

Talking to Smart AI NPCs in Unreal Engine 5: The Future of Gaming & Artificial Intelligence

NVIDIA ACE for Games

Spark Life Into Virtual Characters
With Generative AI



It's alive! The rise of generative non-player characters in video games

BY MARTIN SAS AND VICTORIA HENDRICKX - 10 MAY 2024

AI in Games

Generative AI is transforming games, moving from static experiences to "living games" that dynamically adapt and create content in real-time based on player interactions.



Jack Buser

Director for Games, Google Cloud

**Generative AI Could Create the
Video Games of Our Dreams**

**How I Built an LLM-Based
Game from Scratch**

One Trillion and One Nights

An experiment using LLMs to procedurally generate browser-based JRPGs

13 min read · Jan 22, 2025

Terms / Context

Human :)


Large Language
Models (LLMs)???

AI???


Generative AI
(GenAI)???

Chatbots???

Terms / Context



Terms / Context



Generative AI

Transformational Technology

The world will operate much differently in ways we may not be able to predict

Enabling Technology

Make humans more efficient

Create starting drafts of text or text-based code

Creating simple artwork (pictures, music, etc.)


Disruptive Technology

Replaces jobs

More successful at customer service in one service center

Predicting Human Behavior

- Predict future events for individuals (e.g. future arrests, age of death, future loan default)
- Marketed as objective and fair
- Algorithms are nothing more than **opinions embedded in code**




What is a predictive algorithm?

-




Characteristics of Mathematics / Computers

- Facts
- Objective
- Truth
- Logic



Weapons of Math Destruction

How Big Data
Increases Inequality and Threatens Democracy





Weapon of Math Destruction

Widespread

- Affects a large *number* of people
- The decisions *affect their lives substantially*

Mysterious

- The scoring algorithms' formulas are not available
- People may not even know they are being scored

Destructive

- Individual: They ruin people's lives unfairly
- Systemic: They make the original problem worse



Weapon of Math Destruction

Widespread

- Affects a large *number* of people
- The decisions *affect their lives substantially*

Mysterious

- The scoring algorithms' formulas are not available
- People may not even know they are being scored

Destructive

- Individual: They ruin people's lives unfairly
- Systemic: They make the original problem worse


Why aren't algorithms objective?

(Example: Teacher Assessment)

Two-decade war on teachers

- Must be the teachers' fault
- Let's find the bad teachers, get rid of them

- | |
|-------------|
| • Facts |
| • Objective |
| • Truth |
| • Logic |



1st Gen: Student Scores

- Look for teachers teaching large number of students who did not pass standardized tests.
- **What limitations do you see?**

Limitations of “Student Scores”

- **Assumption:** Score is entirely caused by teacher quality / performance.
- But poverty is correlated with performance.
- Therefore, you target teachers of students in poverty, not “bad” teachers.

2nd Gen: “Value Added”

- Don't blame teachers for students who start out behind
- Let's make a model for the expected score for each student in the class
- Teacher held accountable for **difference** between expected and actual score
- **What limitations do you see?**


Limitations with “Value Added”

- If a teacher taught 20,000 students, this might work.
- Class of 35: Too small of a sample size
- It is a noisy model – test scores affected by temperature, time of day, medication, etc.
- Kids in poverty forget more in the summer.
- If a teacher cheats and makes student scores higher one year, the next teacher is punished.
- **Scores being used to deny teachers tenure**


2nd Gen: Assessing Model

- How can we ***assess*** this model?
- Teachers who teach both 7th and 8th grade math had two scores.
- If this is a good model, then the same teacher should get a similar score in both classes.

Good Model Results



Actual Model Results





Weapon of Math Destruction

Widespread

- used in many urban districts nationwide

Mysterious

- algorithm is proprietary, FOI request was denied

Destructive


- Wrong individuals fired
- Urban national teacher shortage: good teachers fired, quitting, retiring, not going into teaching because of perceived arbitrariness of evaluation

Using data / computer provides a façade of objectivity behind which bias hides

More Examples

- Hiring process
- Parole system
- Life insurance
- Auto insurance
- Qualifying for a mortgage
- Interest rates on loans
- Voting availability
- Targeted marketing / misinformation
- ...

Terms / Context



What are (large) language models?

- AI systems that are trained on large (massive!) datasets of text (e.g., books, articles, websites)
- They “learn” existing patterns in their training datasets and “predict” the most likely next word in a sequence.

The language model problem...

Rank these sentences in the order of plausibility?

