Leveraging AI for Computer Graphics

A Practical Perspective



Dalton Omens — CS 148 — November 2026

My Background





EPIC GAMES

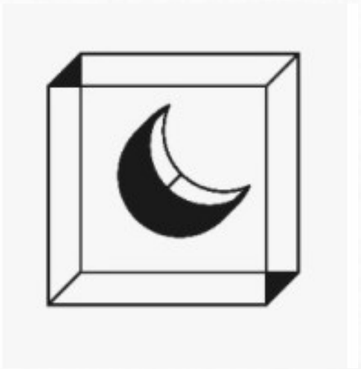

My Background

- Multiple sub-disciplines of graphics and graphics-adjacent work
- Rigging+Animation, 4D Reconstruction, Characters, Creative Tools, Simulation, Generative Modeling, Robotics
- Core principles: Enable artists to realize their vision, and use the best tool for the job (whether it be traditional graphics or genAI)



What are LLMs good at?

(As of November 2026)

- Running sweeps and ablations
- Following specific directions if there is a measurable goal (you have to believe in the goal)
- Creating throw-away visualization UIs
- Suggesting methods and surface-level research for something you have no idea about (you have to tell it to do “deep research”)
- Working overnight
- It seems like, with a lot of time and lots and lots of money, they can solve many problems. But we don’t have OpenAI resources.

What are LLMs bad at?

(As of November 2026)

- Specific Knowledge
- Using a sane amount of money if you’re tokenmaxxing
- Diagnosing what is going wrong in a 3D environment (more on that later)
- Getting the point across
- Being correct, especially if you know what you’re talking about.
    - It’s like reading a reddit comment about something you actually know about and then realizing it’s the same for every other topic you took at face value

How to Learn

- Ask LLMs for sources.
- Ask LLM for most influential results, papers, blogs, companies about a topic
- Argue against it. Play devil’s advocate. Confident incorrectness is rampant, and AI sycophancy will lead you down the wrong tracks.
- If you generate results faster than you can consume and understand them, you will not learn.

Words from the Wise

“It’s like having someone who knows a lot and is … slightly drunk and is sort of throwing out ideas”

