

CSE 190 – Intro to Deep RL

What is an agent?

Prithviraj Ammanabrolu

What is an agent?

Who is asking the question?

- FBI → agent = human who does stuff
- Insurance → agent = human who does stuff
- Human centered research → agent = human who does stuff

See the theme?

What is an agent?

Who is asking the question?

- FBI → agent = human who does stuff
- Insurance → agent = human who does stuff
- Human centered research → agent = human who does stuff

See the theme? jk

- Person who attends NeurIPS → agent = AI (that does stuff?)

What does an AI agent do?

What would you like an AI agent to be able to do for you?

What does an AI agent do?

What would you like an AI agent to be able to do for you?

Some modern applications:

- File your reimbursements (workflow automation)
- Software development
- Do your laundry (household or commercial robotics)
- Get your degree for you (personal assistant)

What does an AI agent do?

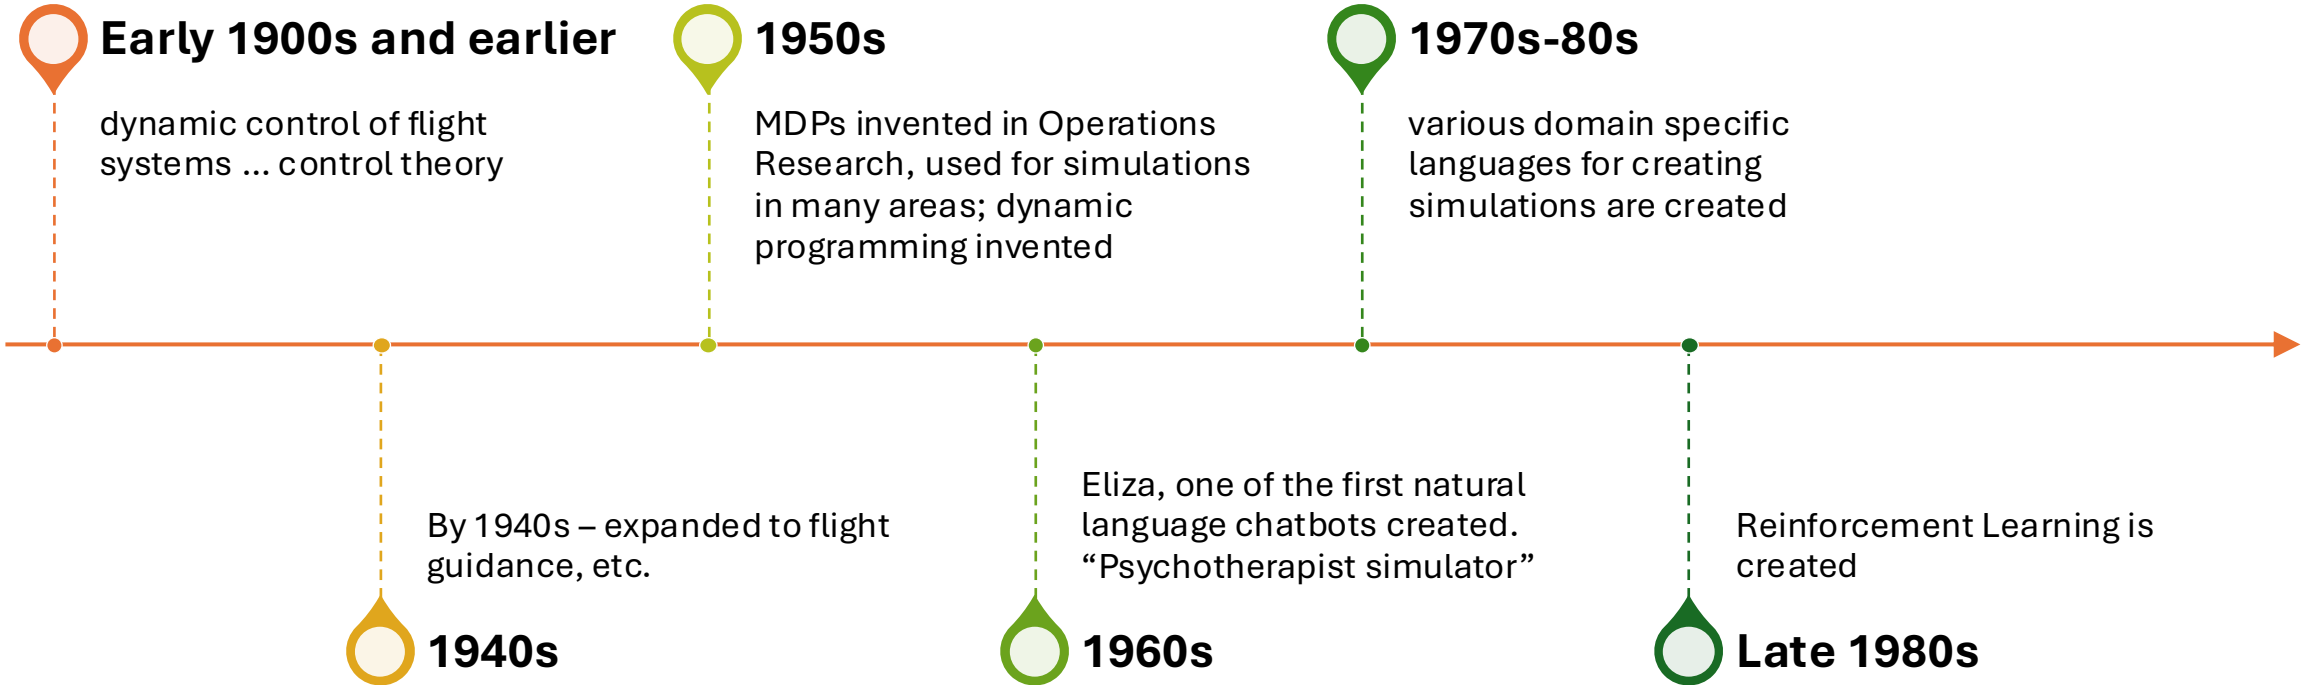
What would you like an AI agent to be able to do for you?