1. Mandy went to the store.
2. store to Mandy went the.
3. Mandy went store.
4. Mandy goed to the store.
5. The store went to Mandy.
6. The food truck went to Mandy.


How **probable** is a piece of text? Or what is **$p(\text{text})$** ?



Mandy says Hiiii!!!

The language model problem...

A language model answers the question: **What is $p(\text{text})$?**



Language models in the news (these days, ChatGPT)

FL

Hi, I'm writing an article about you and all of the cool things you can do. Could you say hello to my readers and give us a quick self-introduction? I'd like it to be friendly and casual, include a silly joke, and end with a note about how great the rest of my article is.

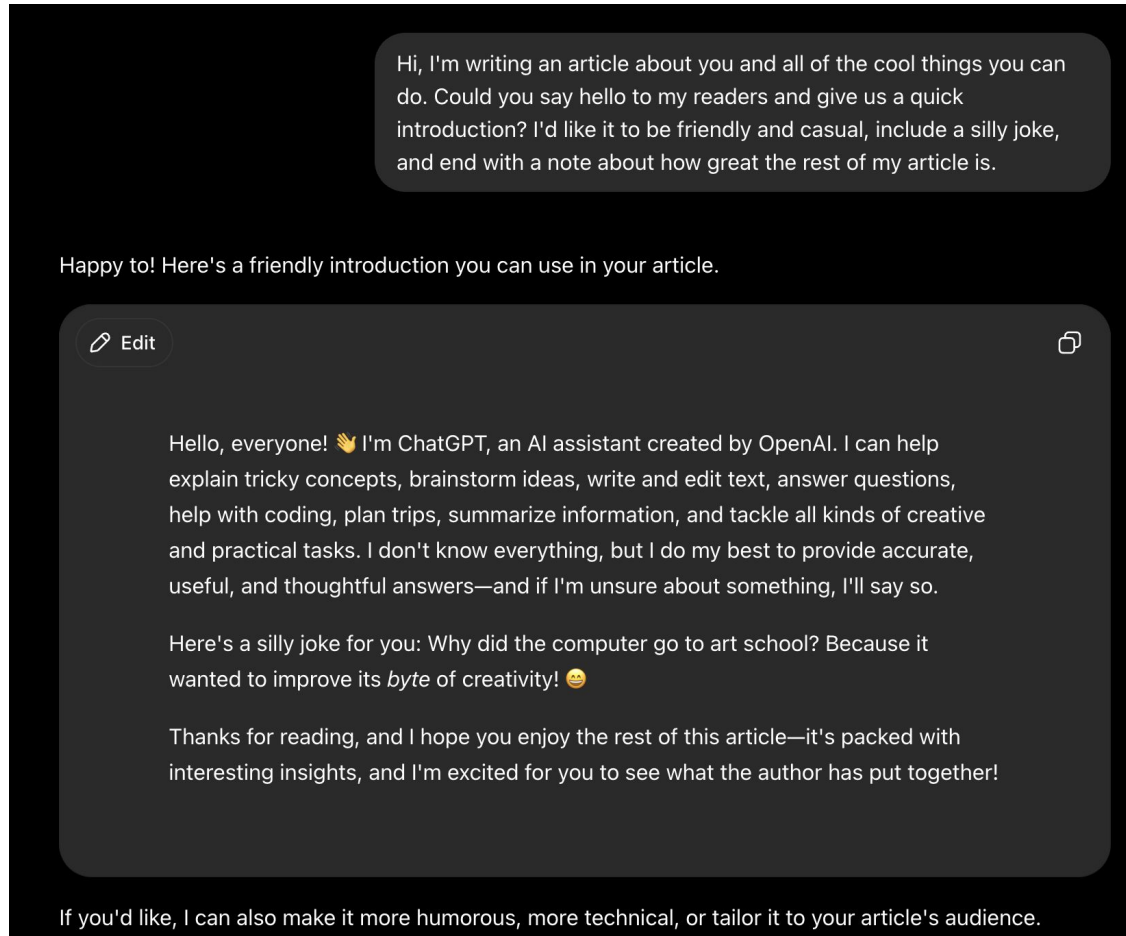


Hello readers! My name is Assistant, and I'm a large language model trained by OpenAI. I'm here to help answer your questions and provide you with information on a wide variety of topics. I'm a bit of a jokester, so here's a quick one for you: Why was the math book sad? Because it had too many problems! But seriously, I'm excited to be featured in your article and I hope you enjoy learning more about me and all the cool things I can do. And trust me, the rest of the article is going to be great!






