Game Agents

CS6960 MultiModal LLM Agents

Kenneth Marino

Announcements

- Projects
  - Feedback Returned
  - Reminder: you can gain back points on this by updating the relevant sections when you submit milestone 1
  - Project Milestone 1 Due Friday
  - See instructions
  - Show some progress

Any Questions

Game Agents

- Also kind of a bucket term
  - It just means an agent specifically for something we normally call a game
  - Generally, the goal here is to get an LLM to play some specific game really well

Can AI Play Pokemon?

Can AI Play Pokemon?

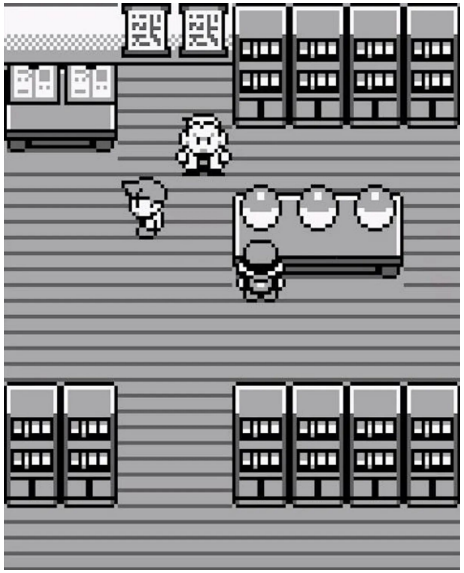

- Strategy game
  - Choose which pokemon to use
  - Choose what moves to use to defeat opponents
  - Decide where to go next
  - Use items in the overworld
- Extremely long-term decision making
  - In the first few minutes you make a choice that affects the entire rest of the run (which starter)
  - Beat 8 gym leaders, Team Rocket, collect HMs to progress, beat the elite four

Yes (kind of)

“We equipped Claude with basic memory, screen pixel input, and function calls to press buttons and navigate around the screen, allowing it to play Pokémon continuously beyond its usual context limits, sustaining gameplay through tens of thousands of interactions. In the graph below, we’ve plotted the Pokémon progression of Claude 3.7 Sonnet alongside that of previous versions of Claude Sonnet, which didn’t have the option for extended thinking. As you can see, the previous versions became stuck very early in the game, with Claude 3.0 Sonnet failing to even leave the house in Pallet Town where the story begins. But Claude 3.7 Sonnet’s improved agentic capabilities helped it advance much further, successfully battling three Pokémon Gym Leaders (the game’s bosses) and winning their Badges. Claude 3.7 Sonnet is super effective at trying multiple strategies and questioning previous assumptions, which allow it to improve its own capabilities as it progresses.”

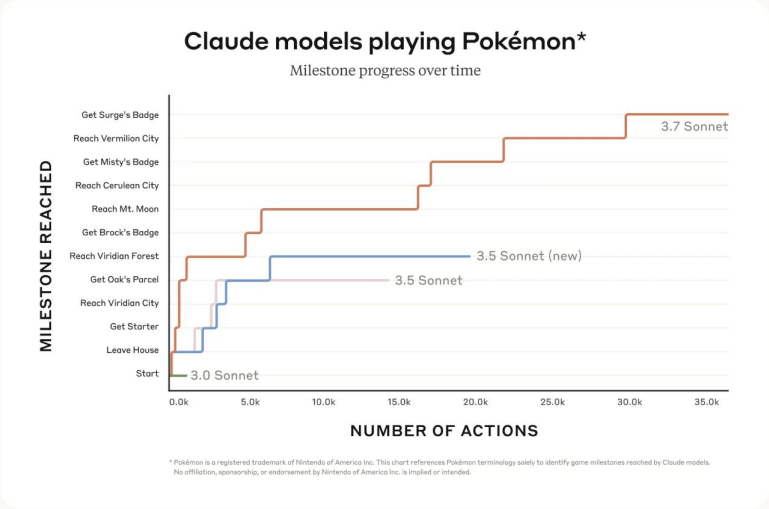

Claude 3.7 Sonnet demonstrates that it is the very best of all the Sonnet models so far at playing Pokémon Red. On the x-axis is the number of interactions Claude completes as it plays the game; on the y-axis are important milestones in the game involving collecting certain items, navigating to certain areas, and defeating certain game bosses.