- ACT!

Agent or not?


- Let's look at some examples:
 - Auto pilot
 - Automated JavaScript buttons
 - Retrieval Augmented Generation
 - Warehouse robots
 - House robots
 - Voice assistants (early Alexa / Siri)
 - Voice assistants (Gemini, GPT + Apple)

A (very) brief history of “agents”




Sequential Decision Making

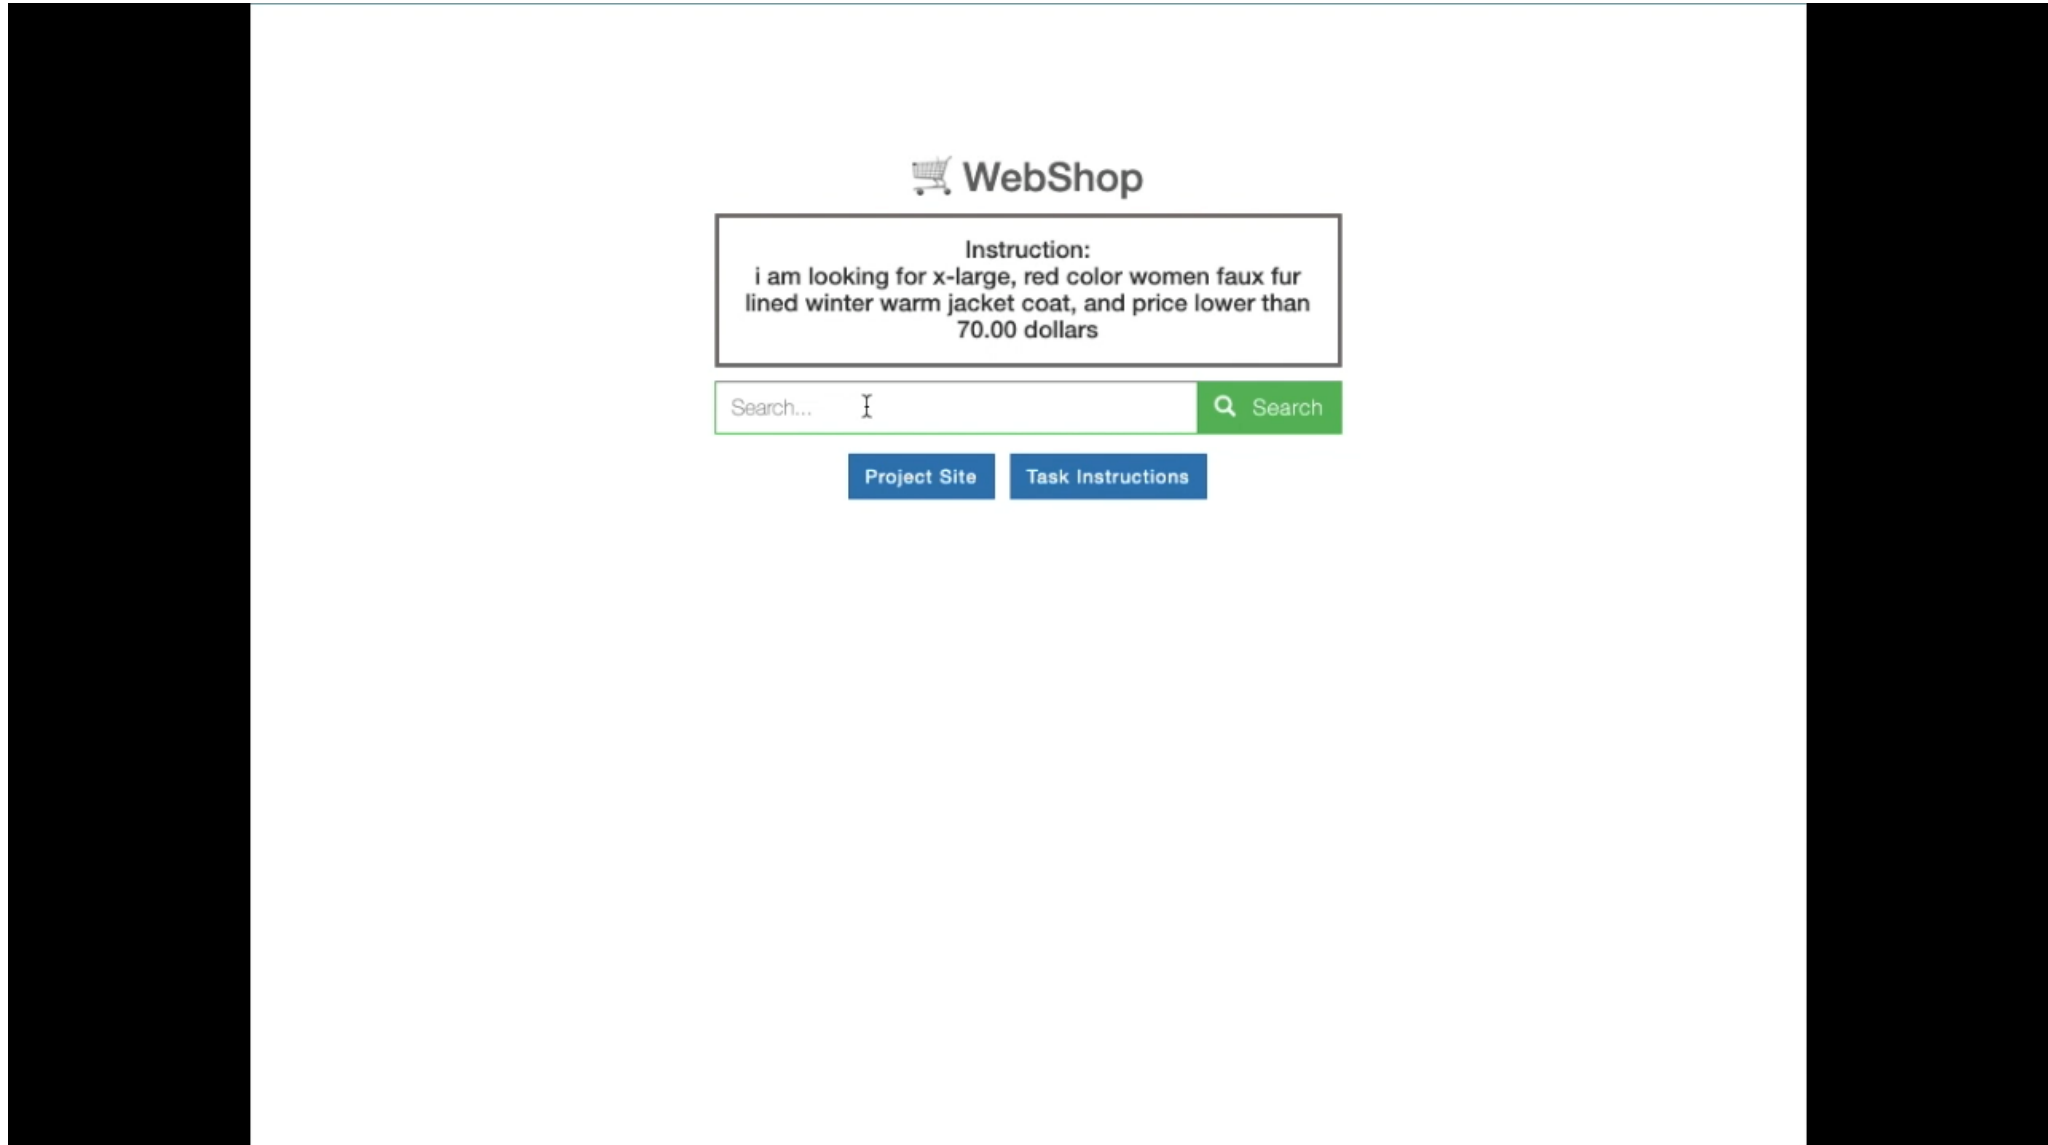
- Most of the tasks we'll cover are sequential decision making
- Why is this important? Different from other ML problems like classification?



