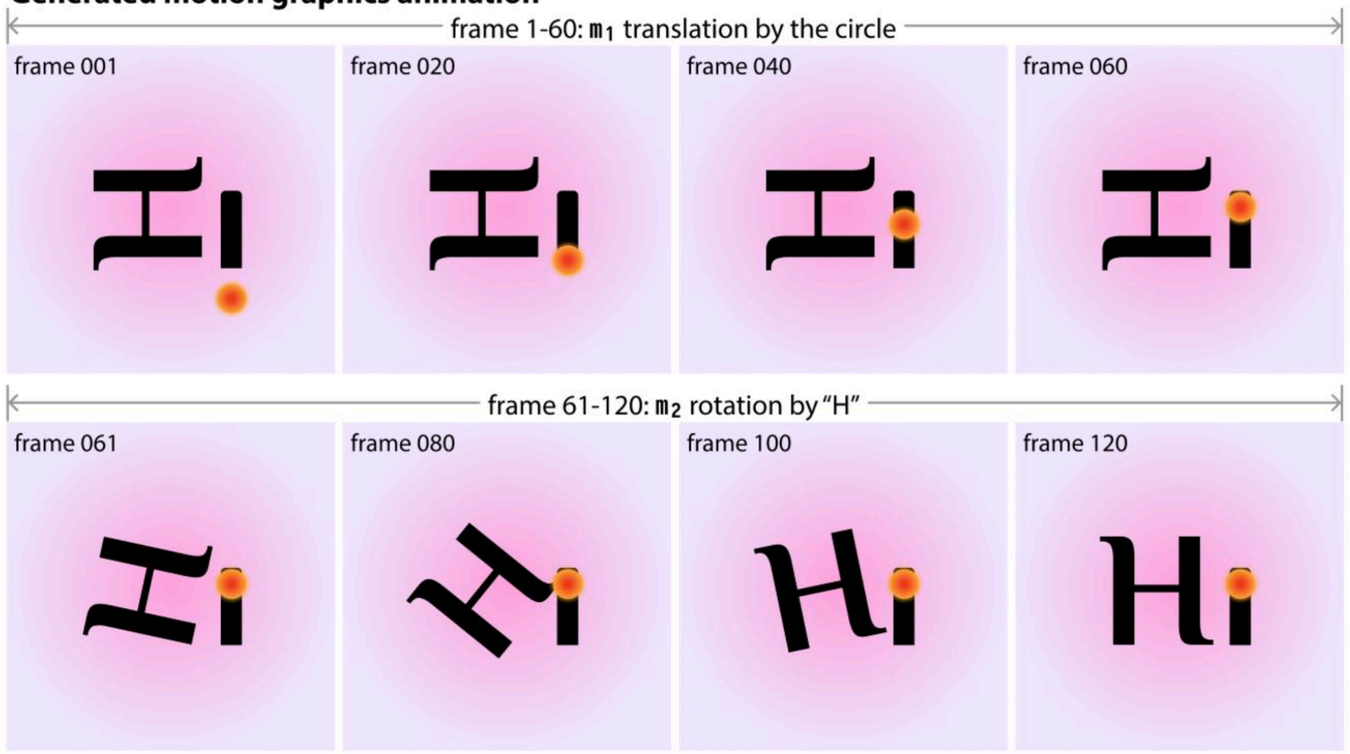

Relighting images

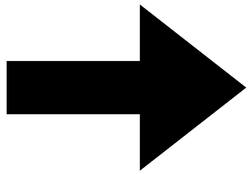

Prompt

man in a disco

Initial Latent

None | Left Light | Right Light | Top Light | Bottom Light

Ron Fedkiw: AI accelerated performance capture and real-time physical simulation

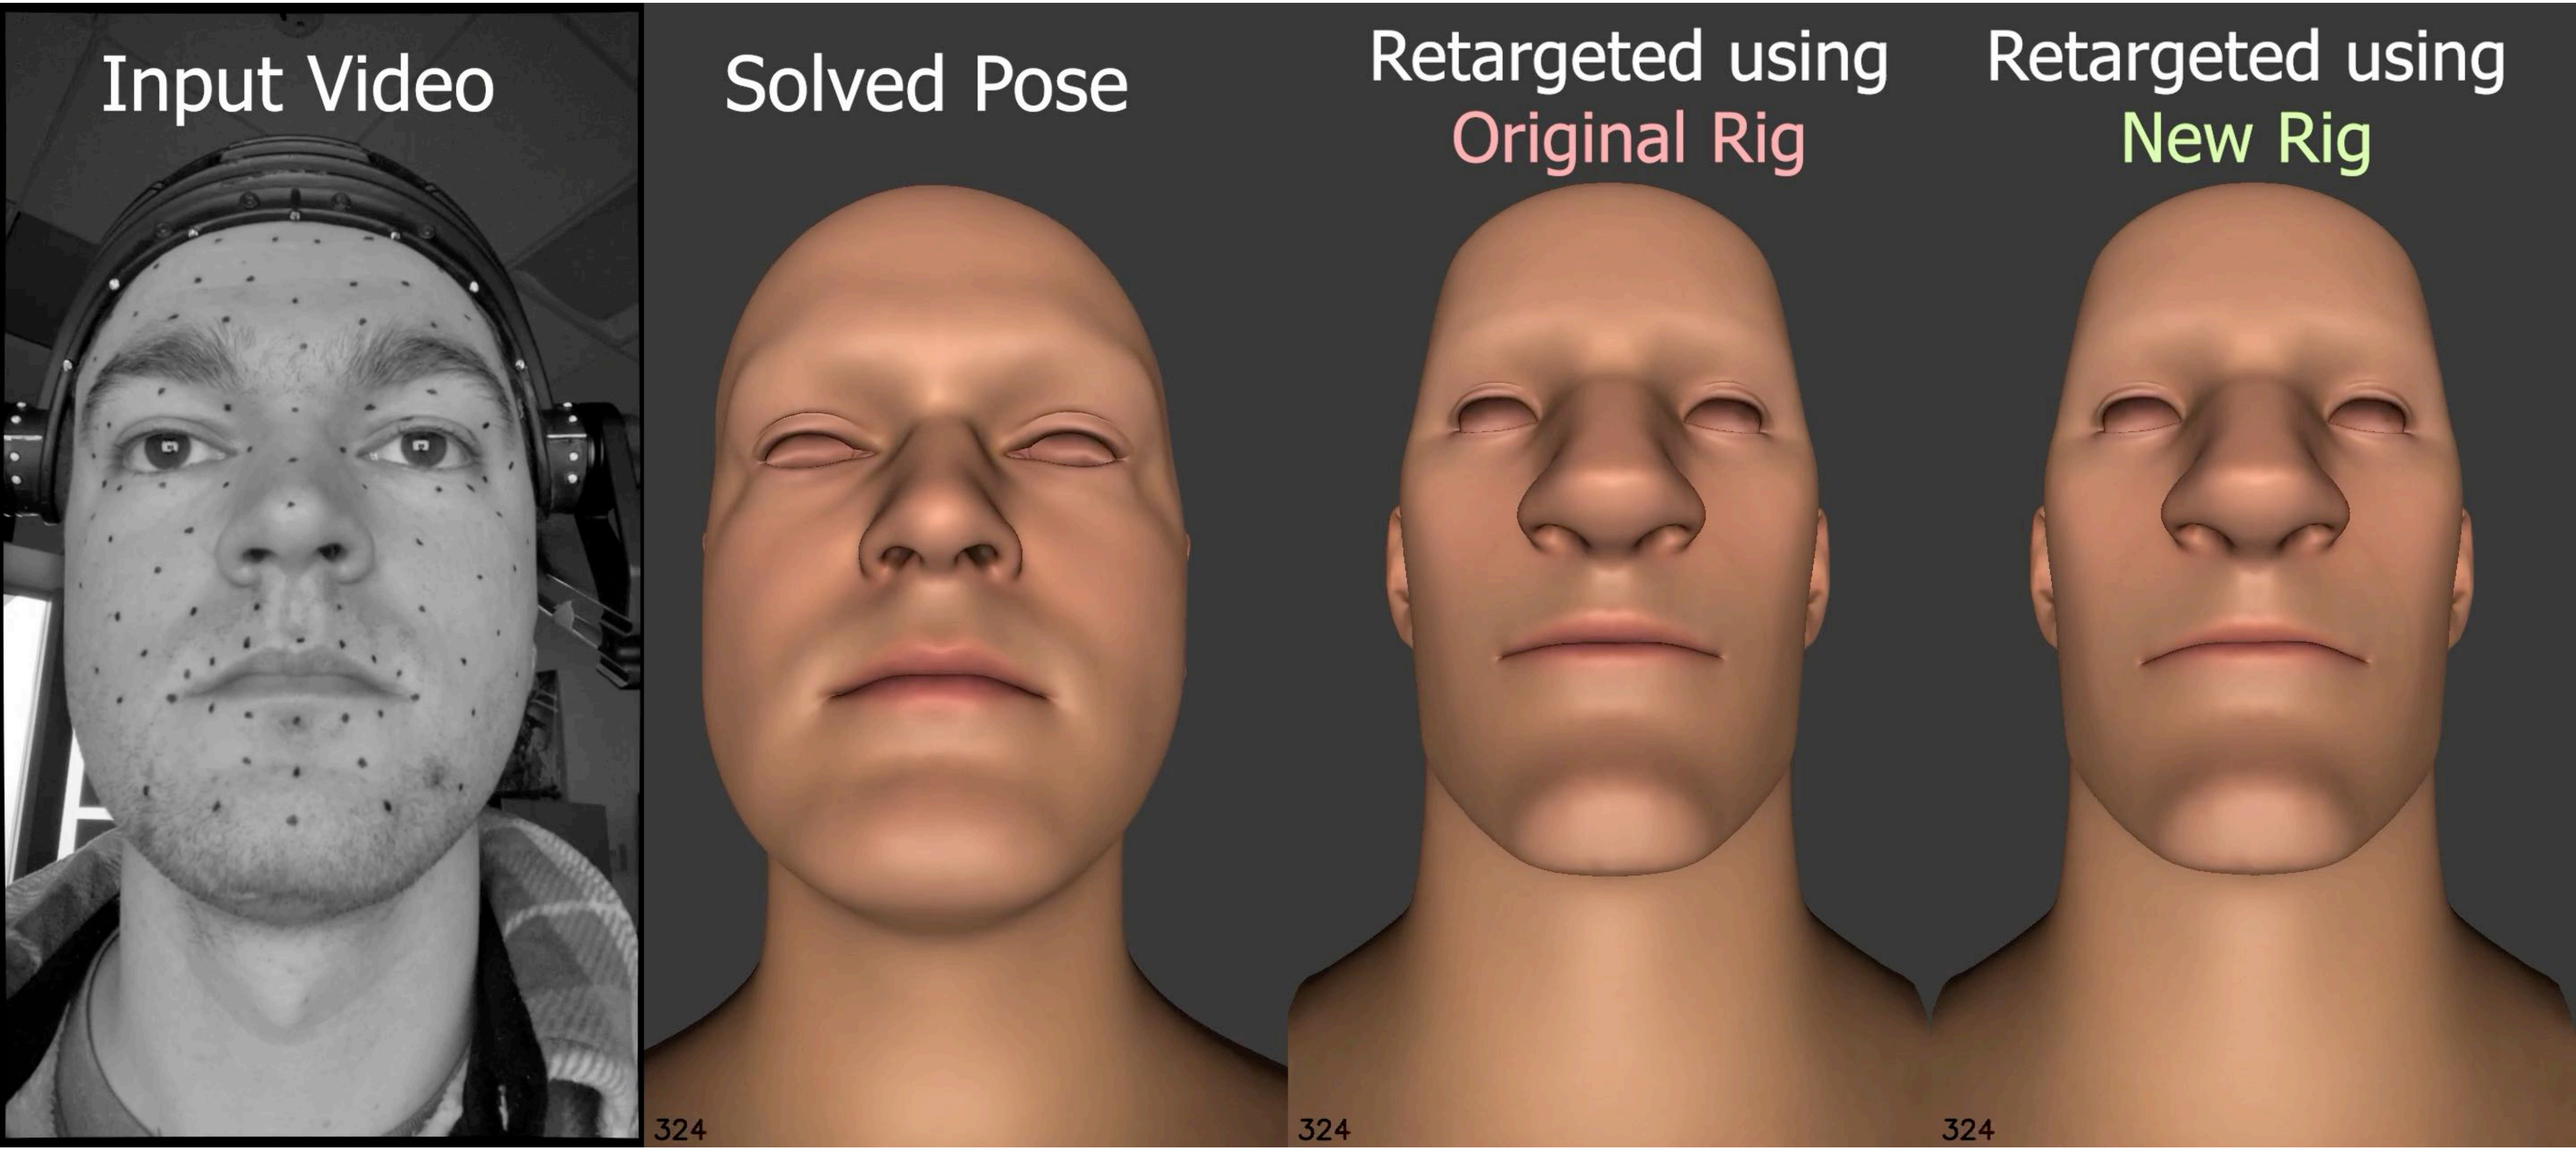

Karen Liu

Human Motion Synthesis

Scene-aware H**uman** Behaviors

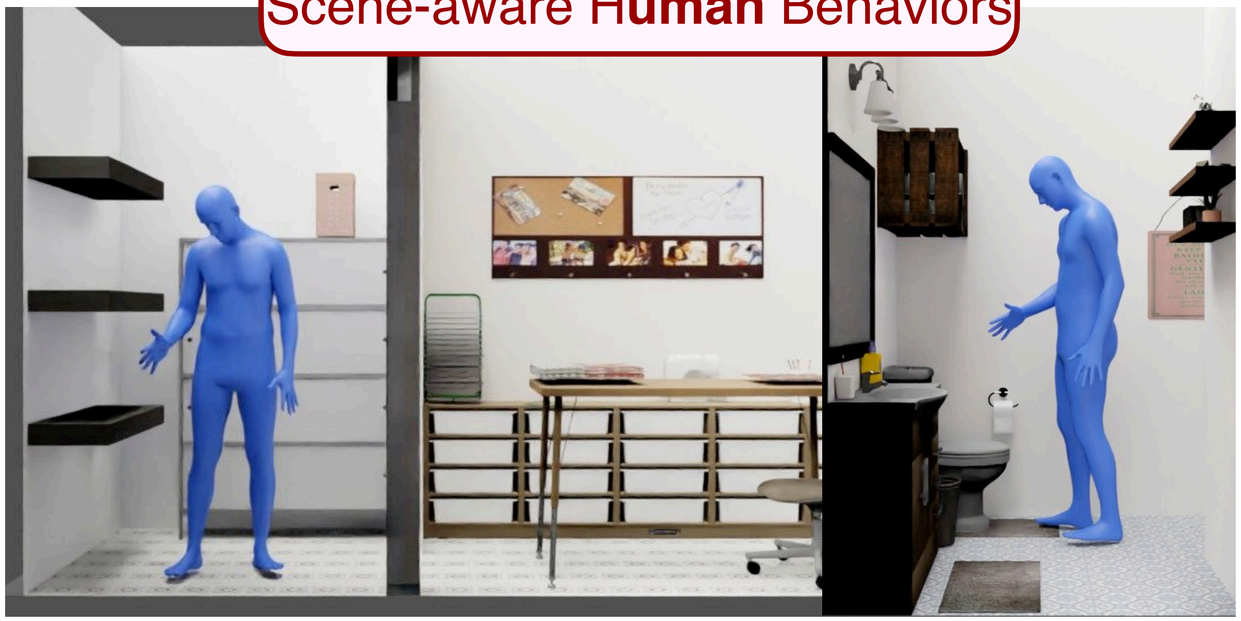

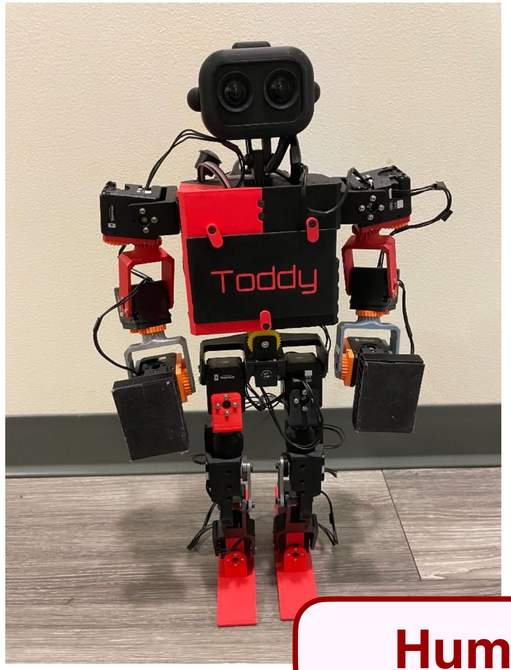

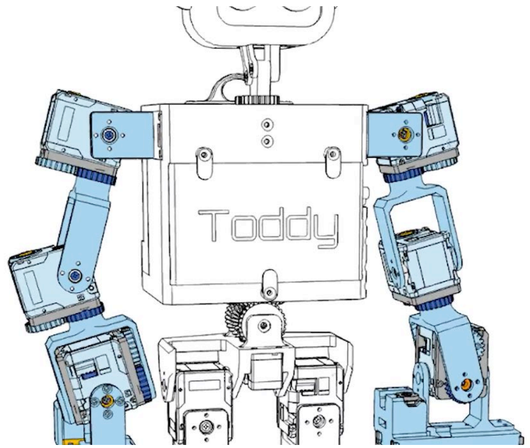

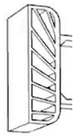

Humanoids!