The Game continues

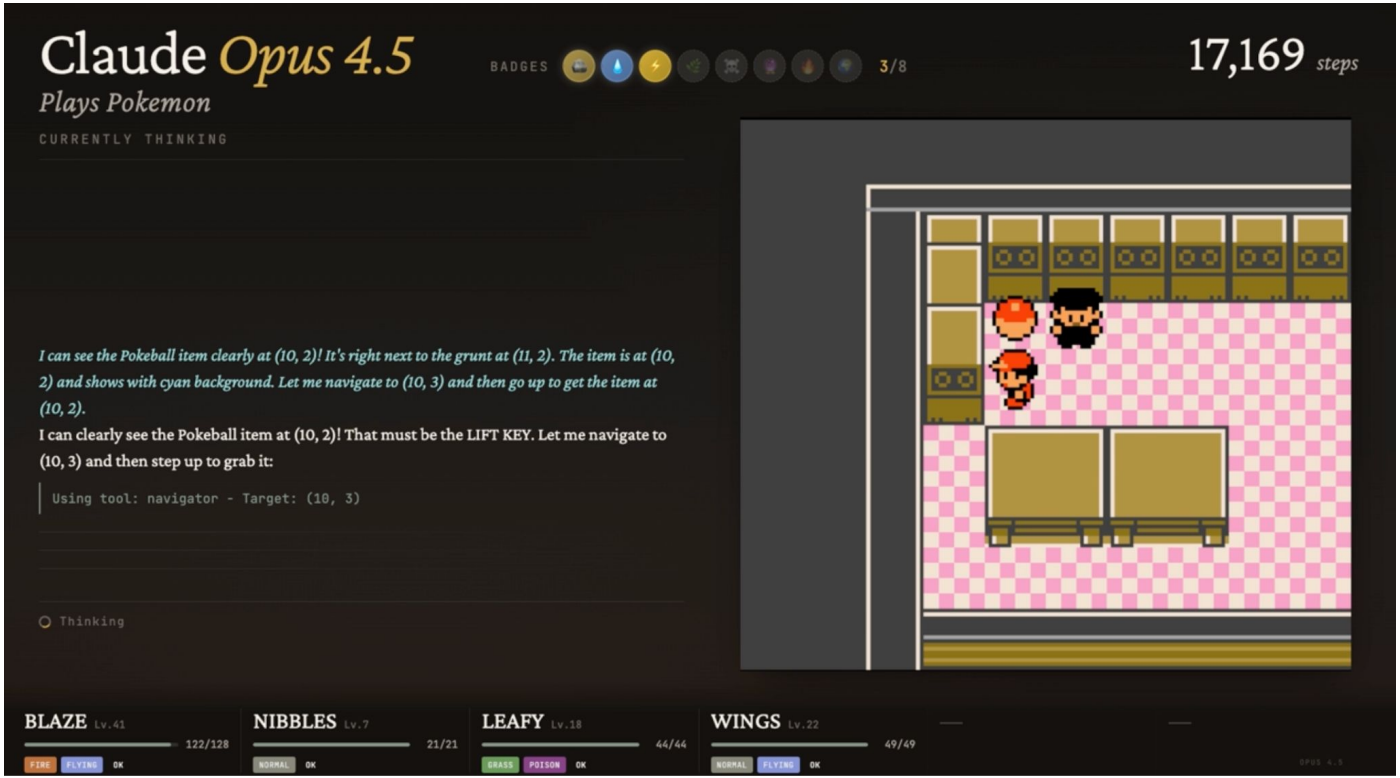

Currently Offline (but got to Team Rocket Hideout)