Past image taken from Springboard

Language models in the news (these days, ChatGPT)



More recent image, taken on July 20, 2026

We use language models everyday!




AI in Games

AI in video games includes the computational systems that:

- simulate strategic decision-making
- generate dynamic game behavior
- control non-player characters (NPCs)

AI has a long history in game development!



AI HAS A LONG HISTORY OF BEING “THE NEXT BIG THING” ...



Timeline of AI Development

- **1950s-1960s:** First AI boom - the age of reasoning, prototype AI developed
- **1970s:** AI winter I
- **1980s-1990s:** Second AI boom: the age of Knowledge representation (appearance of expert systems capable of reproducing human decision-making)
- **1990s:** AI winter II
- **1997:** Deep Blue beats Gary Kasparov
- **2006:** University of Toronto develops Deep Learning
- **2011:** IBM's Watson won Jeopardy
- **2016:** Go software based on Deep Learning beats world's champions

AI in Games (1950 – 1980)



A B C nim-sum


3 4 5 $010_2=2_{10}$ I take 2 from A, leaving a sum of 000, so I will win.
 1 4 5 $000_2=0_{10}$ You take 2 from C
 1 4 3 $110_2=6_{10}$ I take 2 from B
 1 2 3 $000_2=0_{10}$ You take 1 from C
 1 2 2 $001_2=1_{10}$ I take 1 from A
 0 2 2 $000_2=0_{10}$ You take 1 from C
 0 2 1 $011_2=3_{10}$ The normal play strategy would be to take 1 from B, leaving an even number (2) heaps of size 1. For misère play, I take the entire B heap, to leave an odd number (1) of heaps of size 1.
 0 0 1 $001_2=1_{10}$ Y

Logic-based approaches (1950s)




Move	Value
Kc6	Win in 5
Qa6+	Win in 5
Qc6+	Win in 8
Qg6	Win in 8
Qa5+	Win in 8
Qc5	Win in 9
Ke5	Win in 10
Kd4	Win in 10
Qg1	Win in 10
Ke6	Win in 11
Qf2	Win in 13
Ke4	Win in 14
Qd4	Win in 16
Kc5	Draw
Qxb5	Draw
Qe6	Lose in 28
Qf6	Lose in 15
Ob6	Lose in 15

☒ White to move
☐ Black to move

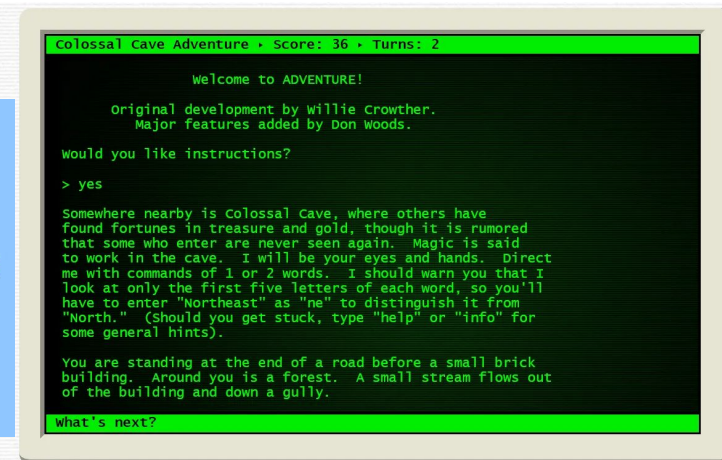



Selective search on data trees (1960s)

AI in Games (1950 – 1980)





Arcade: Movement patterns (1970s)




AI in Games (1950 – 1980)




AI patterns (tracking + following)
(late 1970s - early 1980s.)



AI in Games (1980 – 2000)



Early Pathfinding:
Dijkstra's algorithm (1981)



Finite State Machines (early 90's)

AI in Games (1980 – 2000)

AI “Stupidity”: originally meant to constraining AI capabilities to adjust difficulty levels. This might inadvertently

- repetitive behaviors
- abnormal actions in situation unforeseen by developers
- reduced player immersion

“Cheating” AI: AI opponents often cheating (e.g., AI’s actions sometimes seem overly smart)

- Notable in early versions of the Civ

Deep Blue vs. Kasparov



Deep Blue
IBM chess computer

Garry Kasparov
World Chess Champion

First match

- February 10–17, 1996: held in [Philadelphia, Pennsylvania](#)
- Result: **Kasparov**–Deep Blue (4–2)
- Record set: First computer program to defeat a world champion in a *classical* game under tournament regulations


Second match (rematch)

- May 3–11, 1997: held in [New York City, New York](#)
- Result: **Deep Blue**–Kasparov (3½–2½)
- Record set: First computer program to defeat a world champion in a *match* under tournament regulations

AI in Games (2000 – pre-LLMs)


An AI simulation where users could hatch norns (small furry animals), teach them how to behave, or leave them to learn on their own.