Markov Decision Process



MDP (Modern LLM version)



MDP (Modern LLM version)

- Bonus: what does a MDP for language generation look like?

Components of an Agent

Required: Grounding,
Agency (ability to act),
Planning, Memory,
Learning

Additional:
Embodiment,
Communication, World
Modeling, Multimodality

Grounding (in an Environment)

- Language is anchored to “concepts” in the world
- Many types of language grounding
 - To other modalities (eg images of words/phrases)
 - To social / cultural norms (eg Do Americans really prefer car traffic over high speed rail?)
 - To action (eg pick up the cup from the right side of the table)

Agency

- *Choices* are required for actions
- If an agent has to select what tools to use but there's always only one tool, is that agency?

Agent vs Environment

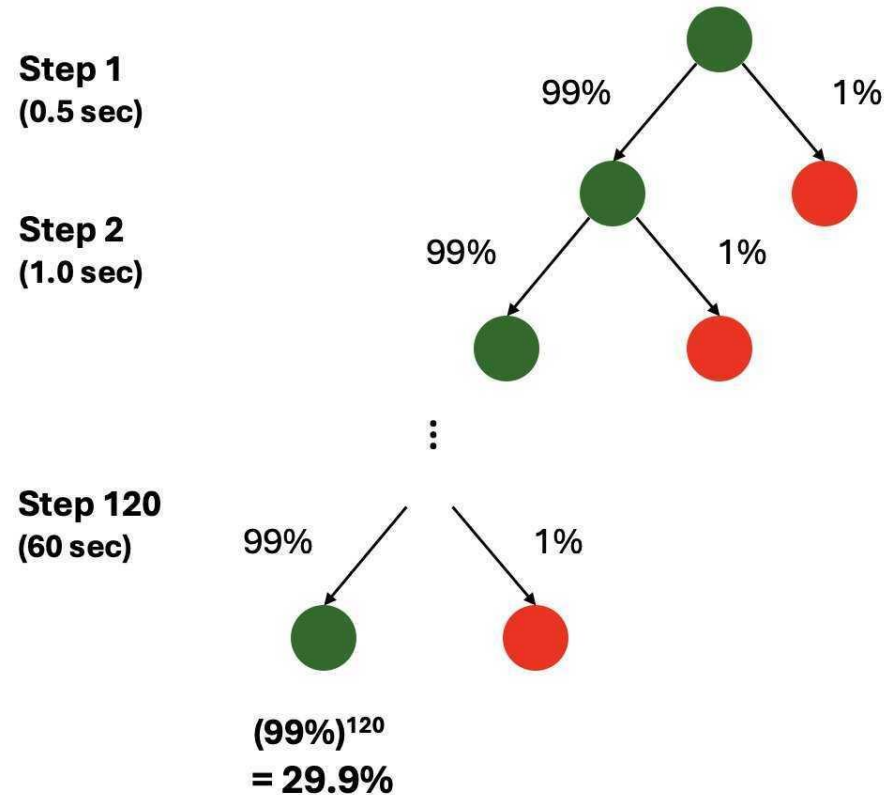
- No good definition (unless you count religious philosophy)
- Some rules of thumb
 - Strictly define the task, e.g. software engineering
 - Think about which one is easier to modify the behaviors of
- We'll talk more about the relationship between agents, simulations, and reality later

Memory

- Short term – what is the relevant information around me that I need to use to act *now*
- Long term – what information have I already learned can be used to help me now
 - Retrieval from the internet is one example of this

(Long Horizon) Planning / Reasoning

- 1 step 99% accuracy over even 120 steps = <30%



Learning (from feedback)

- Many different types of feedback
- Independent of “learning” mechanism
- Doesn't necessarily need to update model weights
 - In context learning (model responding to a prompt “no you got that wrong” is also learning from feedback)

Embodiment

- Robots!
- Physically acting in the real world
- “embodied” hypothesis that says embodiment is necessary for AGI

Communication

- Can the agent communicate its intentions to other agents?
- A necessary pre-req for multi-agent scenarios
 - What is a multi-agent scenario?

World Modeling

- Modeling transition matrix T by MDP definition
- Given the state of the world and an action, predict the next state of the world

Multimodality

- There are currently only a few trillion tokens of “clean” text \approx few terabytes.
 - Clear limit to how much you can scale. Only 1 internet and took us \sim 2-3 decades to get it.
- Youtube has 4.3 Petabytes of new videos a day
- CERN generates 1 Petabyte a day
 - Modalities exist outside vision and language!!
- Efficient use of this data is critical to scaling further


Components of an Agent

- Required: Grounding, Agency, Planning, Memory, Learning
- Additional: Embodiment, Communication, World Modeling, Multimodality

Someone tell me why I listed a few of these as required and the rest as “additional”?

Model vs AI System vs Agent: Rough Intuition

Model	AI System	Agent
GPT-4	ChatGPT	ChatGPT (computer use)
Forward passes of neural net	Mixing models together, model + scaffolding but no agency	Has agency + discussed components











Many software engineering abstractions and definitions exist.
All are roughly correct. Some are useful.

Agent or not?

- Let's look at some examples:
 - Auto pilot
 - Automated JavaScript buttons
 - Retrieval Augmented Generation
 - Warehouse robots
 - Gameplay agents
 - (Current) House robots
 - Voice assistants (early Alexa / Siri)
 - Voice assistants (Gemini, GPT + Apple)

Agent or not?

- Let's look at some examples:
 - Auto pilot 
 - Automated JavaScript buttons 
 - Retrieval Augmented Generation 
 - Warehouse robots 
 - Gameplay agents 
 - (Current) house robots 
 - Voice assistants (early Alexa / Siri) 
 - Voice assistants (Gemini, GPT + Apple) 

It is ok to not use an agent!

- Not every use case needs an agent
 - Most use cases just need models or "AI systems"
- Agents are complicated, if you don't need one then don't use it
- First try the simplest method you have for your task

HW 0 Question Time