Human Dexterity

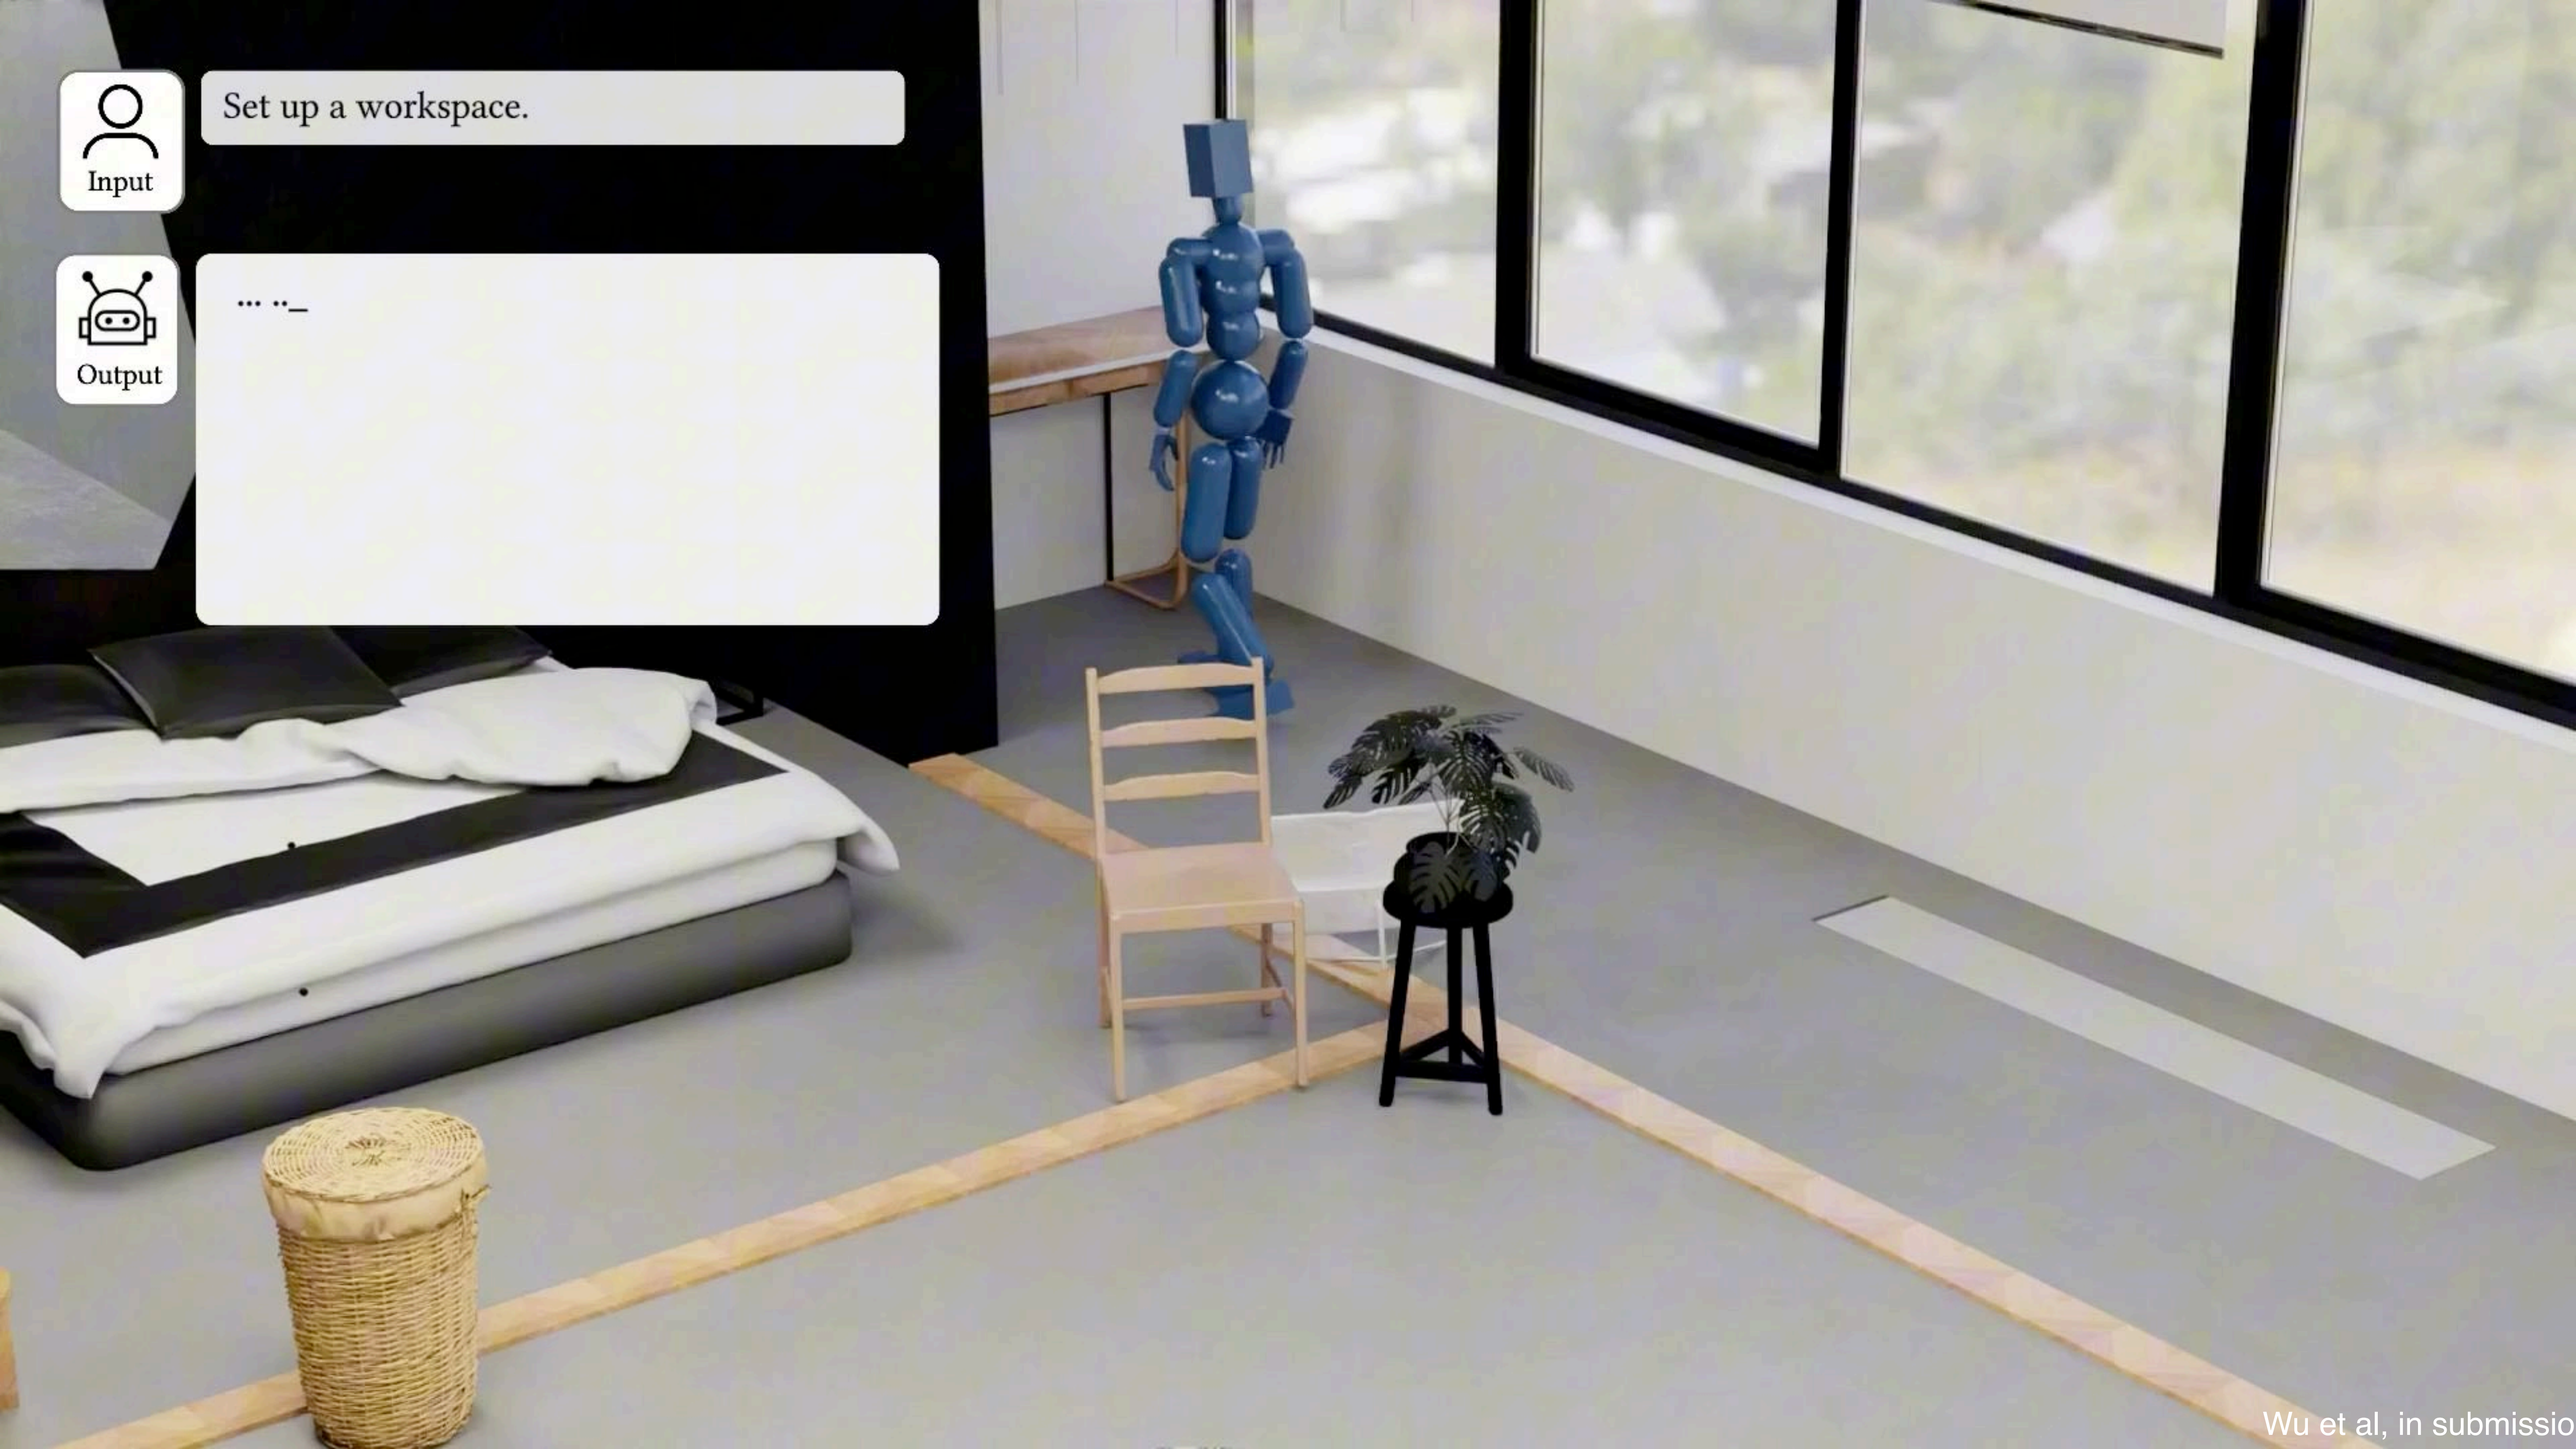
Input
Set up a workspace.
Output
... .._

Jiajun Wu: generating scenes by mixing latents and traditional scene representations

```
(bind "chessboard"
  (lambda (embedding ...)
    (union
      (call "board" ...)
      (call "chess pieces" ...))))

(bind "board"
  (lambda (embedding ...)
    (define (create-squares i) ...)
    (union-loop 64 create-squares)))

(bind "chess pieces"
  (lambda (embedding ...)
    (define (create-pawns i)
      (let* ((white-pawn (call "pawn" ...))
             (black-pawn (call "pawn" ...))
             (white-pose (translate (list (* i 0.125) 0.02 0.125)))
             (black-pose (translate (list (* i 0.125) 0.02 0.75)))
             (white-pawn-t (transform white-pawn white-pose))
             (black-pawn-t (transform black-pawn black-pose)))
        (union white-pawn-t black-pawn-t)))
    (define pieces-order
      '("rook" "knight" "bishop" "queen" "king"
        "bishop" "knight" "rook"))
    (define (create-other-pieces i)
      (let* ((piece-order (list-ref pieces-order i))
             (white-piece (call piece-order ...))
             (black-piece (call piece-order ...))
             (white-pose (translate (list (* i 0.125) 0.02 0)))
             (black-pose (translate (list (* i 0.125) 0.02 0.875)))
             (white-piece-t (transform white-piece white-pose))
             (black-piece-t (transform black-piece black-pose)))
        (union white-piece-t black-piece-t)))
    (union (union-loop 8 create-pawns)
           (union-loop 8 create-other-pieces))))
```