https://www.twitch.tv/claudeplayspokemon

Gemini also getting in on the action

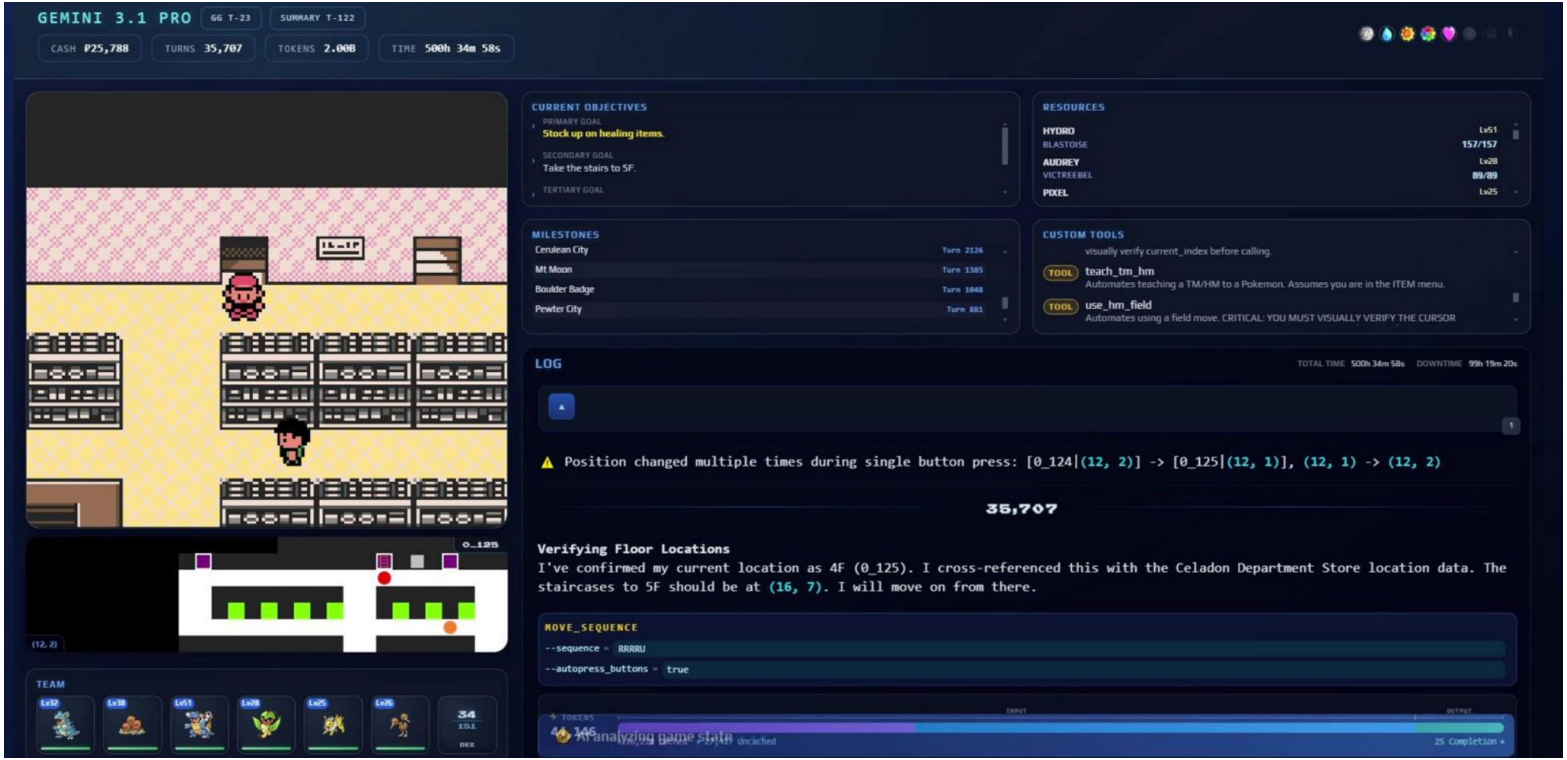

https://m.twitch.tv/gemini_plays_pokemon/home

Why would anyone care?

Games + AI - Long history

1992

TD-Gammon, A Self-Teaching Backgammon Program, Achieves Master-Level Play

Gerald Tesauro
IBM Thomas J. Watson Research Center
P. O. Box 704
Yorktown Heights, NY 10598
(tesauro@watson.ibm.com)

Abstract. TD-Gammon is a neural network that is able to teach itself to play backgammon solely by playing against itself and learning from the results, based on the TD(λ) reinforcement learning algorithm (Sutton, 1988). Despite starting from random initial weights (and hence random initial strategy), TD-Gammon achieves a surprisingly strong level of play. With zero knowledge built in at the start of learning (i.e. given only a "raw" description of the board state), the network learns to play at a strong intermediate level. Furthermore, when a set of hand-crafted features is added to the network's input representation, the result is a truly staggering level of performance: the latest version of TD-Gammon is now estimated to play at a strong master level that is extremely close to the world's best human players.

