

LLMs have become synonymous with chatbots. And that sucks. Many people only experience LLMs as conversational agents like ChatGPT and Claude (and all the other places that chatbots are now jammed into).


Early adopters thought players would want to have open-ended natural language conversation with NPCs, or with the game itself, but that ...didn't happen. It turns out a chatbot that generates a wall of dialog is **exhausting**. Chatbots burn up both computational resources and the player's interest.

But LLMs aren't chatbots. There are other **small, hidden, useful things LLMs can do in games**. And fortunately, generating things other than dialogue can take fewer tokens and use less-advanced models.

I'll show you 3 broad categories in this zine, but I think there are many more waiting to be discovered. We just need to look (and stop making chatbots)

1. ARBITRARILY-FLEXIBLE SELECTION LOGIC FOR NORMAL-ASS GAME CONTENT

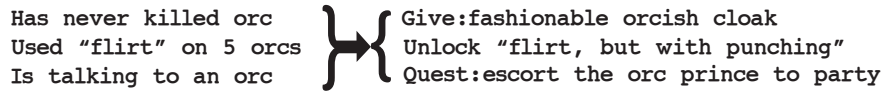
So you have a game with some normal-ass game content. You hand-author what in-game content should be triggered on player actions, for example, when a player picks a dialog choice:



Quests:
kill 5 orcs
escort orc prince
steal orc sword
defeat orc mage
...

But players do **PLAYERCHAOS**. They go off script, do quests out of order, kill the wrong NPCs, do weird stuff. We want to respond in a way that hears them. But we don't want to write infinite new content, or write infinite conditionals. We want to respond to **PLAYERCHAOS....within** existing content and structure.

Solution: Give the LLM the player's actions, and a list of in-game content options, and ask it for its selection.



Bonus: this can respond arbitrary things we couldnt before: natural language and avatar movement and mouse tracking and eye movement and.... and respond to all of them at once (called "sensor fusion")

Does the response make sense? Do the players understand that it's based on their actions....? Maybe. You can use more LLM to explain the logic ("**the orc king asks flirtatiously, he knows your past**")


2. STRUCTURE TO STRUCTURE, JSON TO JSON

What if we want to handle **PLAYERCHAOS**, but not write content for everything the player can do? There's a pattern:

AI DUNGEON 2 used LLMs to respond to natural player text input. But it quickly lost shape and logic. "I defeat the dragon and turn into pure logic" "ok, you turn into logic"

HIDDEN DOOR used LLMs to respond to a mix of **natural language** and **UI** input. It responded with both narrative and updates to the games logic, adding new characters, inventory, and locations to a persistent game state.

Hidden Door's approach is to mix natural language (for flexibility) with structured world data (for stability and logic and UI display). To do this, feed the LLM the world state, and the actions, and **ask it for a changelist** (which you can sanitize)




If LLMs can generate structured objects, like game state, what other content in a game is structured? Weapons? Quests? Items? Relationships?

INFINITE CRAFT
(Neal Agarwal, 2024)

This game is my favorite JSON-to-JSON example, and favorite LLM game

When you combine 2 items, it checks a shared data table of existing combinations.



If it doesnt find it, it give you a reward for being first, and makes a tiiniy LLM call, of just 6 examples and the two items, and request the tiiiiiny token answer of another emoji and name.

Phenomenal cosmic exploration, ...itty bitty token usage!

3. MAKE THE HUMAN DO ALL THE (CREATIVE) WORK

LLM calls are expensive and ethically questionable. You know where you can get better free content?

Bored humans playing your game!

In party games like Pictionary and Charades, the **fun of the game is in performing** for your friends.

But games like this only work because there is an **in-game motivation** (be guessed, get points) for performing.

A few successful LLM games have noticed that you can separate the “judging” from the performing and **have an AI “judge”** motivate the performance.

SUCK UP! makes you (a vampire) argue with an LLM to get invited inside. Is it performance? This game is a Twitch breakout, its more fun to perform for others, even in single player.

DEATH BY AI asks each human to present your plan for escaping a giant (Little Umbrella, 2023) chicken or sudden blizzard. But, we think it chooses randomly, for fairness, and justifies it afterwards.

1001 NIGHTS The human playing as Scheherazade types a story to the (Ada Eden, 2022) LLM king, who responds with more story. Your goal is to get him to mention weapons, which become real.

HEY ROBOT Does the game-creator need to host the LLM? This game (Everybody House, 2019) uses whatever source of stochastic spam (Alexa, etc) you have. You win by guiding it to say particular words.

Technically these are chatbots, but their job is to shut up and get the humans to talk more. This is AI judo - the human provides the energy and content, and the AI can get away with almost no content of their own.

Notice that you don't want much feedback, or even reasonable feedback. Players want to feel witnessed, but not analyzed. You can even mix in **randomly-generated judgement!**



"haha I'm only dancing to get us points"

This lets players **blamewash** their silliness.

Icons from Freepik, Font is Labours/Baloo

© COPYRIGHT 2026, KATE COMPTON, CREATIVE COMMONS

LOW-TOKEN USES OF


LLM*s

IN

GAMES

THAT AREN'T A

F***ING CHATBOT



LLMs can be a thrifty gameplay gluedon't keep them in chatbot hell!

*"Large language models"**, a form of statistical text completion most commonly seen in tools like ChatGPT

** pronounced "largle langle mangles"
(Johnathan Pagnutti)

GALAXYKATE @BLUESKY
@GMAIL
.COM