— Terrence Tao

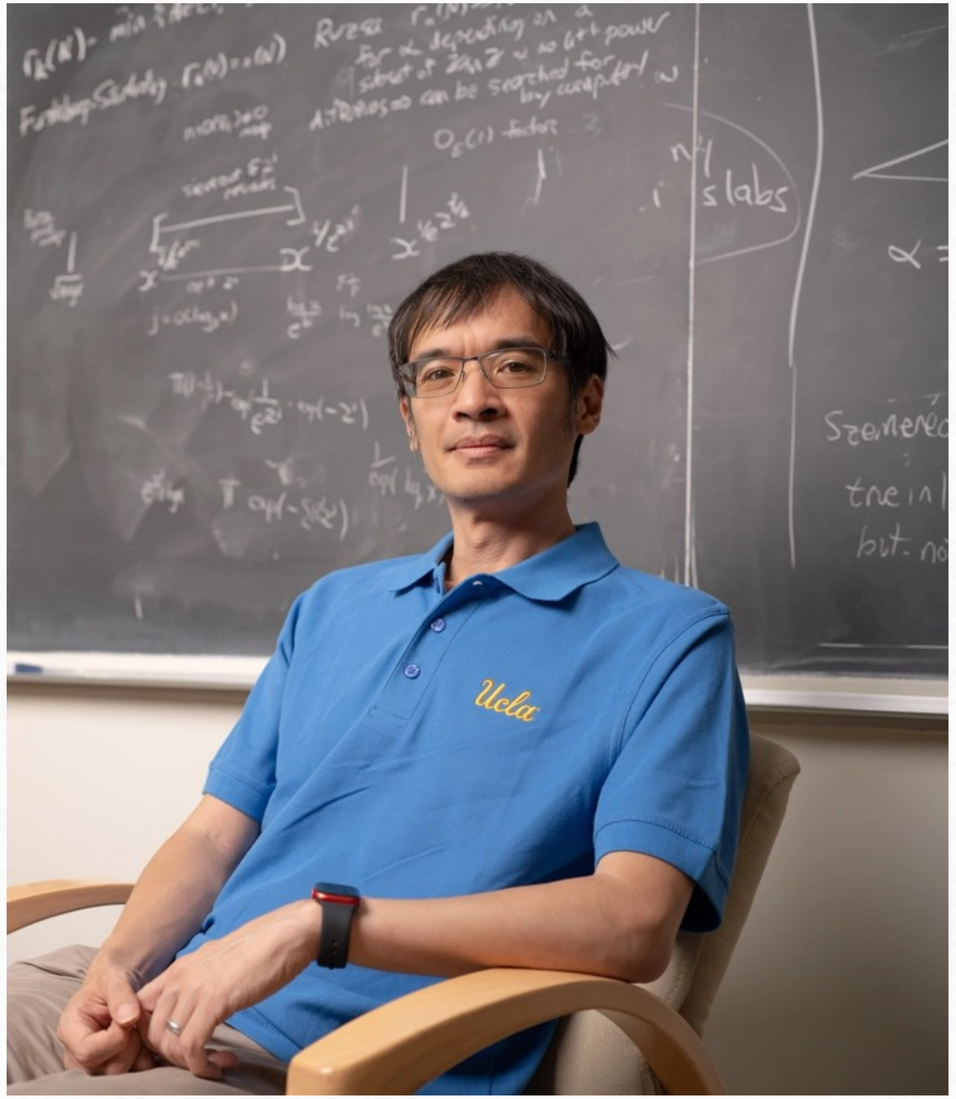

AI Use if a Professional Did This Project

- How has AI use changed how a professional artist would tackle each of the steps of this project?