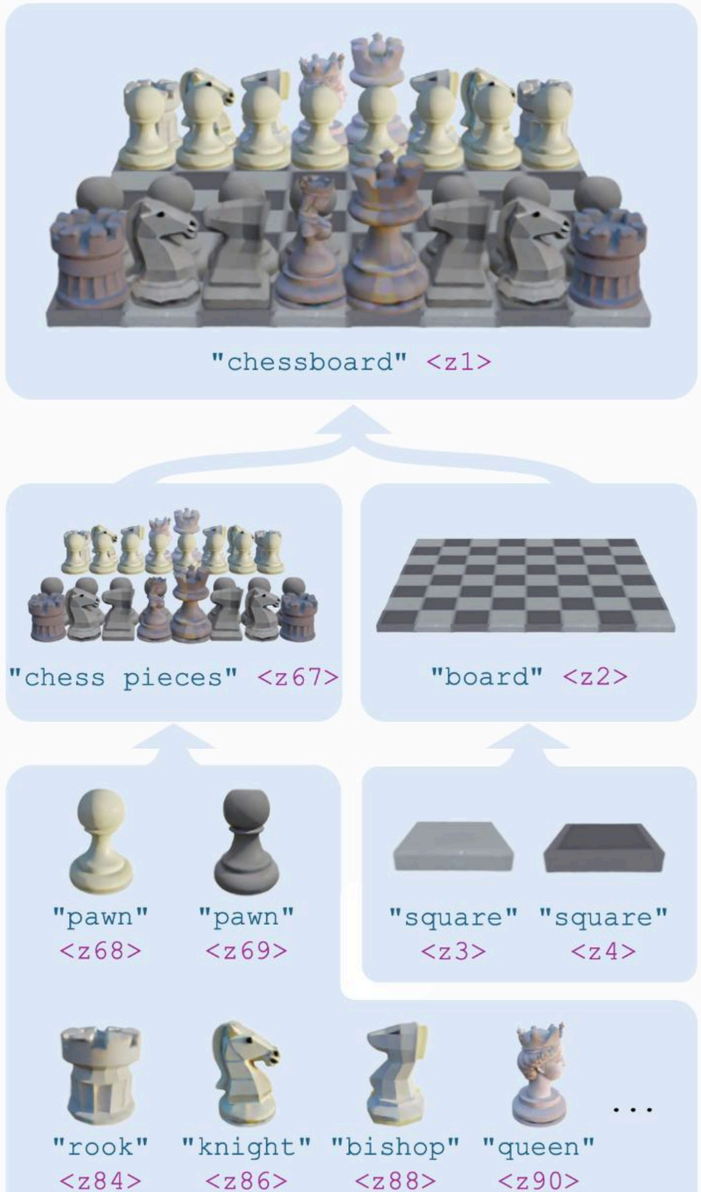

Jiajun Wu: conditioning motions on scenes / and conditioning scenes on motions

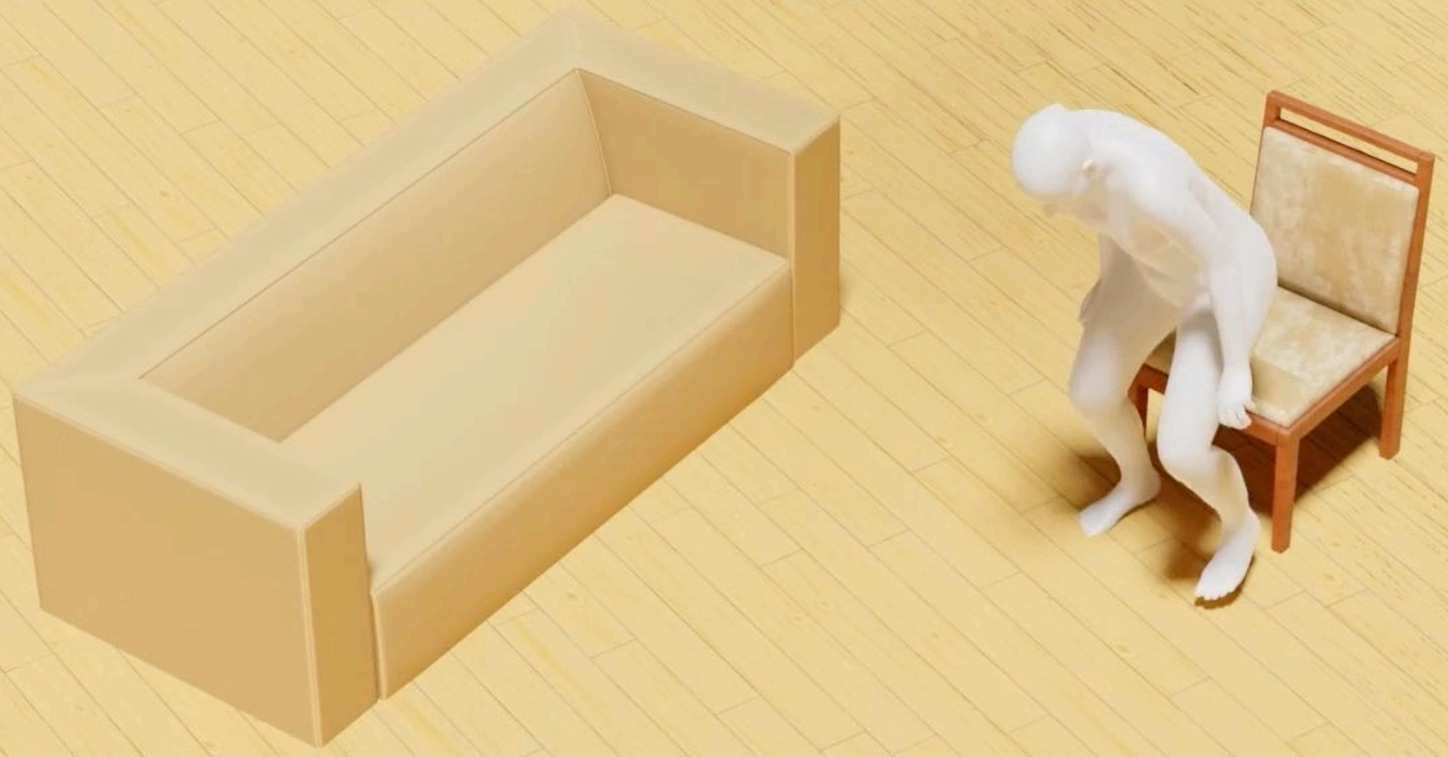

Jiajun Wu: combining physical knowledge with AI to generate plausible dynamics in video

Recent Leo 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

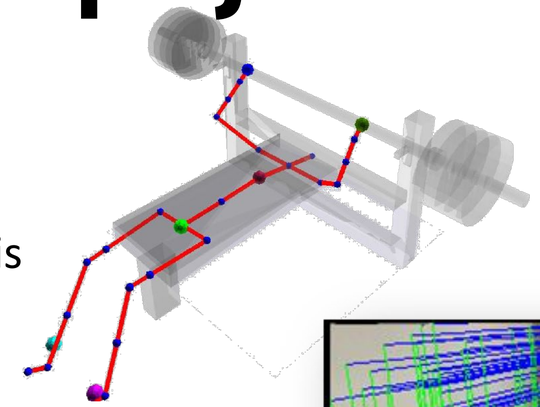

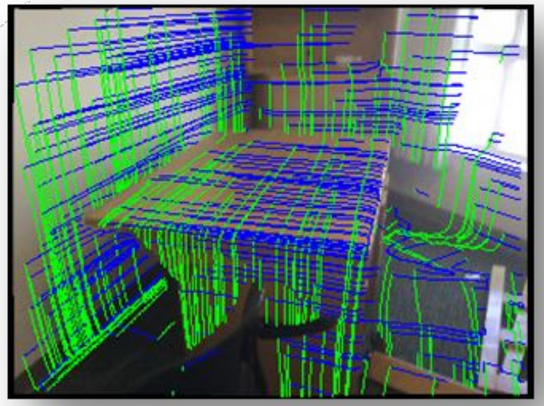

Algorithmic problems in modeling physical objects and phenomena in vision/graphics/robotics