DeepBlue & Chess - 1997

- Basically a Tech Demo for IBM's supercomputers
- Showed advances in "Good-old-fashioned" AI
- Tree-search, Alpha-Beta pruning
- Beat grandmaster Gary Kasparov

IBM Watson

- Showed off new advances in NLP
- DeepQA
- Beat* Jeopardy! Champions Ken Jennings and Brad Rutter

*Hard-core Jeapardy! fans dispute this, is it cheating if the AI can always hit the buzzer first?

DQN Atari - 2013

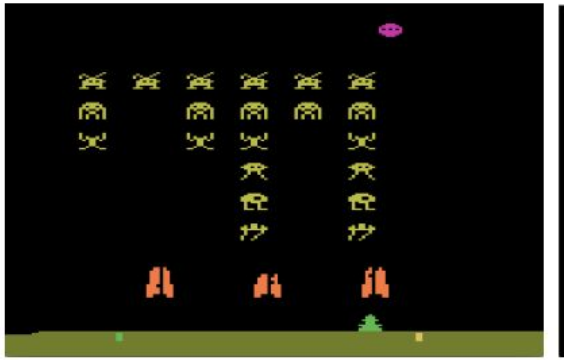

- DQN
- Showed great progress on Atari Games

	B. Rider	Breakout	Enduro	Pong	Q*bert	Seaquest	S. Invaders
Random	354	1.2	0	−20.4	157	110	179
Sarsa [3]	996	5.2	129	−19	614	665	271
Contingency [4]	1743	6	159	−17	960	723	268
DQN	**4092**	**168**	**470**	**20**	**1952**	**1705**	**581**
Human	7456	31	368	−3	18900	28010	3690
HNeat Best [8]	3616	52	106	19	1800	920	**1720**
HNeat Pixel [8]	1332	4	91	−16	1325	800	1145
DQN Best	**5184**	**225**	**661**	**21**	**4500**	**1740**	1075

AlphaGo - 2016

- Advanced MCTS and Neural networks
- Beat Go masters Lee Sedol, Ke Jie and Fan Hui
- Later became AlphaZero algorithm for general self-play with MCTS

What Games give AI

- Give us challenging problems
- Games are interesting to humans **because** they are hard for humans
- Lets us study specific problems in isolation / controlled environment
- Different games come into prominence as AI develops new techniques that let us tackle new problems
- Similarly, drive to win benefits AI as it gives us a goal to explore the frontier of what is possible

What is the next grand challenge for LLMs?

- What is a game that, if an LLM agent could do well on it, we'd say we've made progress?
- Like past examples, we want a game that shows the progress made so far in the community
- I don't think (personal opinion) we have an answer yet

Possible Answers

Pokemon?