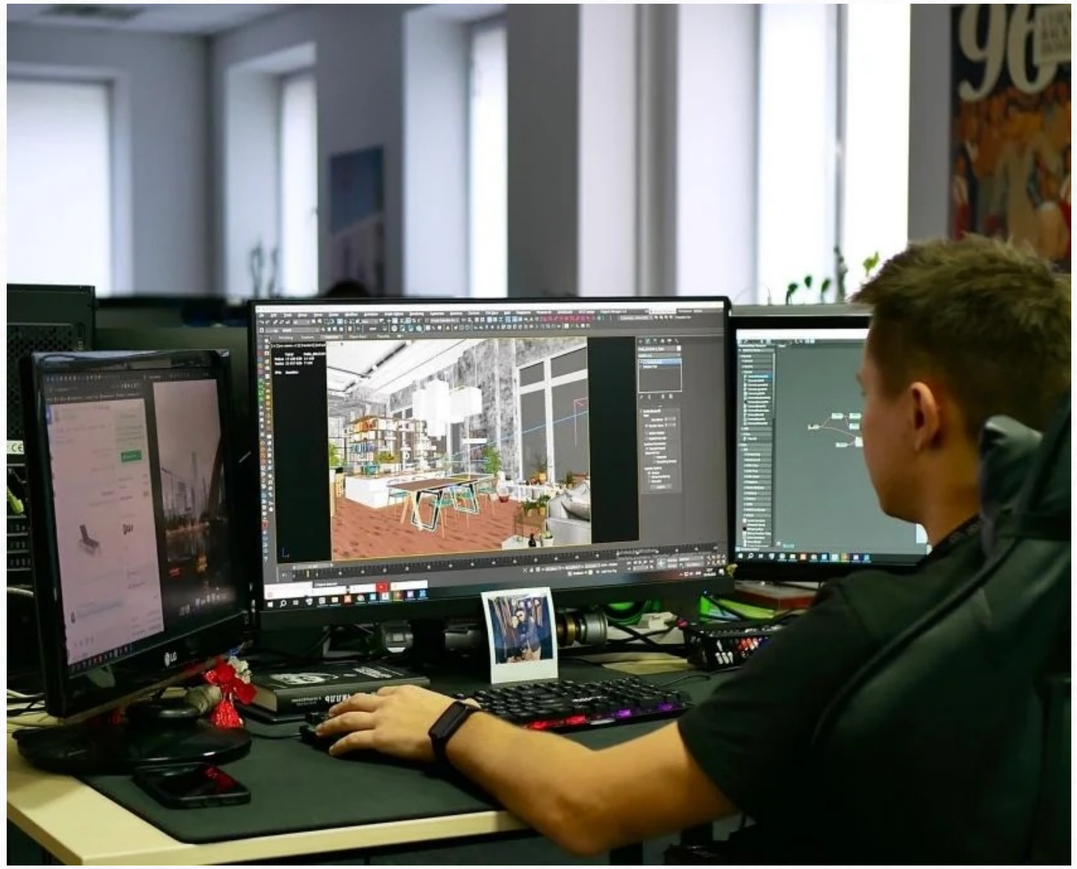

Concepting

- One thing AI has been very good at: generating images
- Concept art can be seen as the “creative soul” of a project.
- This is also the cheapest part of the pipeline





Modeling

- AI mesh generators like Meshy, Tripo, Rodin abound for previz
- Their tessellation and UV mapping is still poor, so currently unusable for deforming objects or texturing. This will improve.
- “Hero” objects more important to model well versus “background” objects
- Code-writing models (LLMs) can also create geometry, especially with rigid articulations.
    - Better for hard-surface modeling, worse for organic modeling



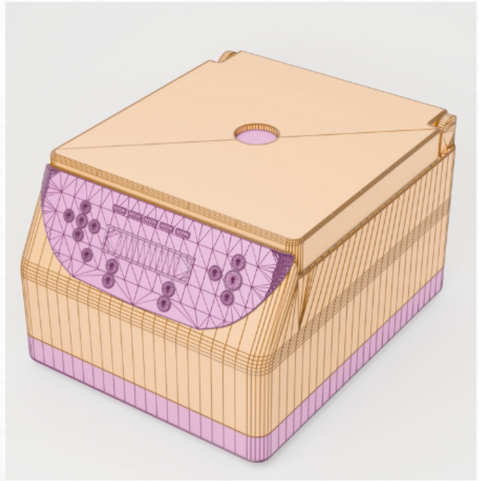

Texturing

- Falls under the category of 2D image generation: mostly accomplishable via AI tools for projects that don’t need to match an exact concept
- “Seamless” textures are important, but existing almost-seamless textures can be modified
- Usually paired with auto UV unwrapping. AI texture generation doesn’t really care about “good” or “clean” UVs
- Hand-painting can still be useful if a really specific look is needed
- Volumetric texturing can be generated via code, e.g. MCP servers connected Blender





Lighting

- AI output is not used heavily here, if you are not replacing the entire rendering pipeline

- AI analysis-by-synthesis to reconstruct a lighting setup from a concept image is still quite poor. This will improve.