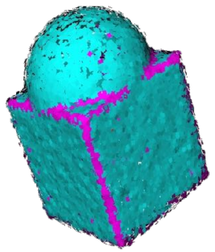

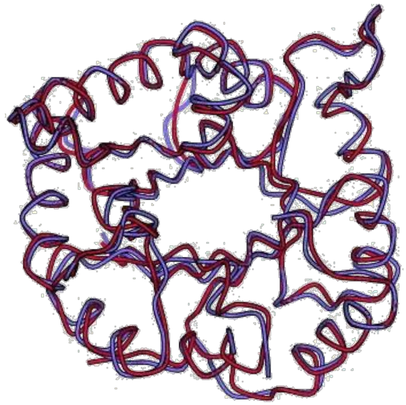

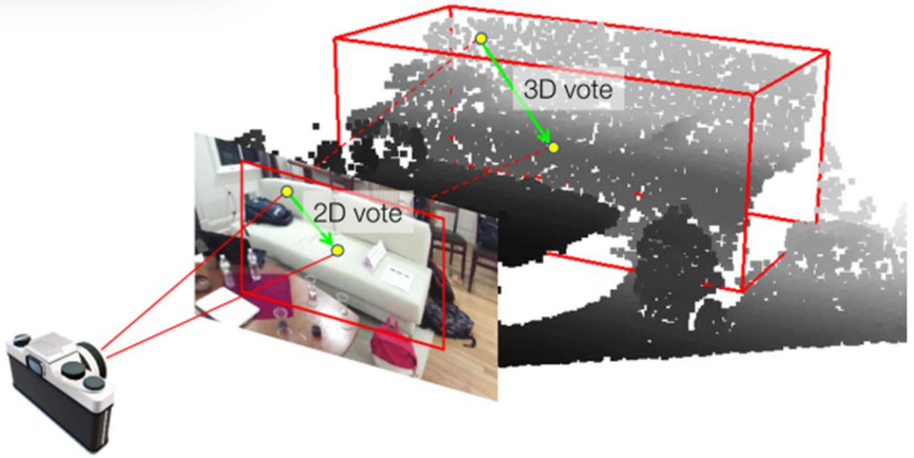

Leo Guibas

Geometry processing and analysis

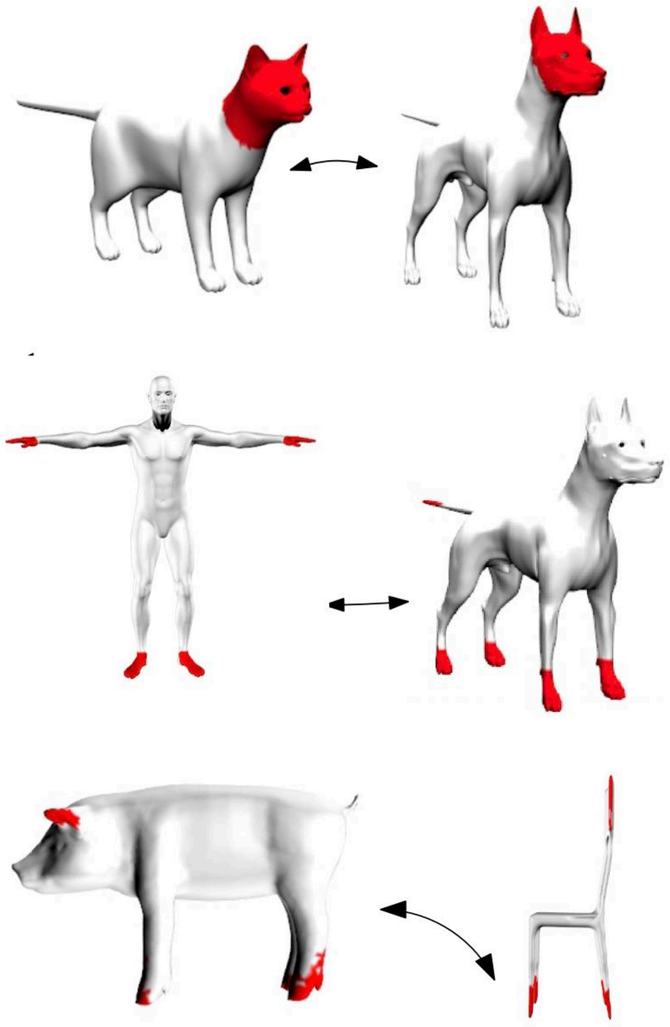

Shape Similarity and Correspondence

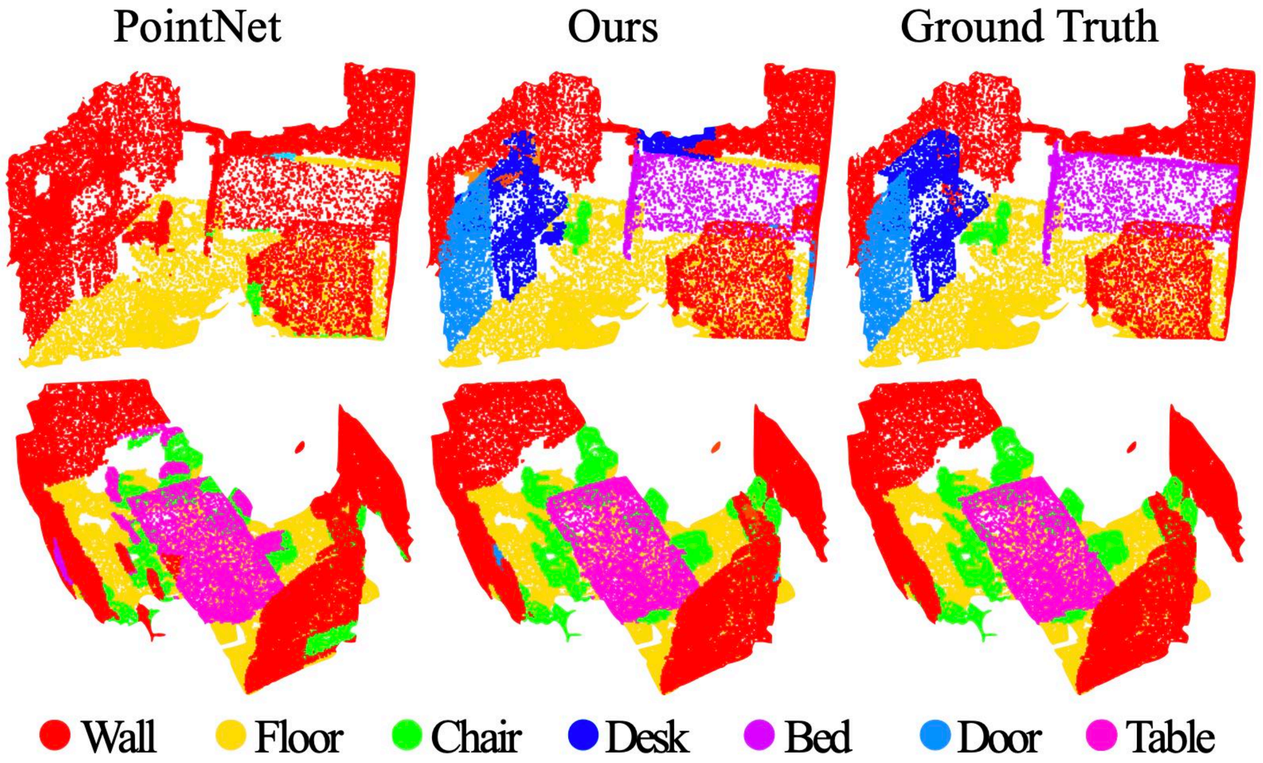

PointNet: Deep Learning on Point Clouds

Gordon Wetzstein: Stanford Computational Imaging Lab

Neural Rendering and 3D generative AI

XR & Wearable Computing

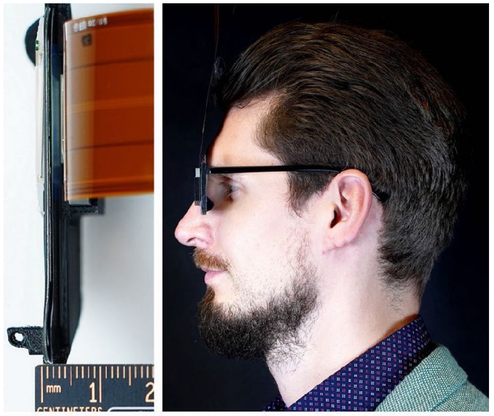

Deep Optics

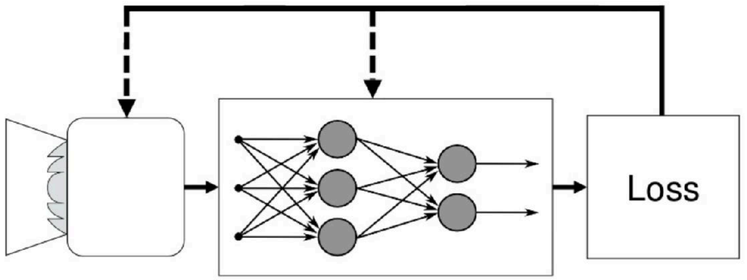

Single-photon Imaging

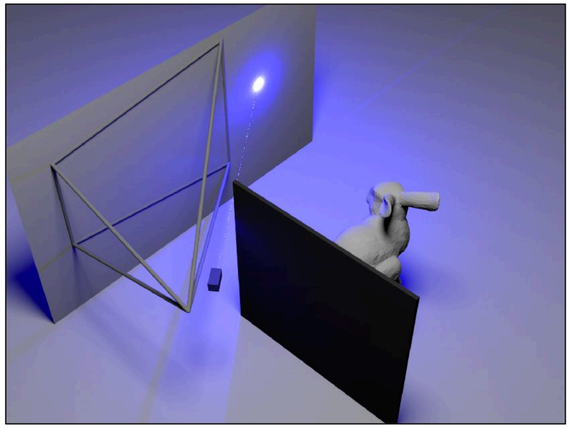

Computational Microscopy

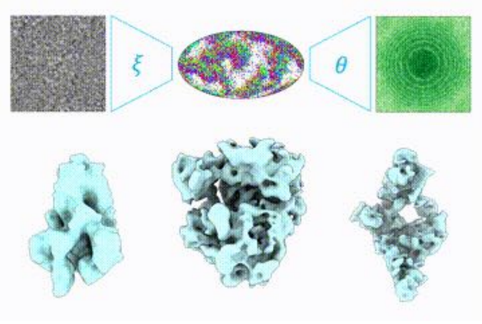

Computational Cameras

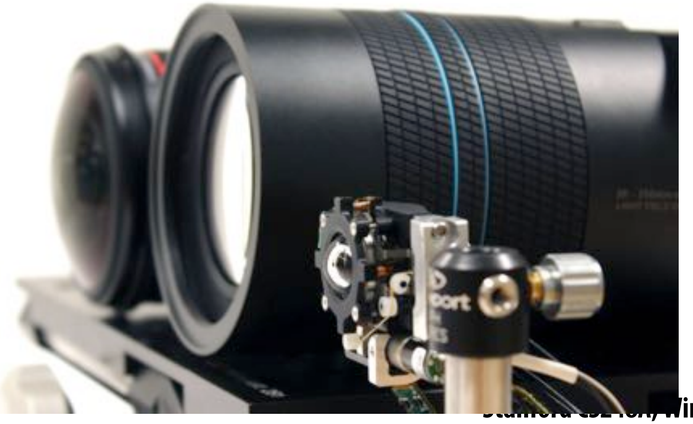

Expanding capabilities of 3D generative AI

Chan et al. CVPR 2022

Improved Water Sound Synthesis using Coupled Acoustic Bubbles

Kangrui Xue, Ryan M. Aronson, Jui-Hsien Wang, Timothy R. Langlois, Doug L. James

ACM SIGGRAPH 2023

A framework for automatically synthesizing bubble-based water sounds for fluid animations.

We model inter-bubble coupling forces to simulate the collective oscillations of bubble clouds.

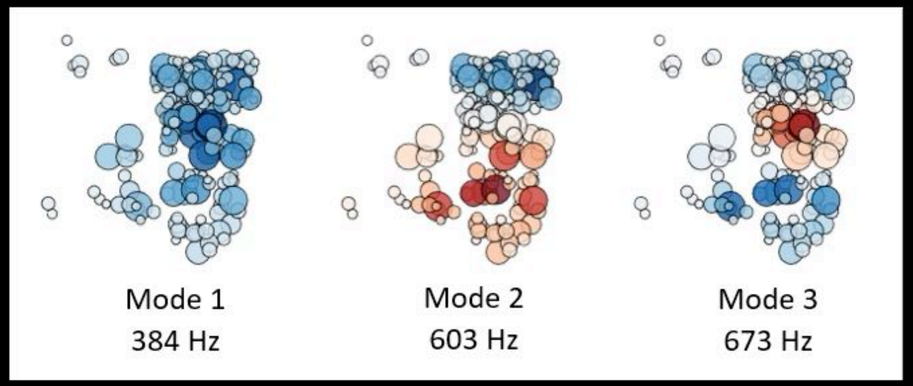

We also improve sound rendering speed and robustness using a GPU wavesolver with sample-and-hold geometry.

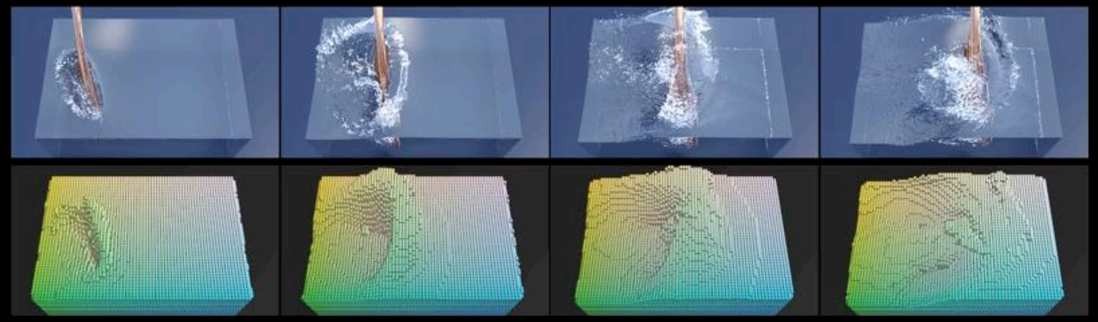

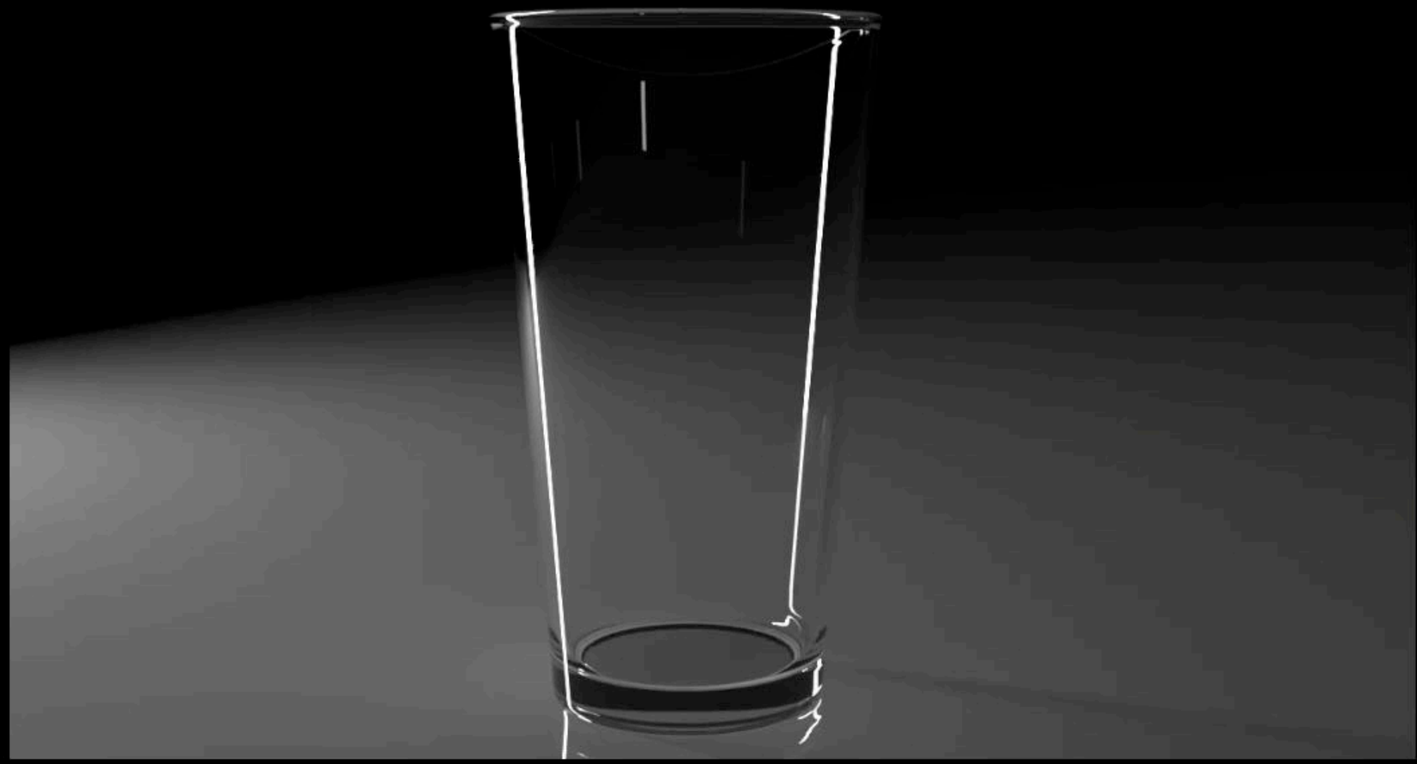

(The following sounds were generated by our method)