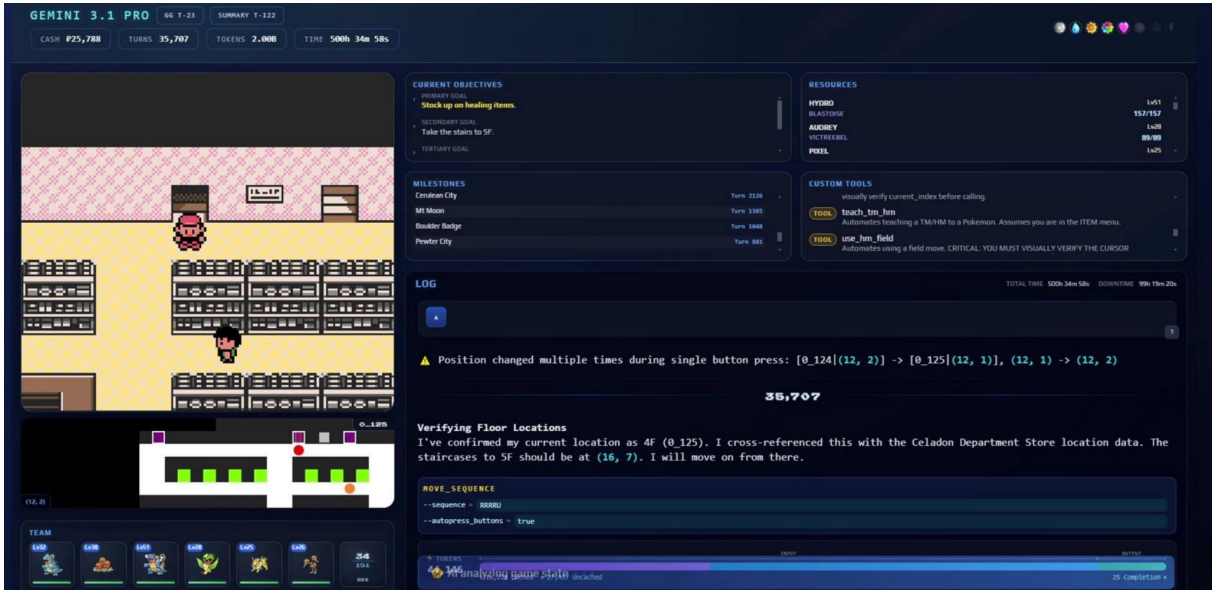

- Show LLM Agent improvements in memory / long-term reasoning and autonomy

NetHack?

- Surprisingly complex game
- Need really good strategy and execution
- Extreme long-term payoffs (essentially you win by getting as many resources as possible and delving as deep as possible)

NetHack (hard takeoff indicator?)

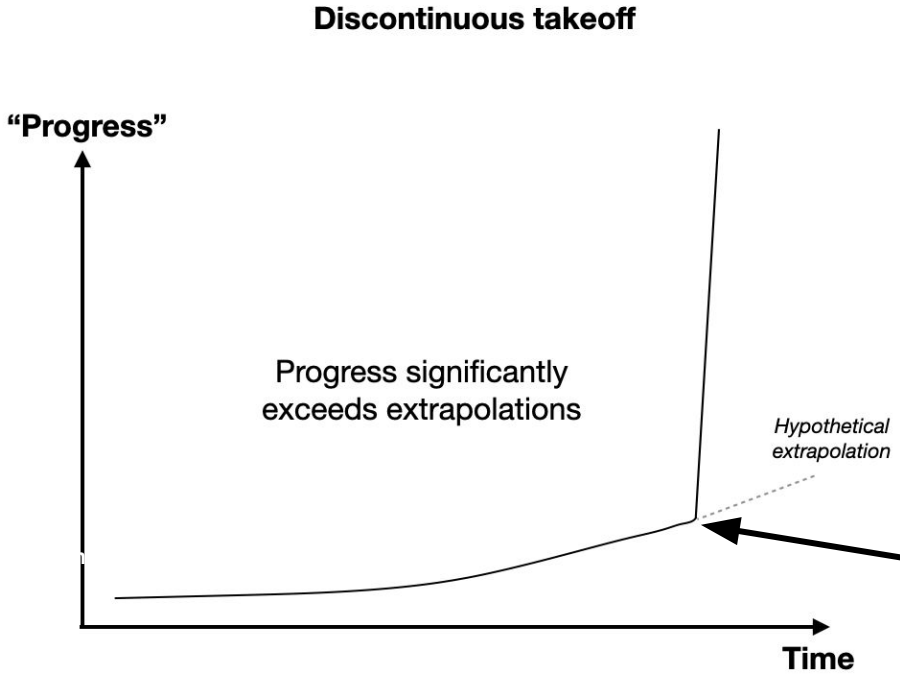

What about the human side?

- Can AI not only learn the rules of game but interact socially with humans and operate under complex human norms?
- This is my personal answer for, if AI could do it, we've made enormous progress

Can AI run a DND Game?

Can AI run a DND Game?

- A Tabletop role-playing game
- Group of players create a character and have an adventure in a fantasy setting where they can use swords, cast spells
- World is determined by rules written down
- Dungeon Master - person in charge of telling the story to the players that they can interact with and acts as the “computer” deciding what happens next based on the players’ actions

Can AI run a DND Game?