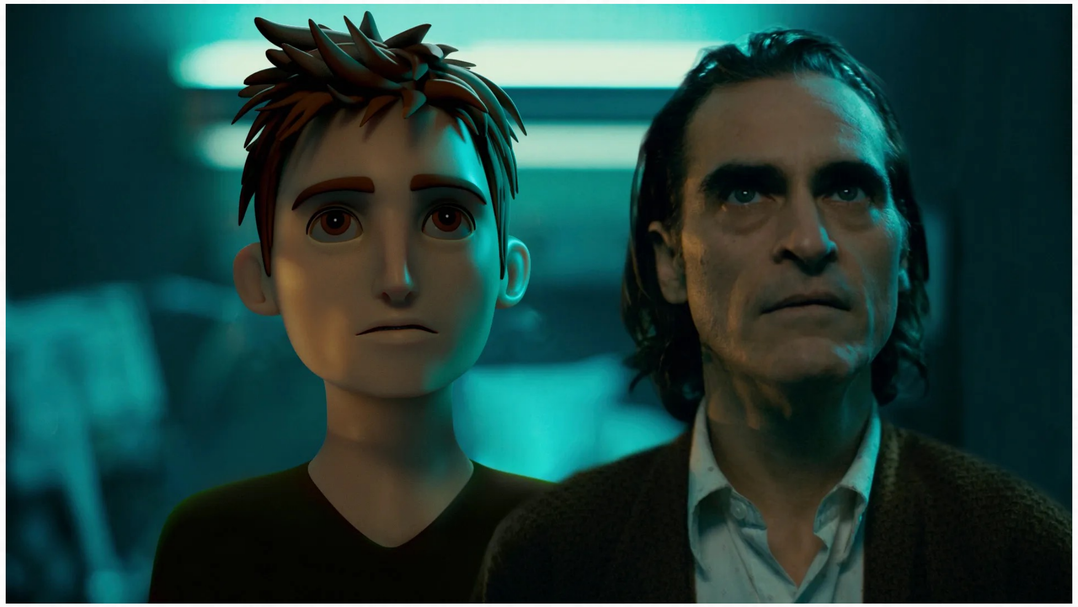

Rendering

- Neural rendering, or AI video models, could theoretically be used to replace all or part of the graphics pipeline
- Can be as light as a re-lighting model like DLSS or as heavy as a World Action Model (WAM)
- The lack of fine-grained control makes these methods unsuitable for matching a specific artistic vision
- Lower-level methods like ML-enhanced light transport are still a domain of research

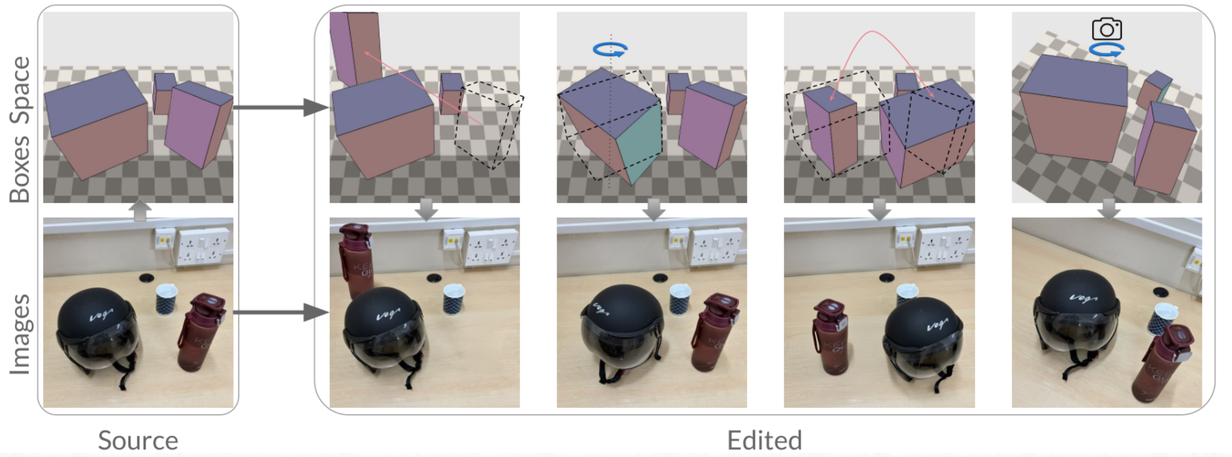



Simulation

- AI coding models are getting better at writing simulations, which previously took a lot of expertise, but still quite bad at it. This will improve.
- Deeply integrating AI components in simulation (to bridge the sim2real gap) is still a research problem.
- AI can set up sims in popular software like blender and Houdini pretty easily now, reducing the expertise barrier for non-high-fidelity results. However, a professional will still not rely entirely on AI.
- AI video can also replace the need for sim, if physical accuracy is not important