Much exciting research on training agents in simulation to perform complex tasks

OpenAI et al, Arxiv 2019

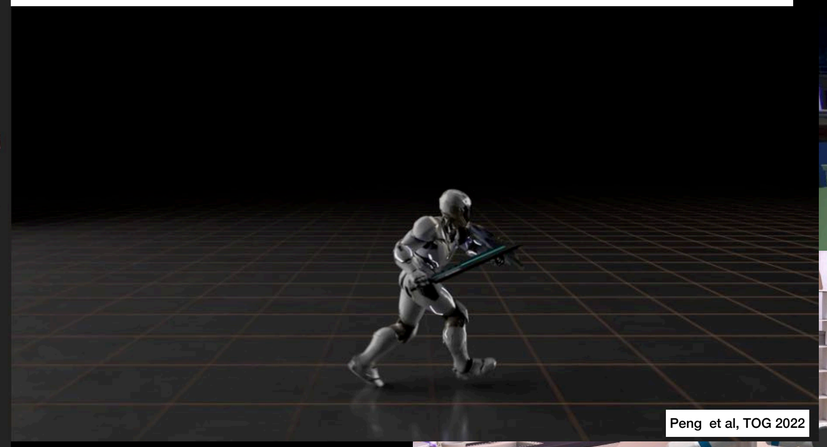
Peng et al, TOG 2022


Zhang et al, TOG 2023

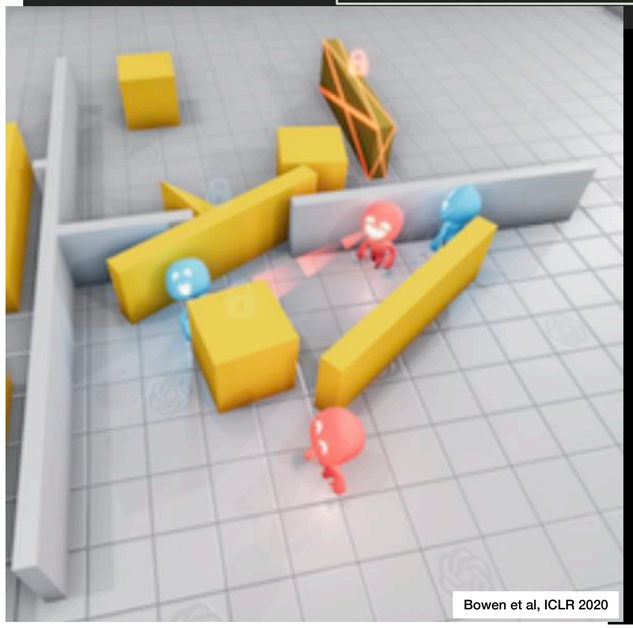
Bowen et al, ICLR 2020

Bellemare et al, JAIR 2013


Szot et al, NeurIPS 2021

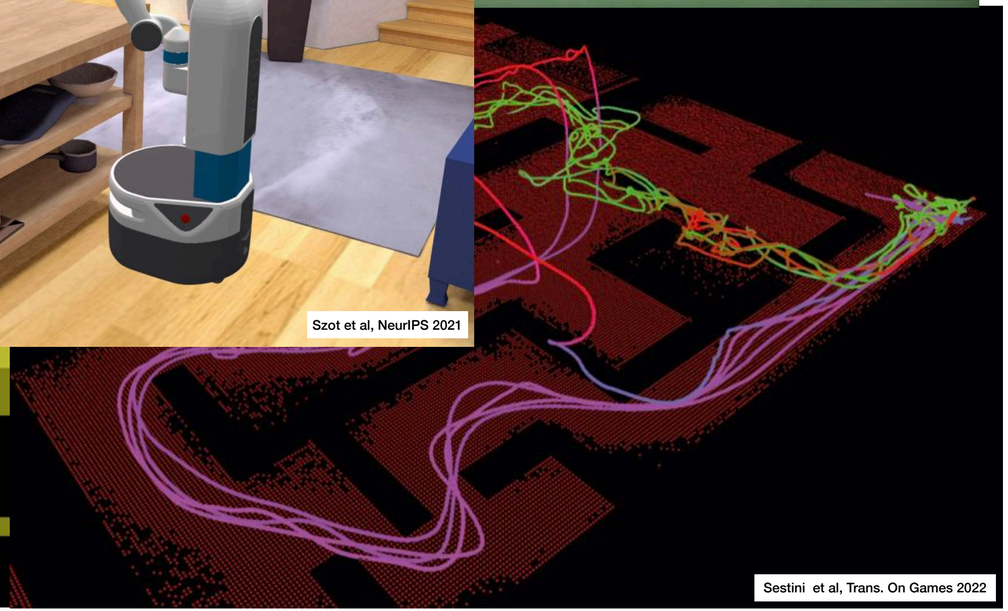
Sestini et al, Trans. On Games 2022

AI agent training loop

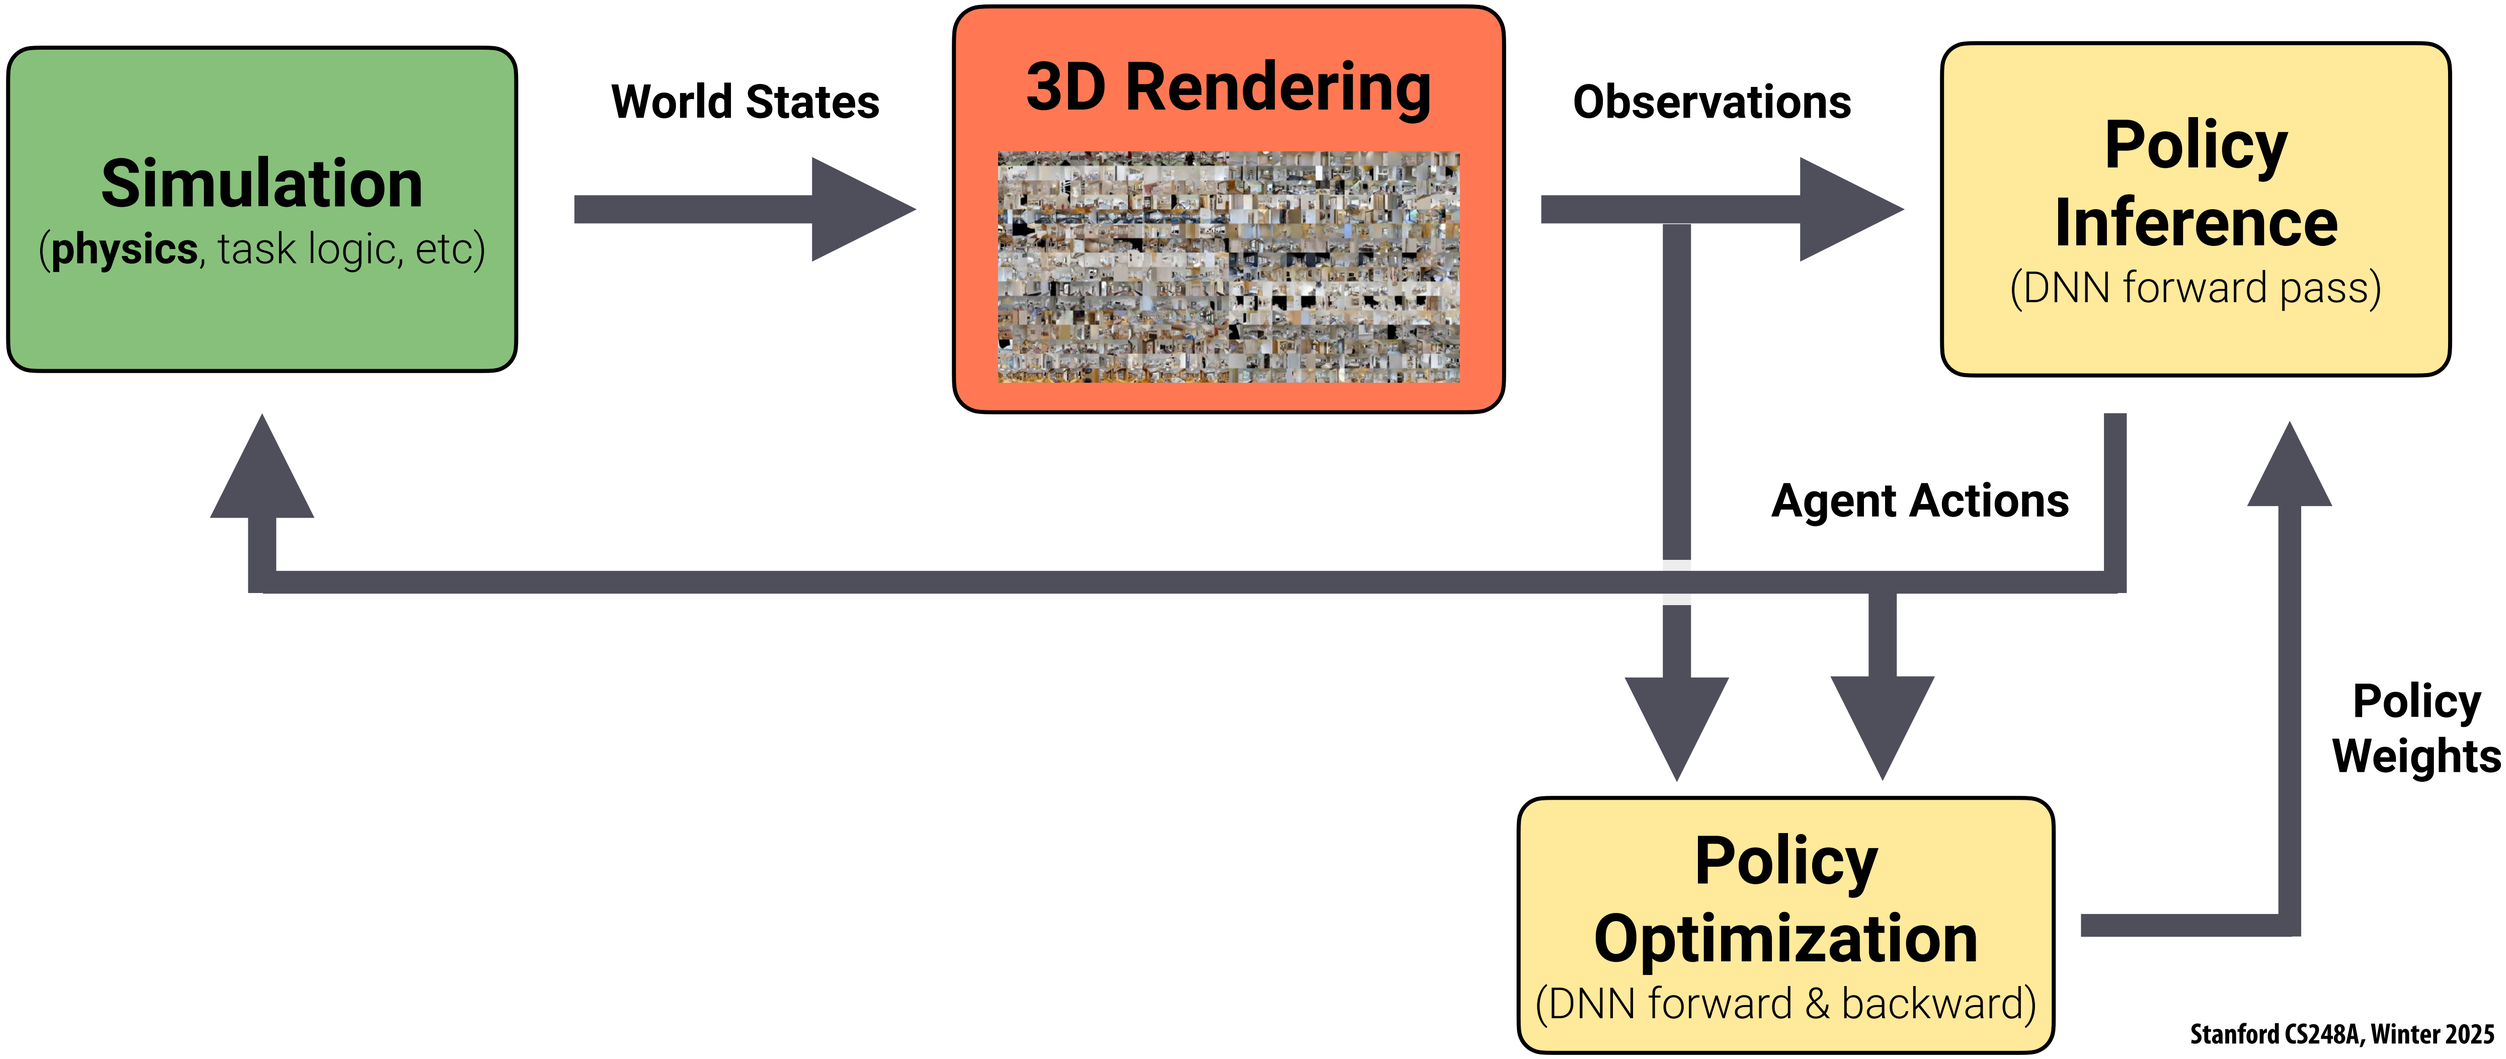

“Batch” simulators: Millions of FPS by executing thousands of environments in parallel on the GPU