Each norn has a several-layer **neural network** representing its brain.




Neural Networks: First successful application after 2nd AI Winter

AI in Games (2000 – pre-LLMs)



Navmeshes: 3D Pathfinding (2000)




Noun Phrases (NP): "the burglar"
Verb Phrases (VP): "robbed the apartment"
Sentences (S): "the burglar robbed the apartment"




Symbolic NLP and Believable Characters (2005)

AI in Games (2000 – pre-LLMs)




Behavior Trees and
Procedural Generation (2008)



Minecraft wizard, and record holder for the *Super Mario World* speedrun [SethBling] is experimenting with machine learning. He built a program that will get Mario through an entire level of *Super Mario World* – Donut Plains 1 – using neural networks and genetic algorithms.

Reinforcement Learning (2015)

The dawn of transformers for Language Modelling (2017-2020s)



Breakout #1 – AI Uses in Games

Group with your neighbor(s) and **pick a game that everyone knows**

Where could you **use AI** in that game?

Where could you **specifically use GenAI**?

GenAI in Video Games

- **Generate game content** (e.g., quests texts, descriptions, texts, character backgrounds)
- **Procedural generation** (automatically generate levels, environments, quests – e.g., MarioGPT)
- **Adaptive game environments** (e.g., dynamic changes in weather, terrain)
- **Localization** (e.g., quickly generate localized content in different languages, cultures, player preferences)

GenAI in Video Games: Smart NPCs

Traditional NPCs were scripted with pre-defined dialogues and behaviors.

Smart NPCs, using LLMs, can respond dynamically to the player's choices.

They could even have their own personality, ethics, and long-term memory!



Breakout #2 – Smart NPCs

First discussion:

- What are some **benefits** of smart NPCs?
- What are some **risks** associated with smart NPCs?

Second discussion:

Now, think of a game you've played.

Why or why not would AI NPCs work in that game?

Smart NPCs

Benefits

- Emergence and replay ability
- More dynamic answers / hints (not struggling if need helps)
- More personal storyline
- Characters can react to more things
- Keep players more concentrated / in flow state
- Make player more emotionally invested
- NPCs have more personalities, make the game more fun

Risks

- Too emotionally invested / reliant on smart NPCs
- Demanding for the hardware
- The player can manipulate the NPCs
- Players may ask for too many hints → lose fun long term
- AI can go off script and break the immersion of the game
- Possibilities of errors in generated texts (e.g., grammars)
- Too many AI bots, less satisfying gameplay experience

Smart NPCs

Benefits

- Boost player immersion (more realistic and emotionally engaging interactions)
- Personalization (custom quests, difficulty, rewards, etc.)
- Scalability or less repetitive experiences (automatically create new dialogues, quests, events, etc.)

Risks – They’re unpredictable!

- AI NPCs could say stuff that accidentally break the game narrative or rules
- AI NPCs could be rude, biased, offensive
- AI NPCs could support player cheating (they could be tricked to spill secrets)
- Some players could have “relationships” with AI NPCs instead of other human players

Fun Video :)



Youtube: I Tried to Convince Intelligent AI NPCs They are Living in a Simulation

Breakout #3 – Game Programming

LLMs have changed the ways people code!

What are some **benefits** of using LLMs for game programming?

What are some **risks** of using LLMs for game programming?

Using LLMs for Game Programming

Benefits

- Faster
- Help programmers find personalized solutions faster
- Make game programming more accessible

Risks

- If something goes wrong, it can be hard to debug
- Hard to identify inefficient code (no immediate error)
- Over-reliant
- Hinder learning for novice programmers

Using LLMs for Game Programming

Benefits

Automatically generating assets, code, content can:

- Increase process efficiency
- Provide endless resources
- Increase profits???

Risks

- Ethical and copyright concerns
- Limit career growth for developers (less skillful people could lose their job)
- Lower game quality (if using LLMs in irresponsible ways)
- Hidden costs???


HIDDEN COSTS OF DEPLOYING LLMS

DISCOVERING LLMS



ENDLESS
OPPORTUNITIES
BRING JOY AND
CURIOSITY!