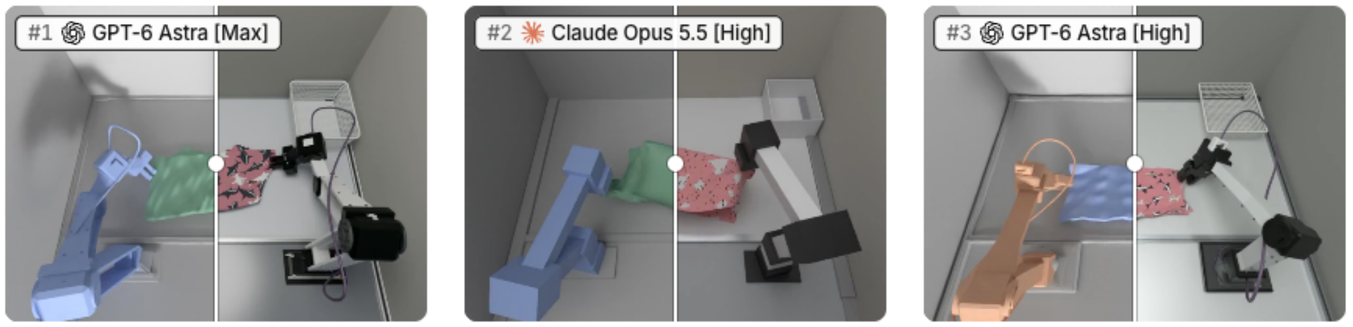

Writing Code

- I don’t know what it’s like to not know how to code, but I don’t write code anymore. I’d imagine it isn’t really that important moving forward (to be a scientist or engineer)

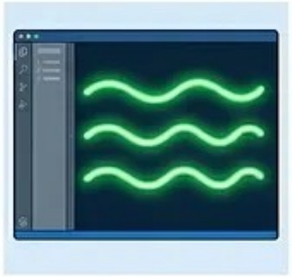

Figure 1



Figure 2



Figure 3



Figure 4



Figure 5



Figure 6



Figure 7



Figure 8

Using AI for Computer Graphics

- DreamWorks rig speedup: Universal Function Approximation
- Small, overfit models can be used to replace computation-heavy operations or add detail from data.
- Synthetic data generation is common