Google MJX:
GPU Physics

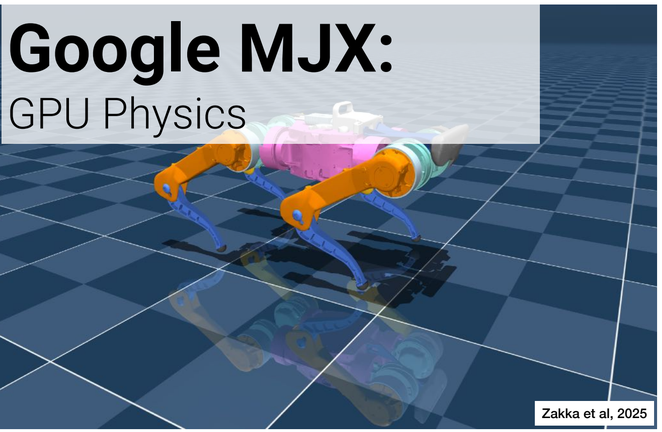

Isaac Gym:
GPU Physics

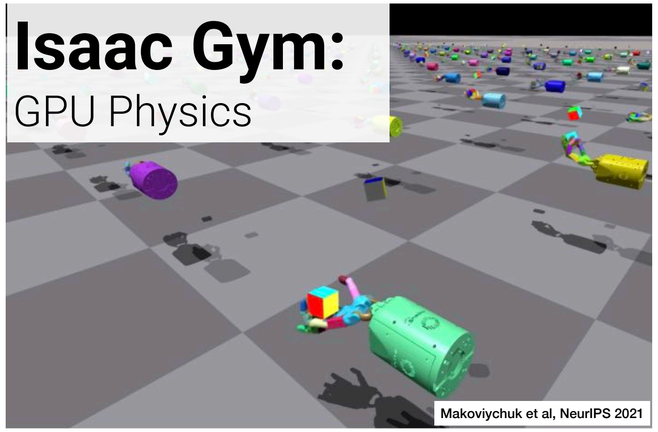

CuLE:
Atari RL Benchmarks

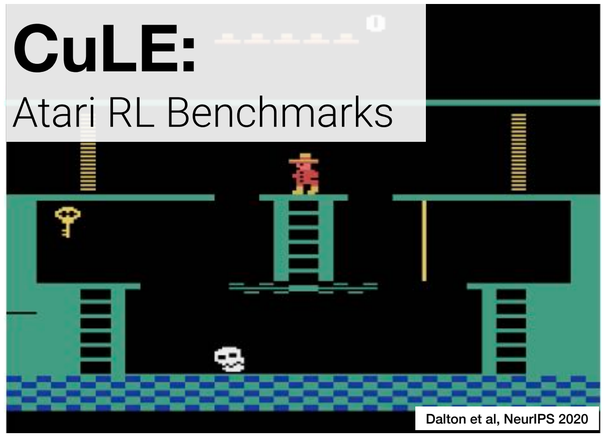

BPS3D:
Home Navigation

Maniskill:
Robotic Manipulation

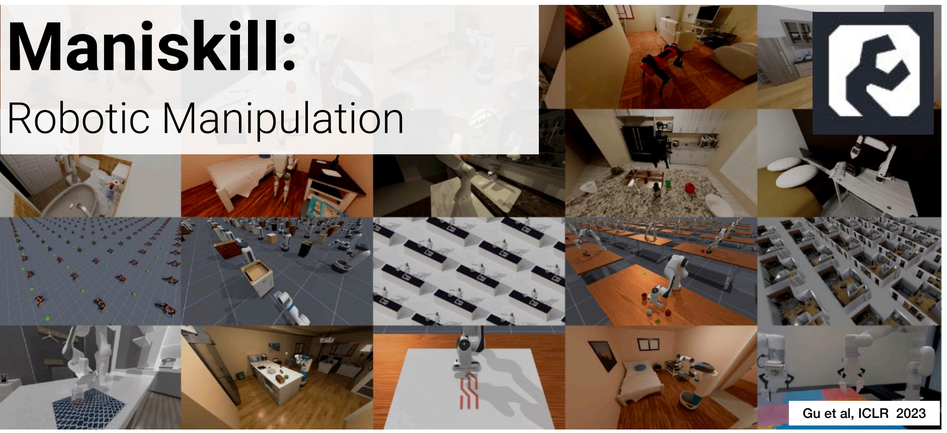

Kinetix:
Physics Based
Problem Solving

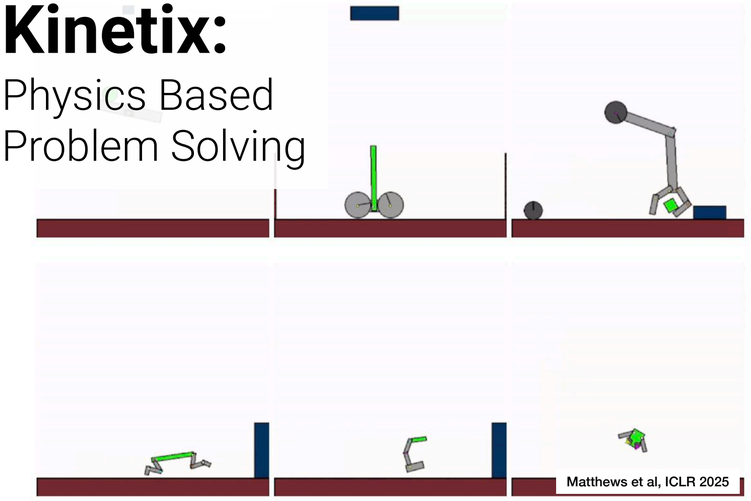

Craftax:
Open-Ended 2D Exploration

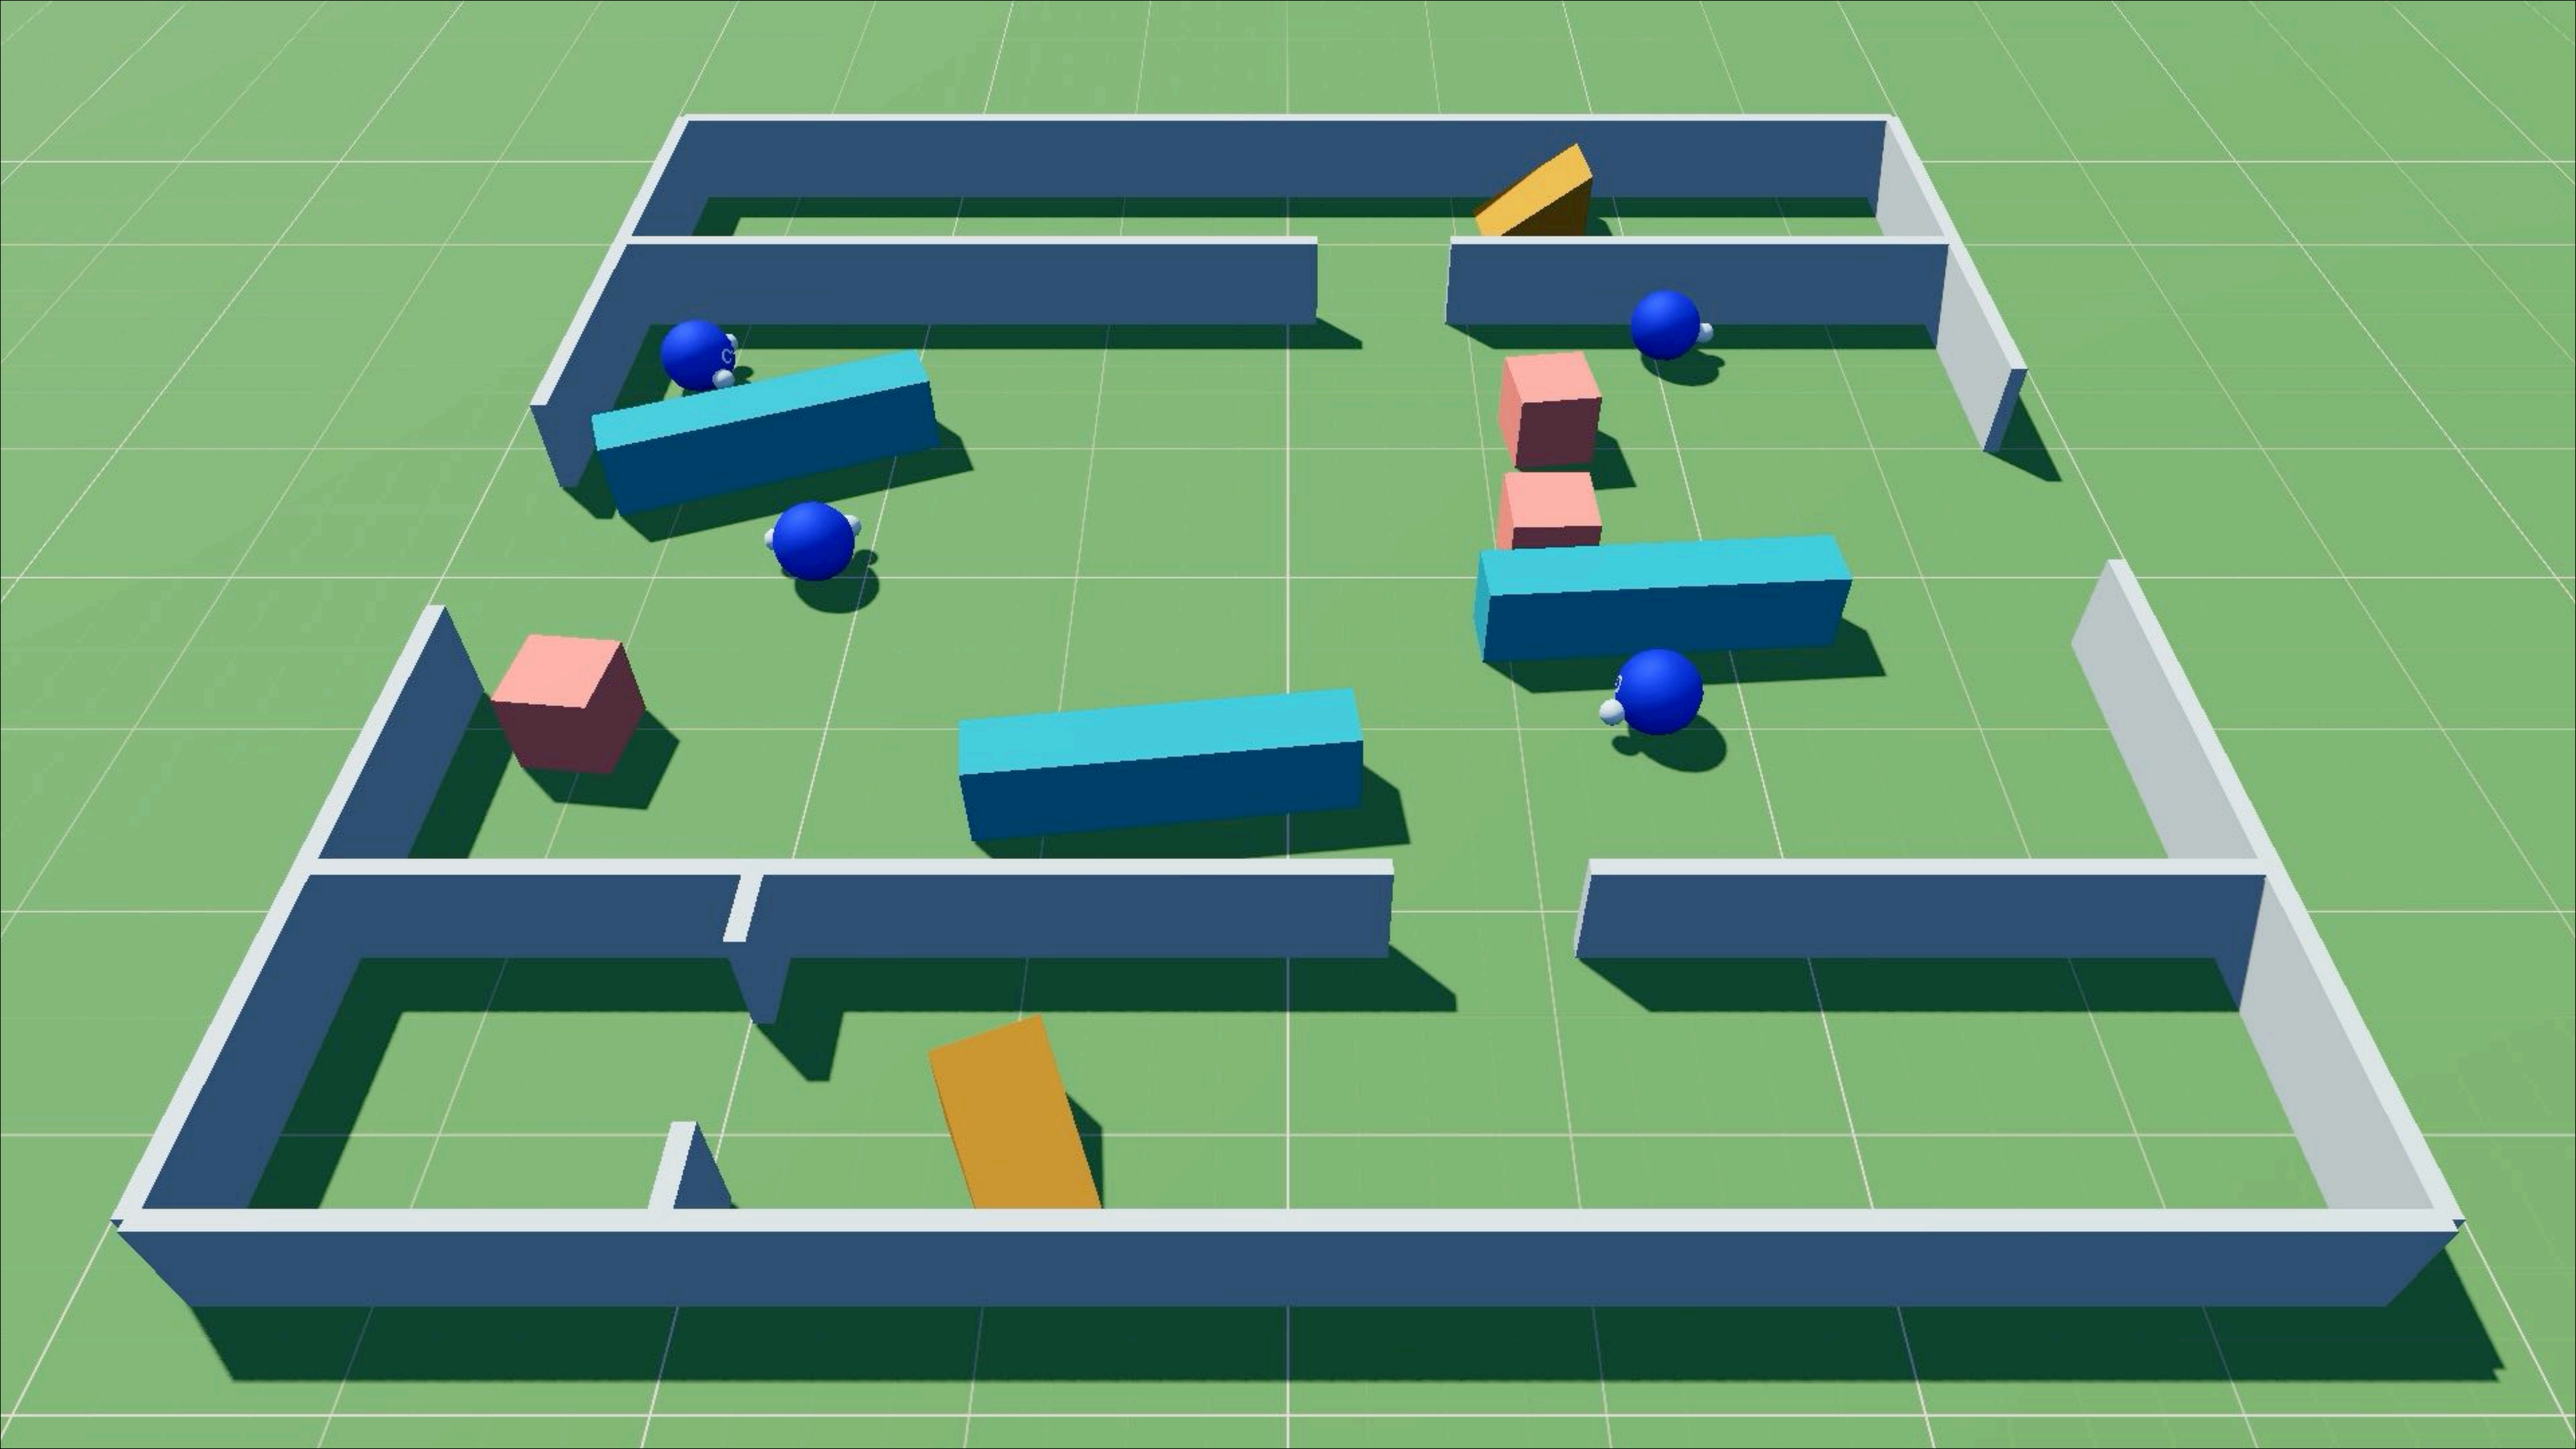

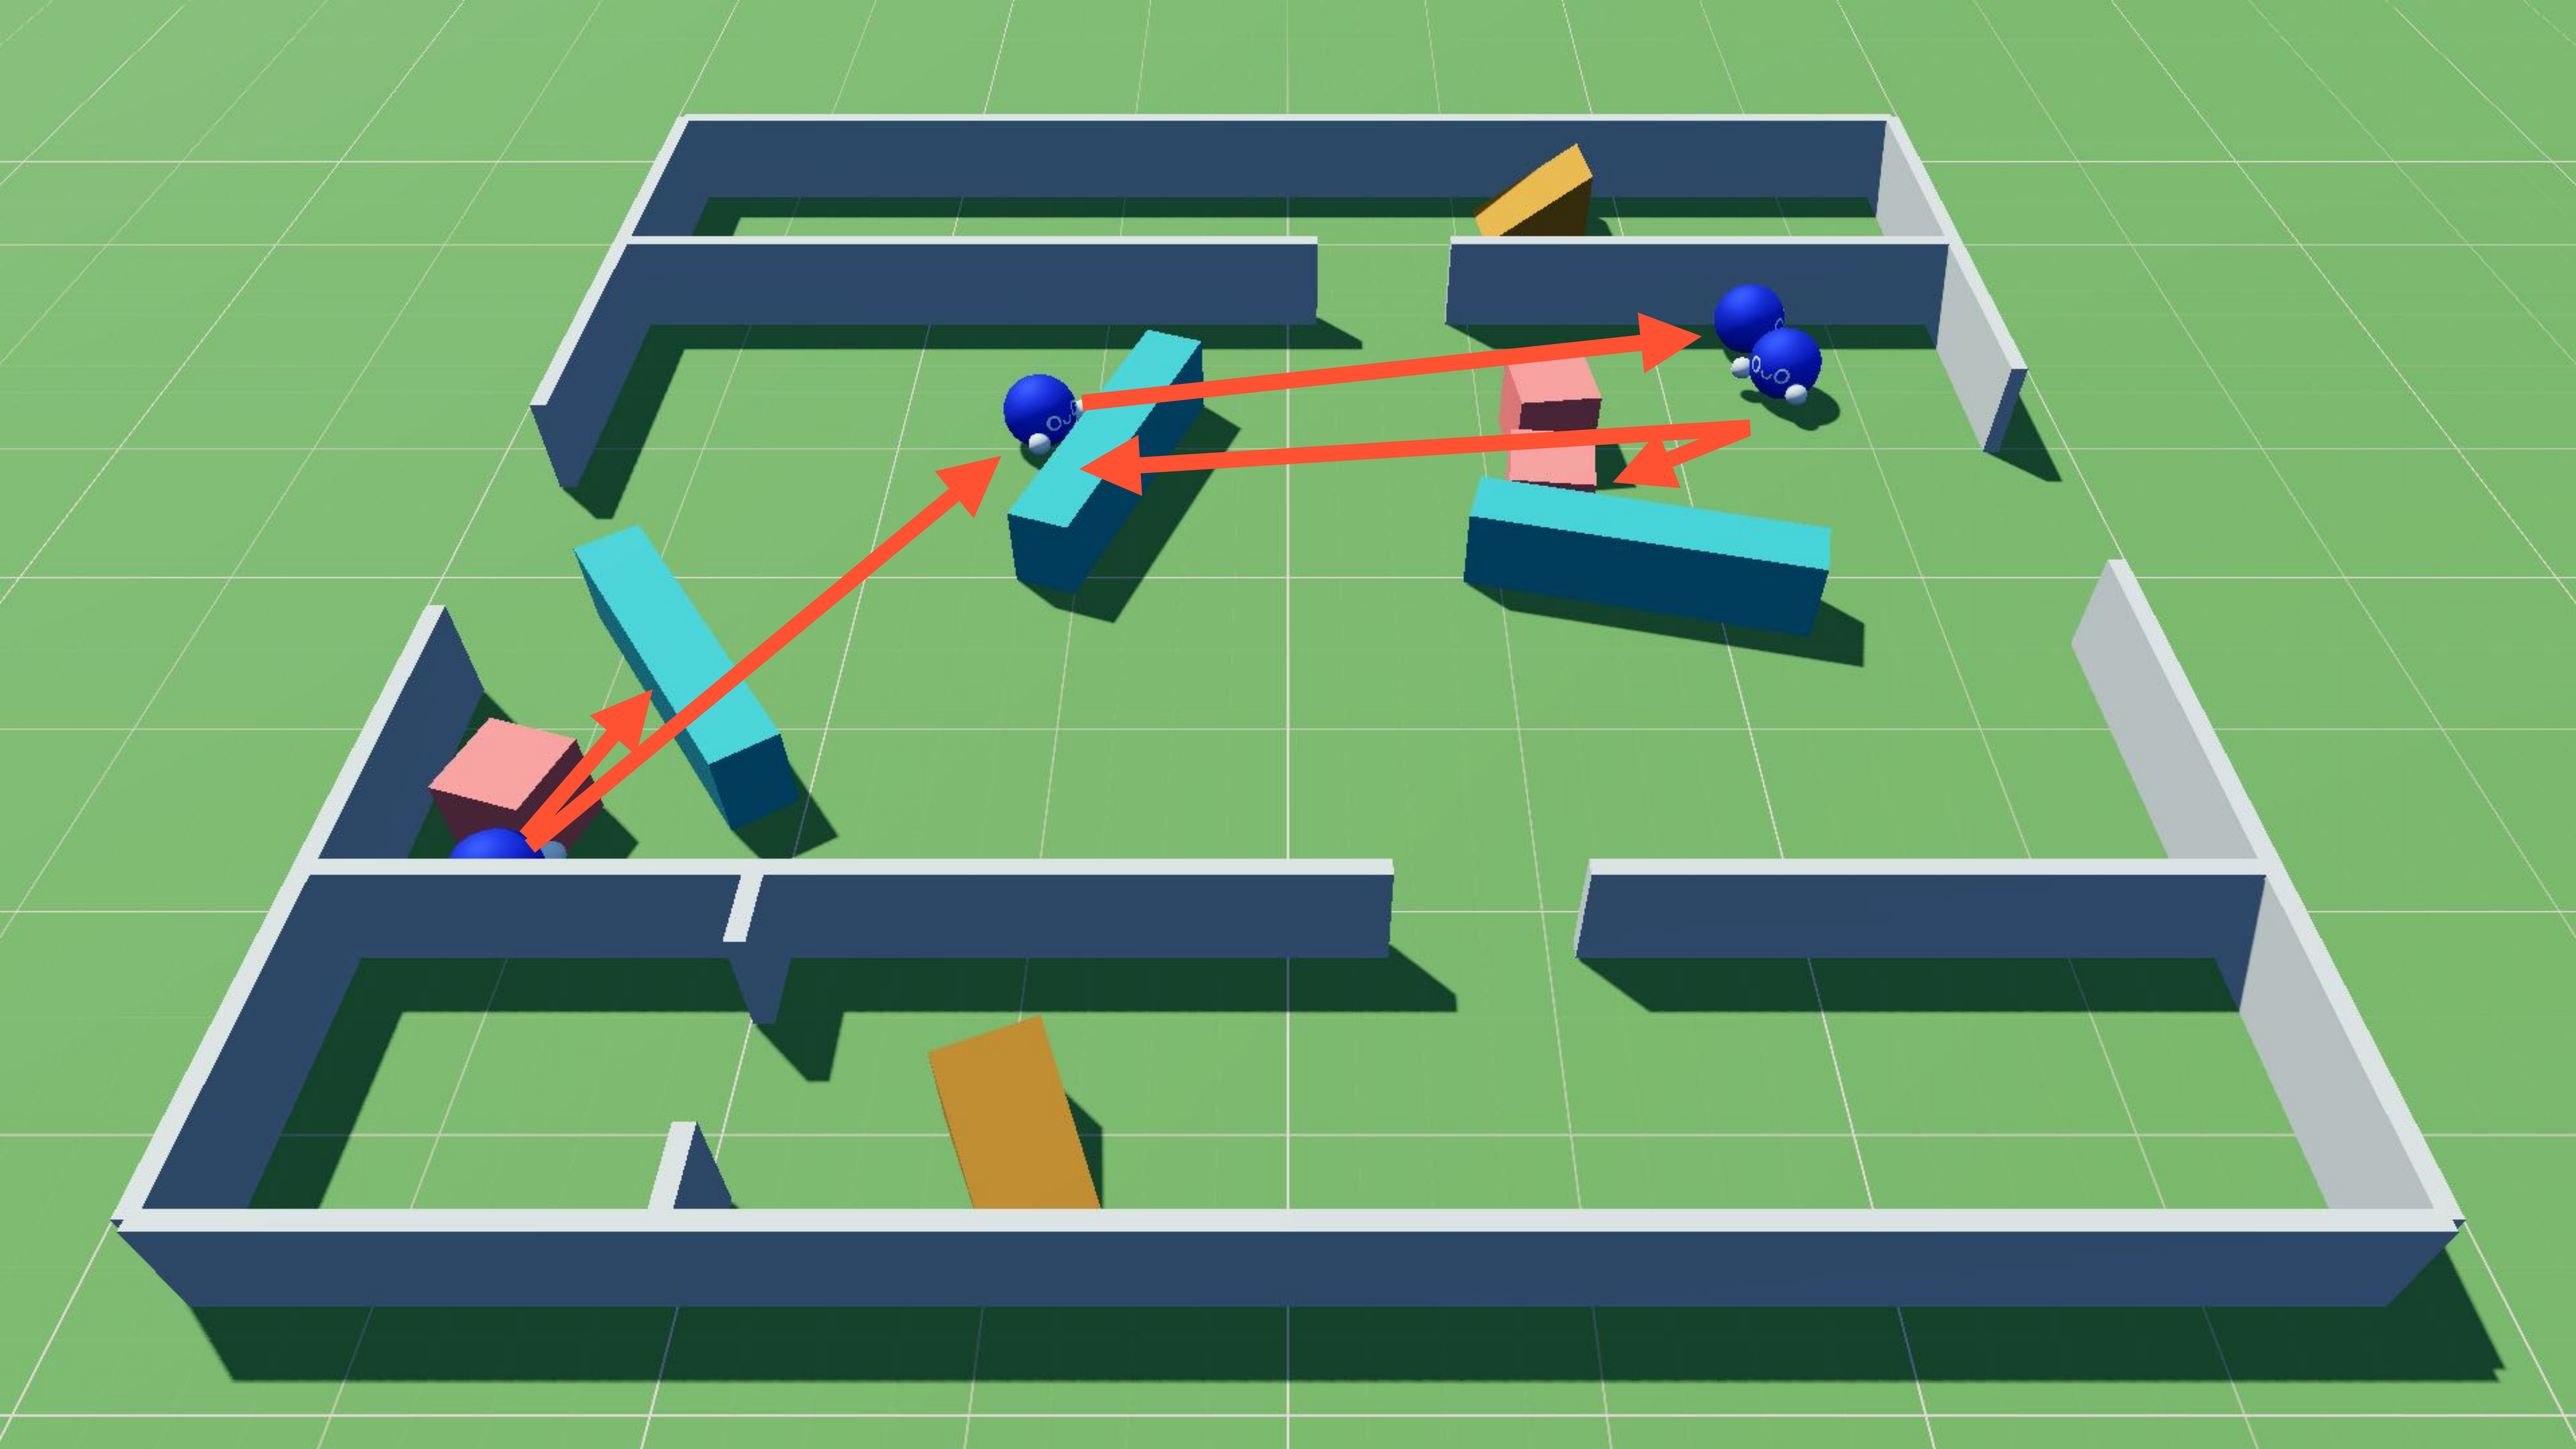

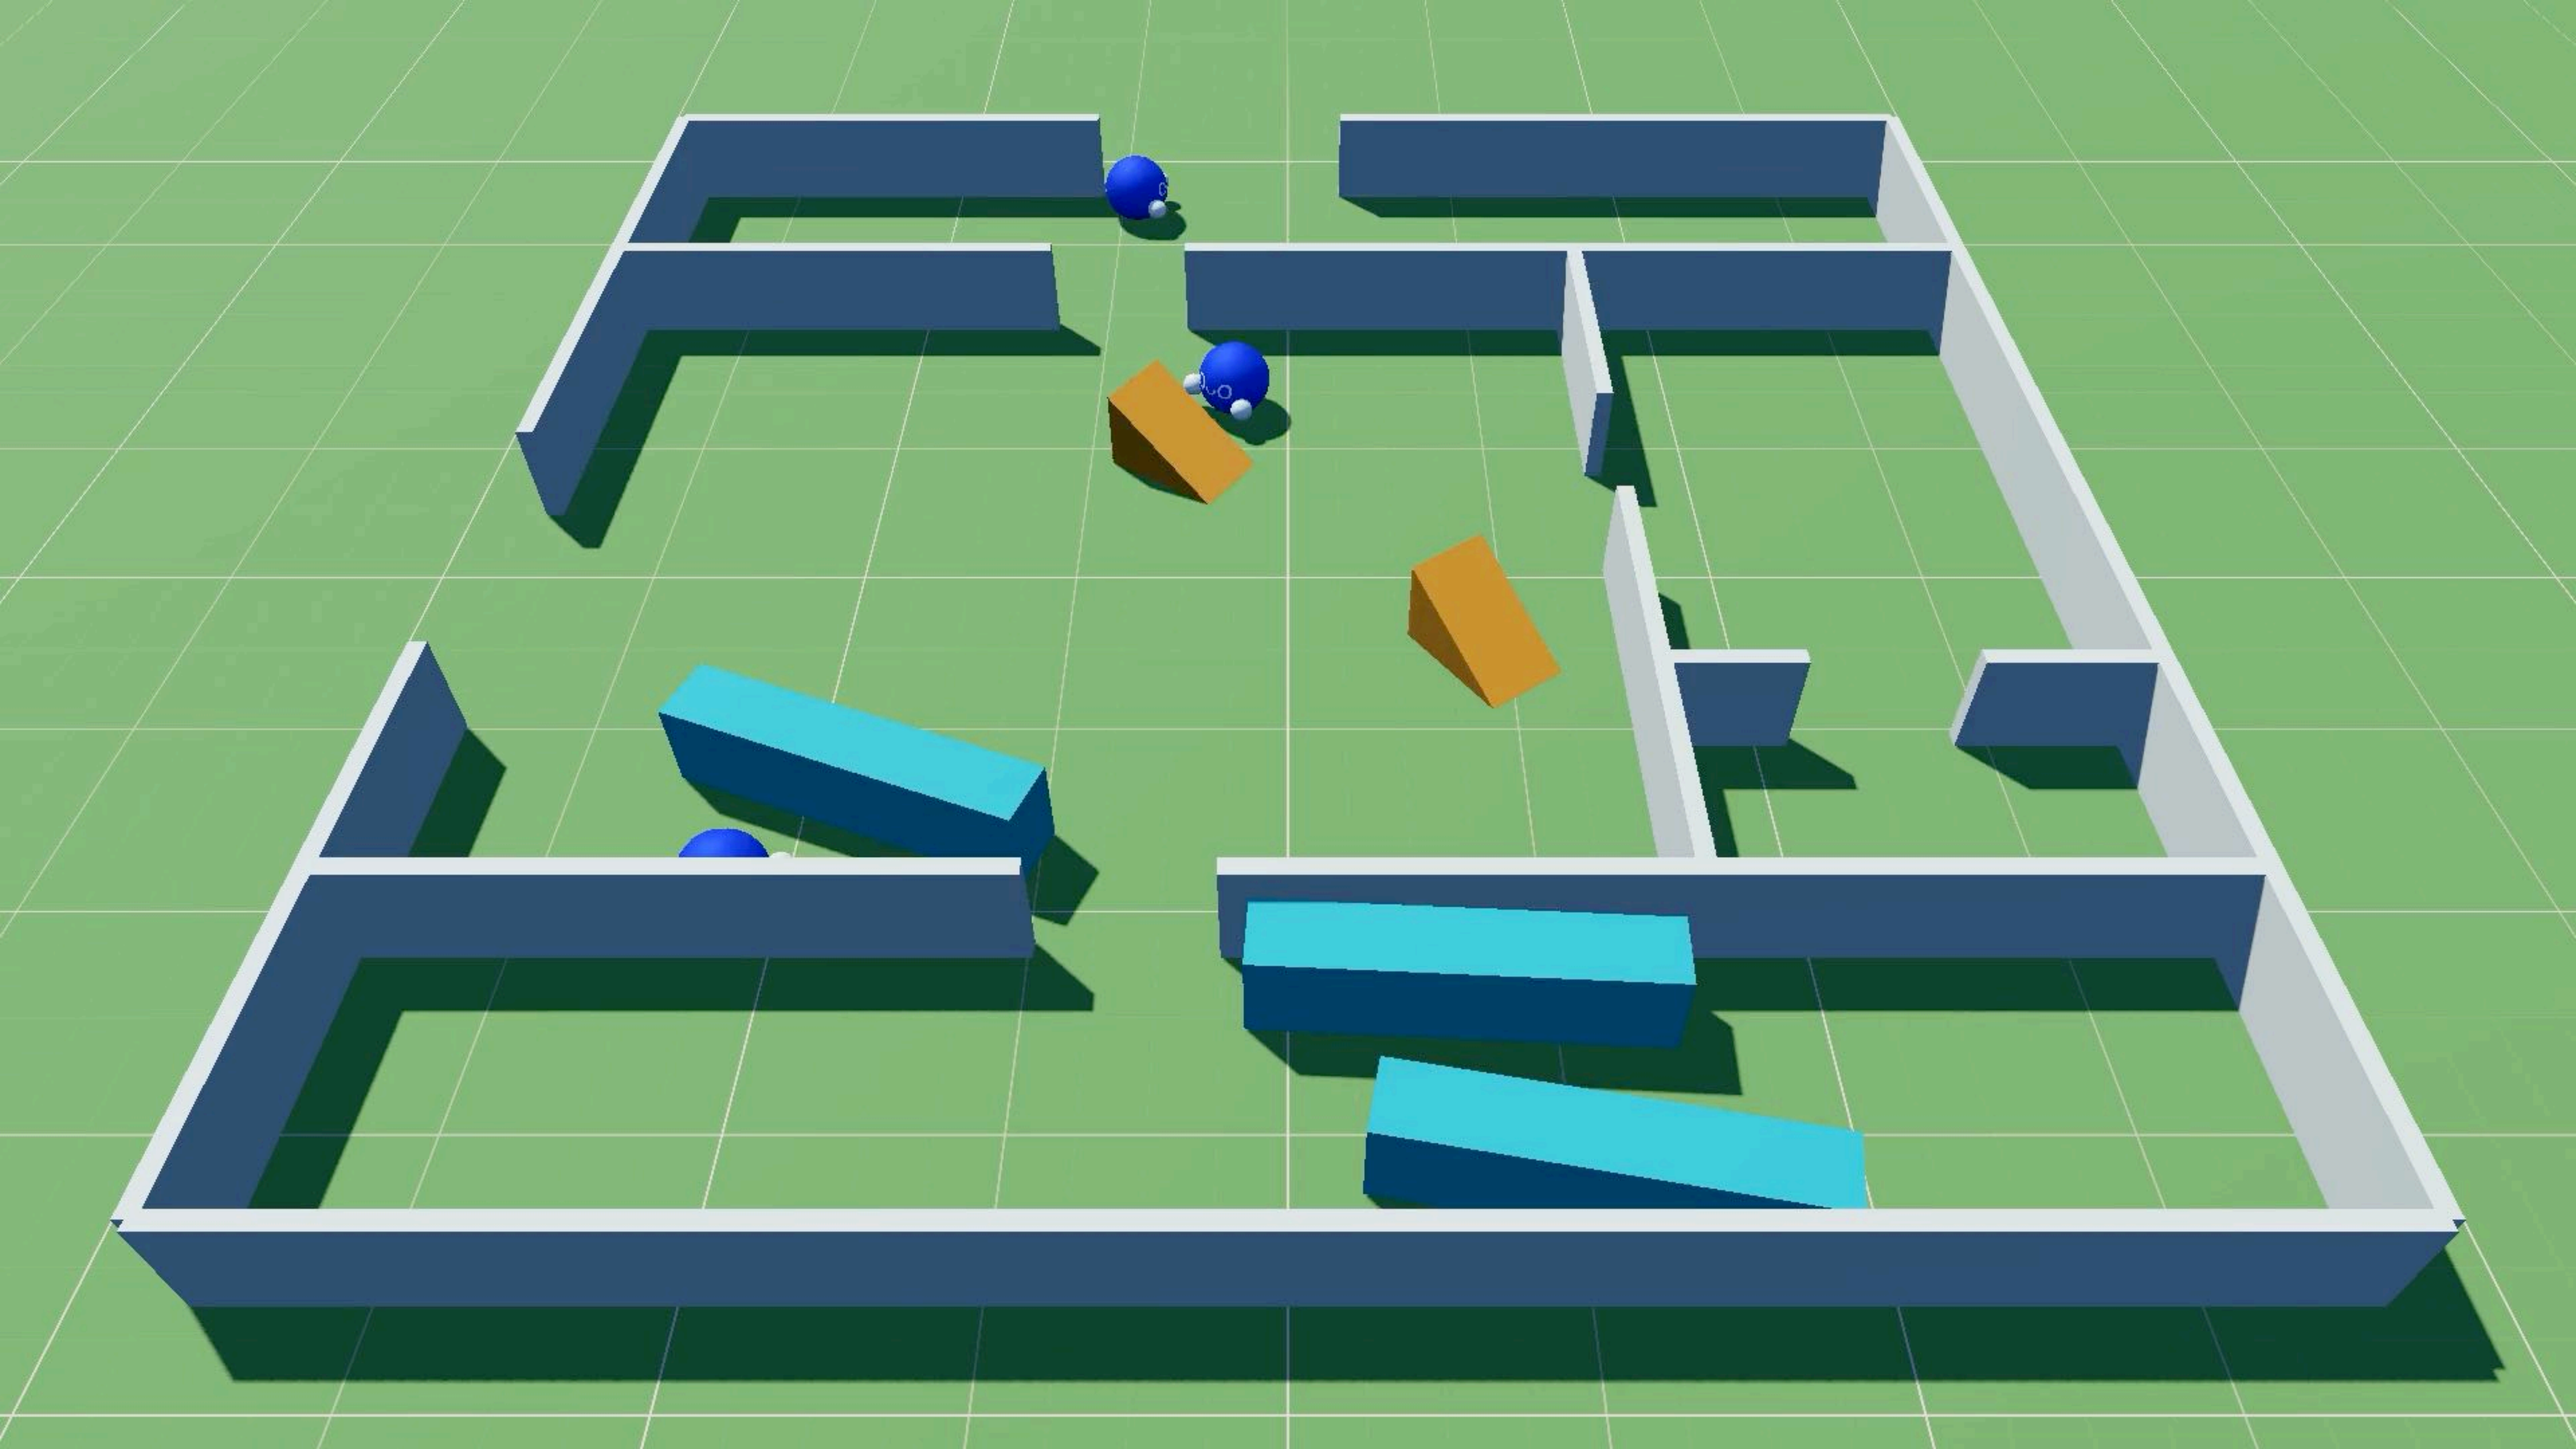

Research idea: We need a "game engine" for building batch simulators!