UH-OH!
LLM COSTS
A LOT



WHAT IF I CAN
COMPRESS
LLM MODELS?



COMPRESSION
PITFALLS

RISING BIAS,
INCREASING
LATENCY




“Vibe Coding”

Creating a software with minimal effort, where the developer:


- Describes a project or task to an LLM and asks it to generate code
- Doesn't review, test, or understand all the code generated

ROSEDALE


More People Are Searching
'Vibe Coding' Than
'Learn to Code'



“Vibe Coding a Video Game” Trend

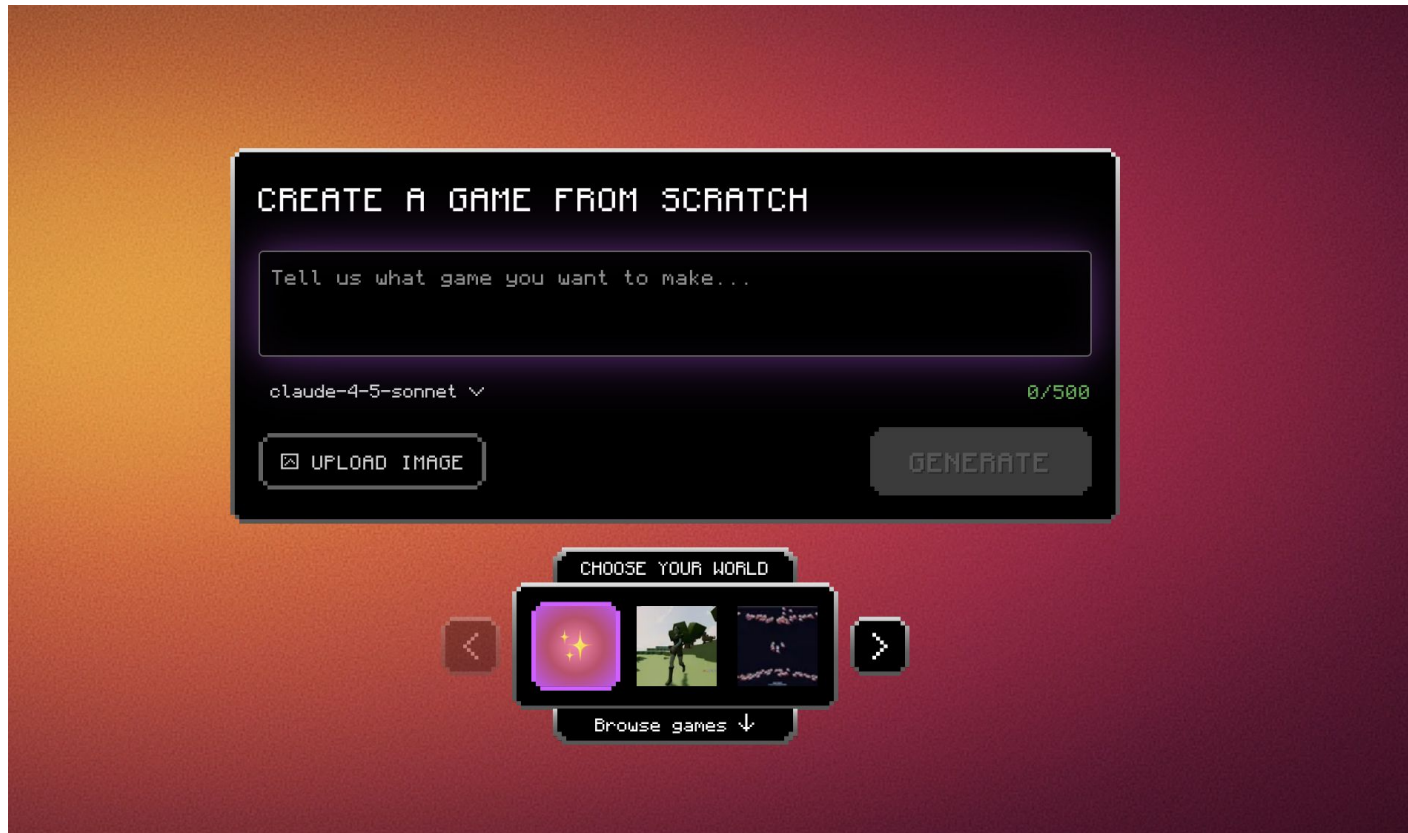


Discussion: What are some issues associated with this trend?



Breakout #4

Let's try something fun :)



<https://rosebud.ai/>