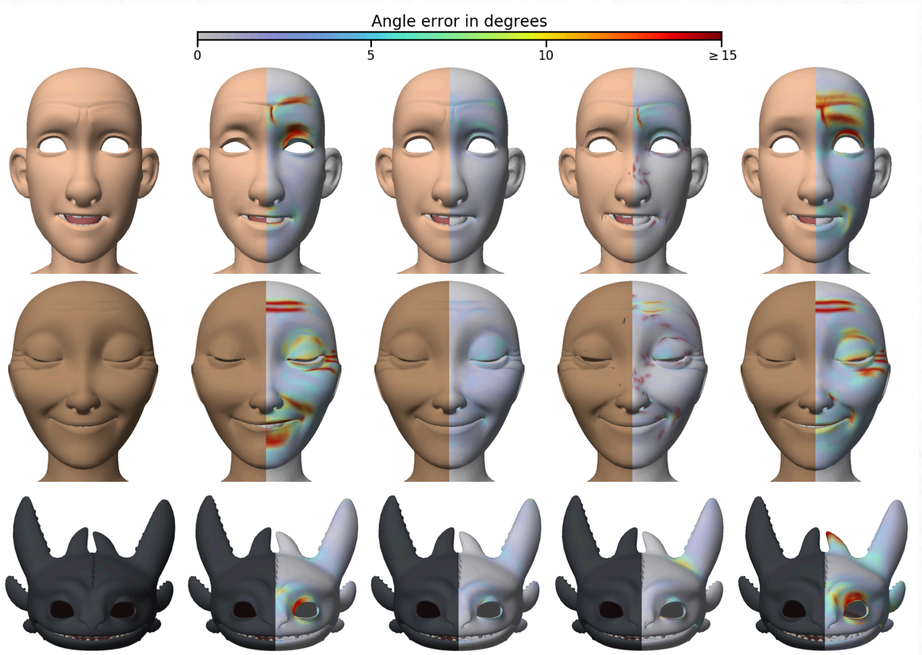

Using AI for Computer Graphics

- Grad School: Feature tracking (dots), neural rigs, depth estimation.
- All small neural networks
- Reconstruction, animation, deformation transfer. Some texturing and lighting too

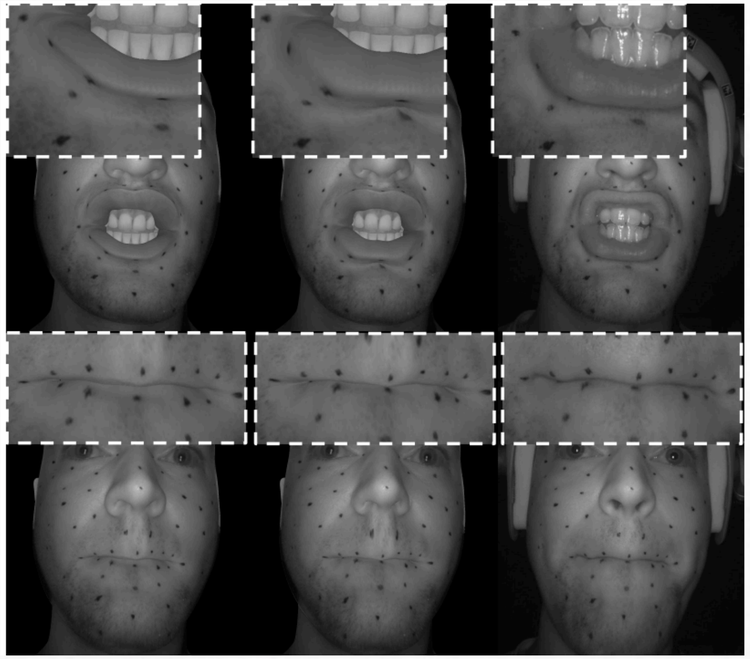

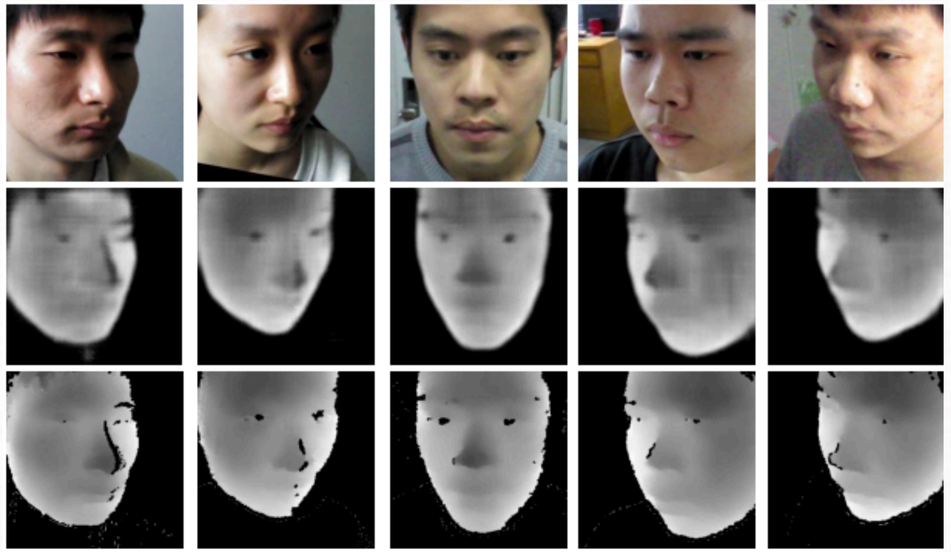

Using AI for Computer Graphics

- LLM-driven AI characters and storylines
- Working with artists and storywriters to shove their preferences into LLMs
- Fine-tuning was possible but required a large data scaling platform
- Agentic flows, right as it was becoming dominant
- As LLMs got better at writing code, they became worse at acting