Maximize throughput: millions of sim steps/sec for simple 3D environments
(When running many environments in parallel)

Programmable: environment creators should be able to author diverse set of worlds, define custom world rules/behavior

Productive: non-parallel systems experts should be able quickly be able to create novel worlds

Batch Hide and Seek batch simulator is orders of magnitude faster than original

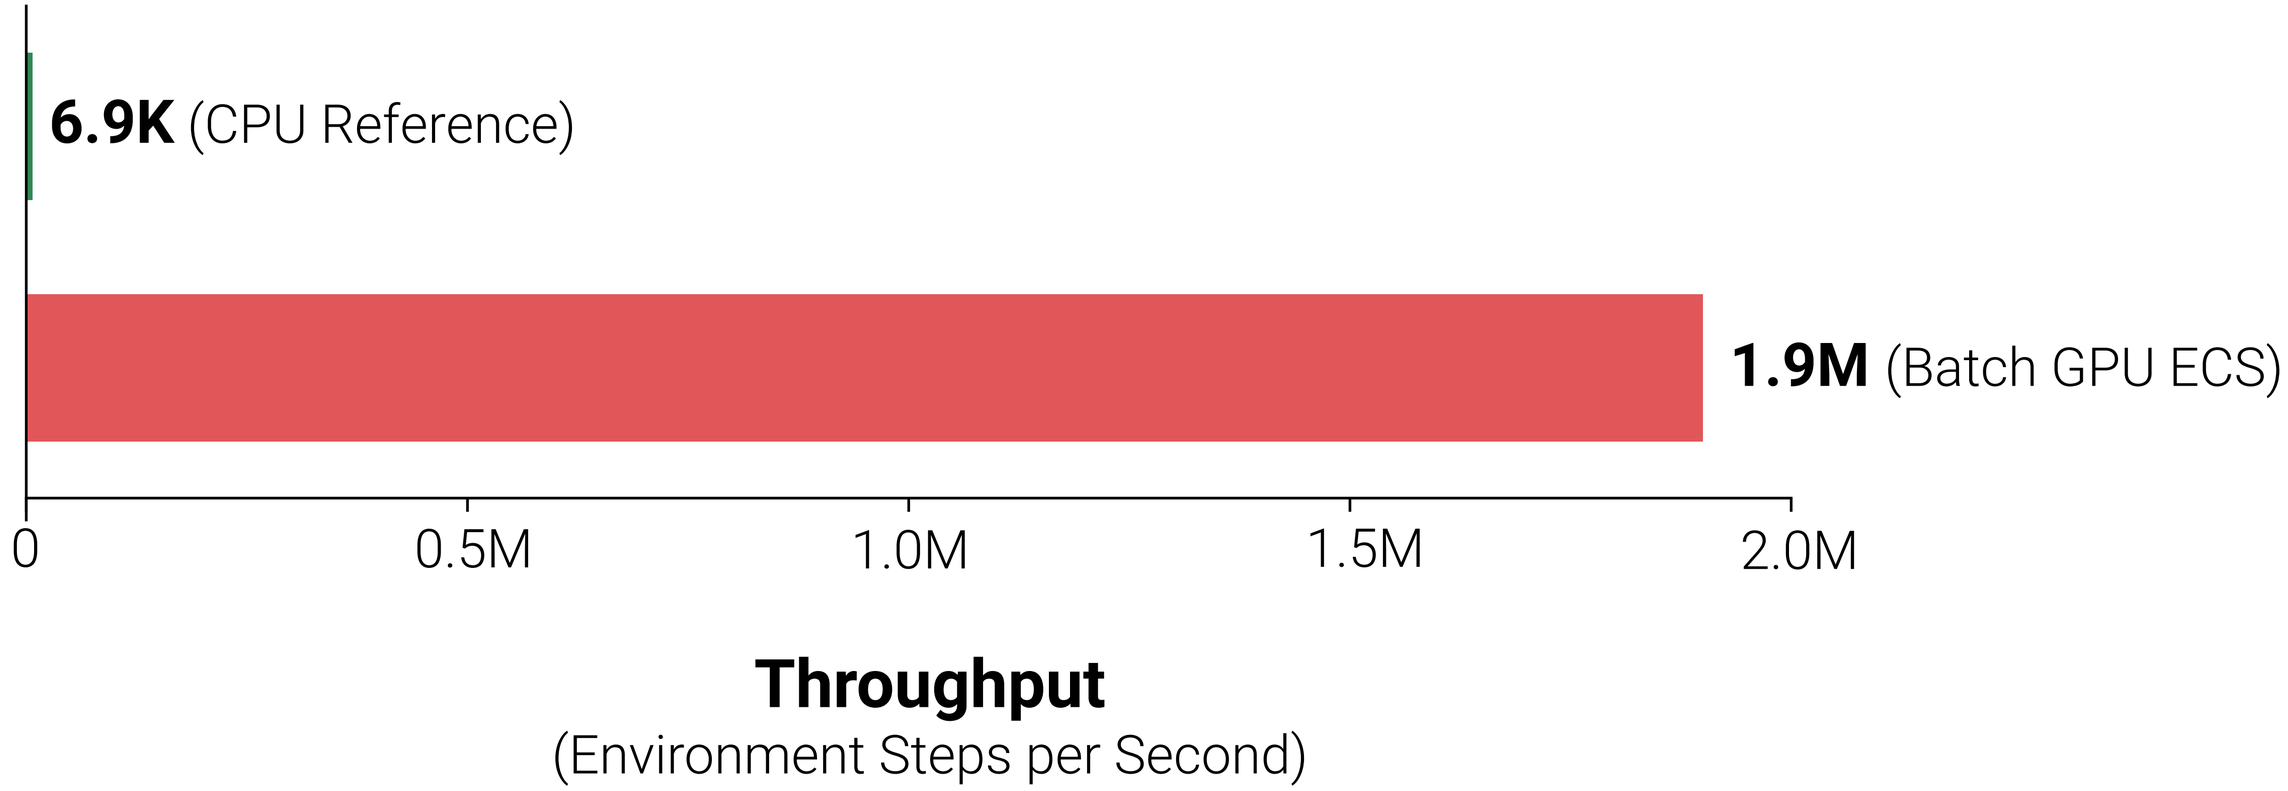

Architecting a High-Performance 3D Batch Renderer

GPUDrive: Multi-Agent High-Level Autonomous Driving Batch Simulator

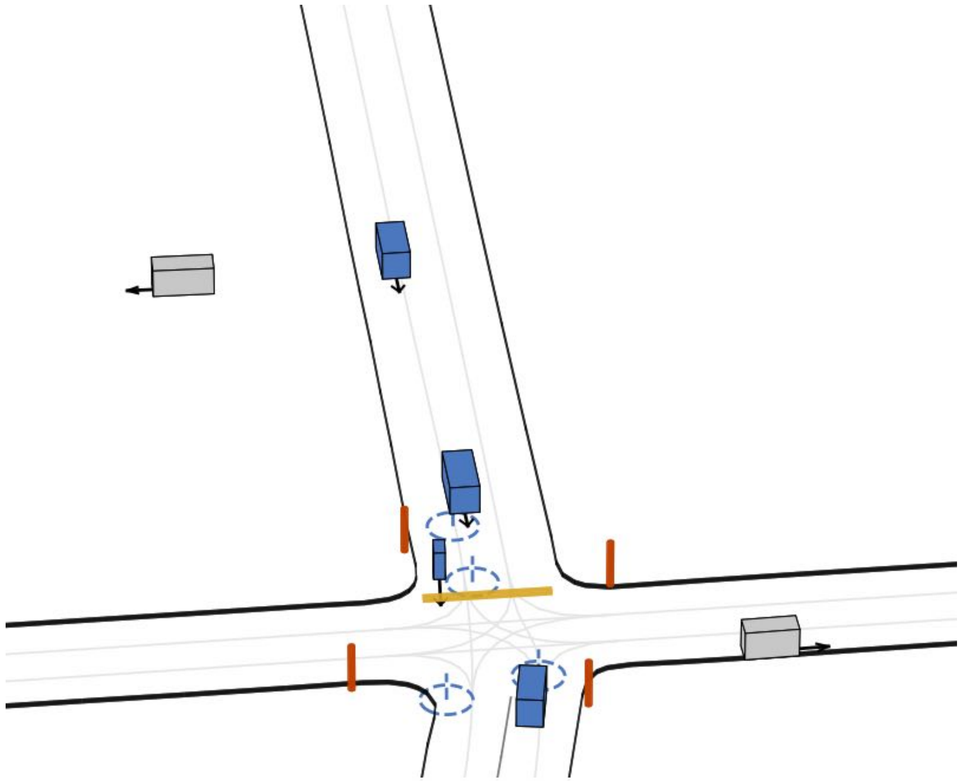

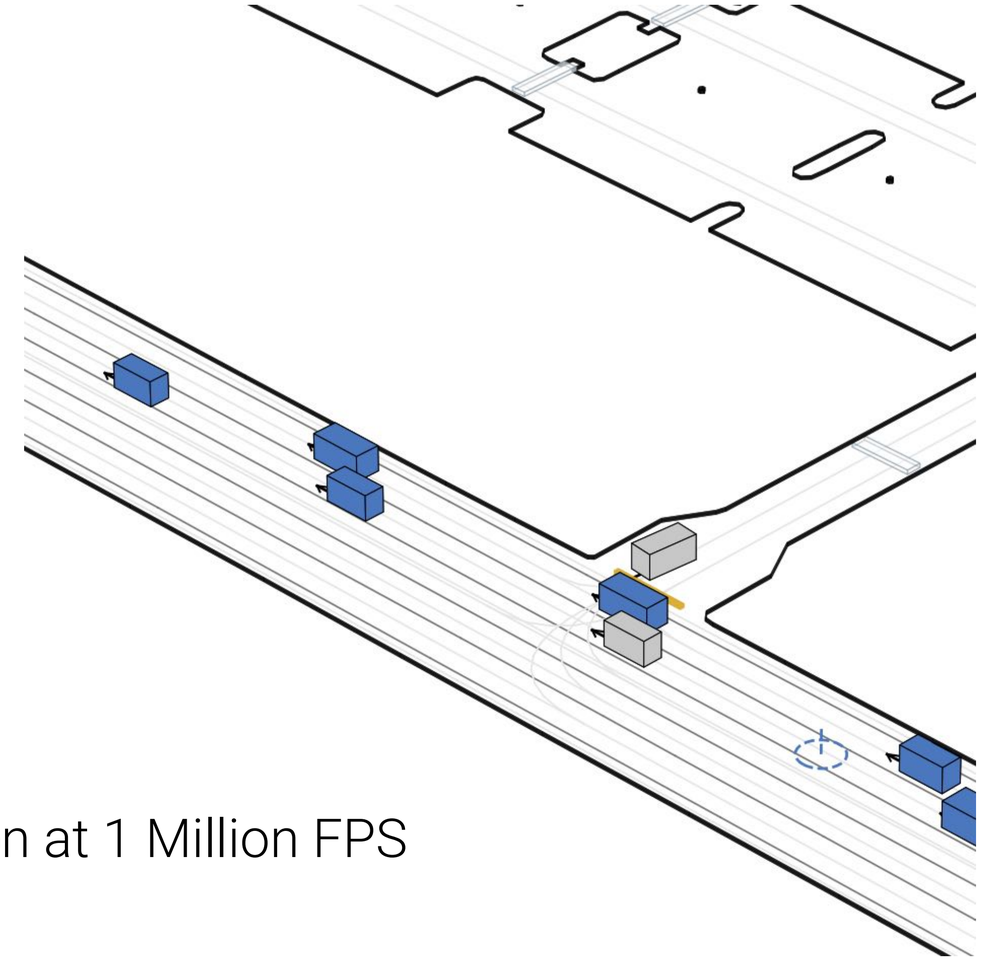

GPUDrive: Data-Driven, Multi-Agent Driving Simulation at 1 Million FPS

Saman Kazemkhani, Aarav Pandya, Daphne Cornelisse,
Brennan Shacklett, Eugene Vinitsky

ICLR 2025

Sim-to-Sim: Policies trained in Madrona can transfer directly to unseen obstacles in Roblox

Maiercik et al. Under submission

Example question: what are the right levels of task abstraction for simulation? And how do we convert between them?

Text-based game

Low-fidelity sim

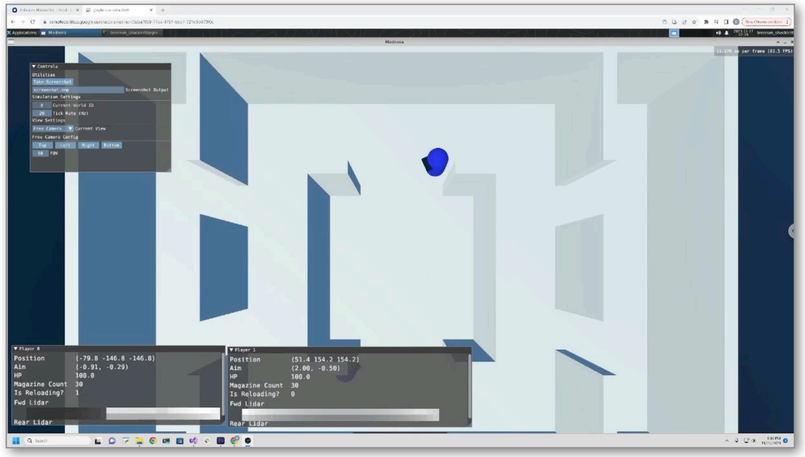

High-fidelity game

LLM policy action:
“You should take cover”

Low-fi agent action:
Move to (x,y)

Game agent policy action:
Game controller input

A fun resource

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

http://kesen.realtimerendering.com/

SIGGRAPH 2024 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 43, Issue 4 (July 2024) Proceedings of ACM SIGGRAPH 2024

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

Accepted Journal Papers

Cyclogenesis: Simulating Hurricanes and Tornadoes
Jorge Alejandro Amador Herrera, Jonathan Klein, Daoming Liu (KAUST), Wojtek Palubicki (Adam Mickiewicz University), Soren Pirk (Christian-Albrechts-University (CAU)), Dominik Michels (KAUST)

Scintilla: Simulating Combustible Vegetation for Wildfires
Andrzej Kokosza (Adam Mickiewicz University), Helge Wrede (Christian-Albrechts-University (CAU)), Daniel Gonzalez Esparza (KAUST), Milosz Makowski (Adam Mickiewicz University), Daoming Liu, Dominik Michels (KAUST), Soren Pirk (Christian-Albrechts-University (CAU)), Wojtek Palubicki (Adam Mickiewicz University)

Interactive Invigoration: Volumetric Modeling of Trees with Strands
Bosheng Li (Purdue University), Nikolas A. Schwarz, Wojtek Palubicki (Adam Mickiewicz University), Soren Pirk (Christian-Albrechts-University (CAU)), Bedrich Benes (Purdue University)

One Noise to Rule Them All: Learning a Unified Model of Spatially-Varying Noise Patterns
Arman Maesumi, Dylan Hu, Krishi Saripalli (Brown University), Vladimir G. Kim, Matthew Fisher (Adobe Research), Soren Pirk (Christian-Albrechts-University (CAU)), Daniel Ritchie (Brown University)

Alignment conditions for NURBS-based design of mixed tension-compression grid shells
Masaaki Miki (The University of Tokyo), Toby Mitchell (Thornton Tomasetti)

Spin-It Faster: Quadrics Solve All Topology Optimization Problems That Depend Only On Mass Moments

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
  - Then peel off and explore your own addition to the 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.

Thanks for being a great class!

Good luck finishing projects.

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