To run a game well you'd need

- Good understanding of natural language rules and how to apply them
- Ability to know how rules are applied in novel scenarios
- Describe and set the scene well for an engaging experience
- Set up an interactive story over a long period of time that's responsive to what players want
- Navigate social norms and expectations
- Have natural dialog (audio) that engages with players

Can AI ru Game?

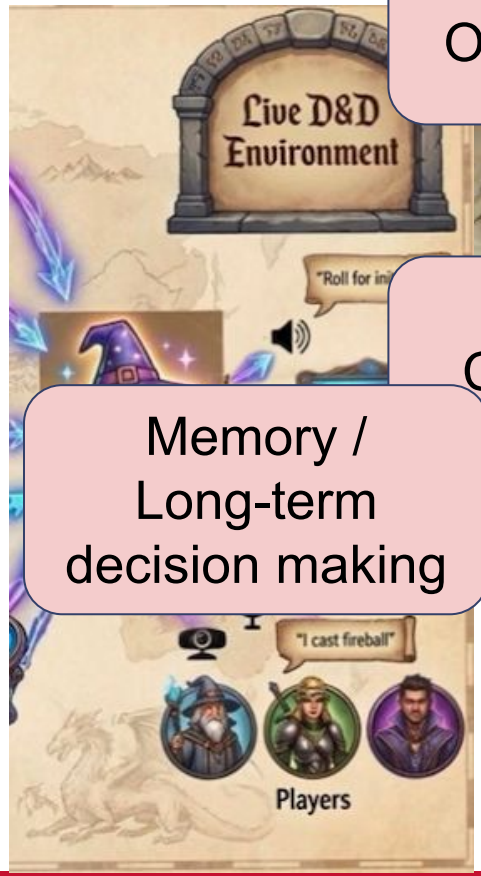

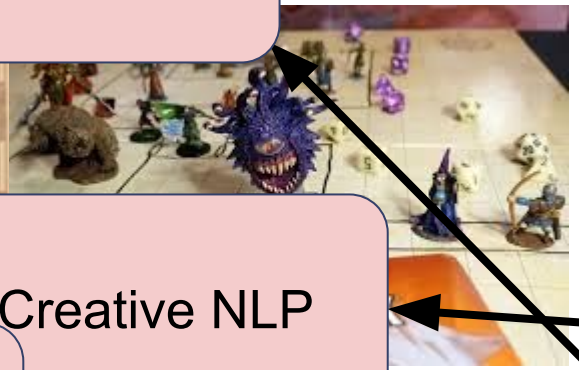

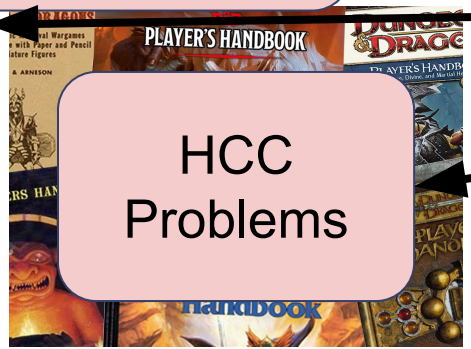

Adaptation Over Time

Creative NLP

Memory / Long-term decision making

HCC Problems

In-the-wild NLP / generalization

To run a game well you'd need

- Good understanding of natural language rules and how to apply them
- Ability to know how rules are applied in novel scenarios
- Describe and set the scene well for an engaging experience
- Set up an interactive story over a long period of time that's responsive to what players want
- Navigate social norms and expectations
- Have natural dialog (audio) that engages with players

Open Problems in VLM Agents

1. Long-horizon problems
2. Making the V in VLMs matter
3. Exploration
4. Continual/Neverending Learning
5. Personalization
6. Evaluation

Lots of what I see as open problems in Agents

Recap

- Games in AI
  - Give challenges to AI
  - Helps show-off/develop techniques to tackle specific problems
  - Lets us study specific problems in isolation / controlled environment
- Past Games in AI
  - Backgammon
  - Chess
  - Jeopardy!
  - Atari
  - Go
- Games for Agents
  - How to test natural language ability
  - Test generalization of LLMs
  - Test the “agent” aspects - long term decision-making

Any Questions

Questions

Now for the presentations!