Aside: At the Mercy of Big AI

- We have to change our entire work paradigm every time a new frontier model is released, e.g. Astra
- When will this stop? Will it accelerate or slow down?
- Try to work on something that isn't on the horizon of being subsumed by Big AI. This is difficult for "coding" work.
- Will "art" become a safe haven, buoyed by the humanists?

Using AI for Computer Graphics

- World-building in the Unreal Engine
- How to make it easier for people to do content creation?
- Tried to get VLAs to understand 3D images — very hard!.
- Giving software MCP for agents to use

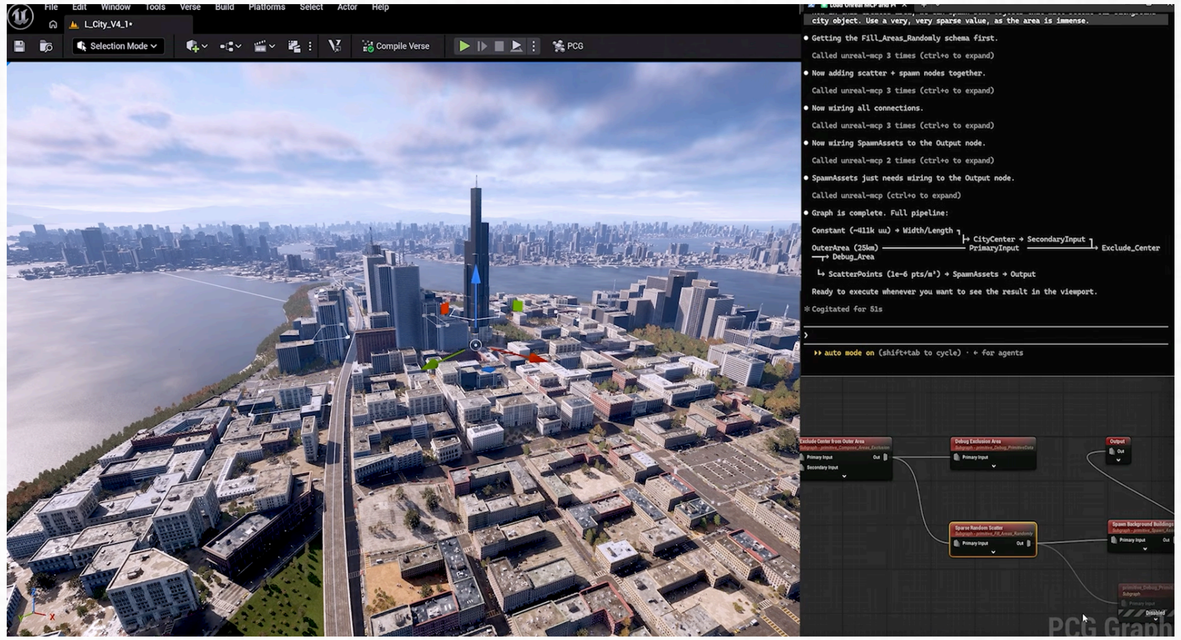

Using AI for Computer Graphics

- Digital Twin Creation for Physical AI
- Vision-Language Models (VLMs) are now good at reconstructing some visual appearance but not good at all at reconstructing physics.
- This will change as physical environments can be used to provide more training data to make VLMs better.
- Lots of LLM harnessing; “riding the wave” of the frontier labs



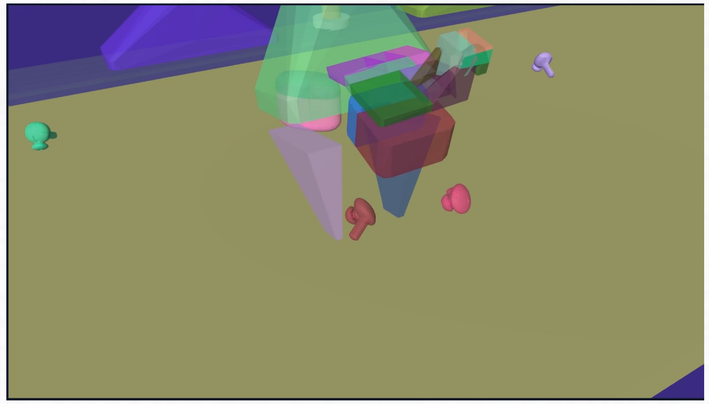

Using AI for Computer Graphics

- Robotics — not exactly graphics, but AI-dominated
- Reinforcement Learning (RL) and Vision-Language-Action (VLA) models abound.
- Interesting new techniques like auto-improving code-as-policy (CaP) and agent-in-the-loop (AITL), leveraging LLMs.
- Modeling, texturing, lighting, rendering from traditional graphics pipeline still in heavy use



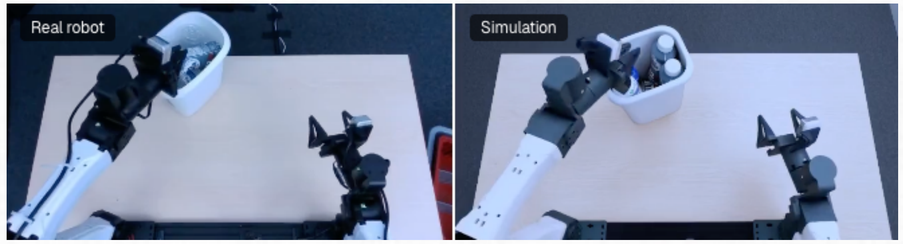

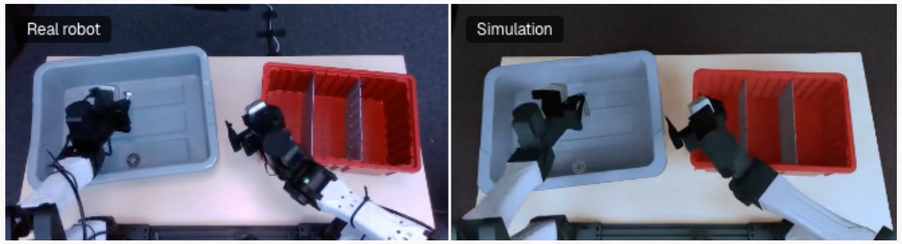

Hiring

- When I hire someone, I no longer really care about how well they can do coding assignments
  - Some people still care about this, but IMO “LeetCode Dead”
- However, I care A LOT about if they have a high-level and deep grasp of the topic and can contribute beyond what an LLM can do
  - If they can’t, I should just increase my token budget
  - I see too many long-form “take-home interviews” where the candidate just prompts and contributes nothing beyond it and doesn’t understand what they did

Takeaways

- AI techbros will scream that the computer graphics pipeline is dead, and that monolithic video and world models are the future
- Controllability remains a central concern
- In practice, due to the controllability issue, AI for 3D has not taken over. But this could change rapidly; it is mostly at the whims of Big AI right now.

Takeaways

- Critical thinking is more important than ever.
- Prompting is the way of life now. Unfortunate for those who liked solving the mini-puzzles of software engineering but overall better for freeing up focus for higher-level tasks and ideas
- Computer Graphics as a tool for art is very much alive (related to the “control issue” from the prior slide)
- Using AI for a 3D CG output still gets you “good” but not “great”
- Learn how to make art! It’s good for the soul!

Q&A