CS11-711 Advanced NLP Agents

Sean Welleck





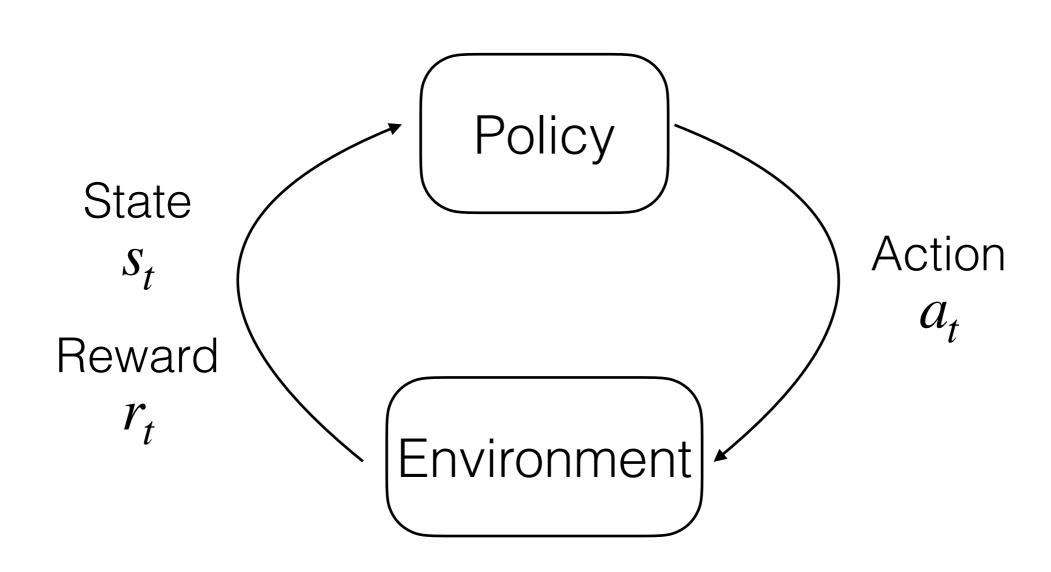


https://cmu-l3.github.io/anlp-fall2025/ https://github.com/cmu-l3/anlp-fall2025-code

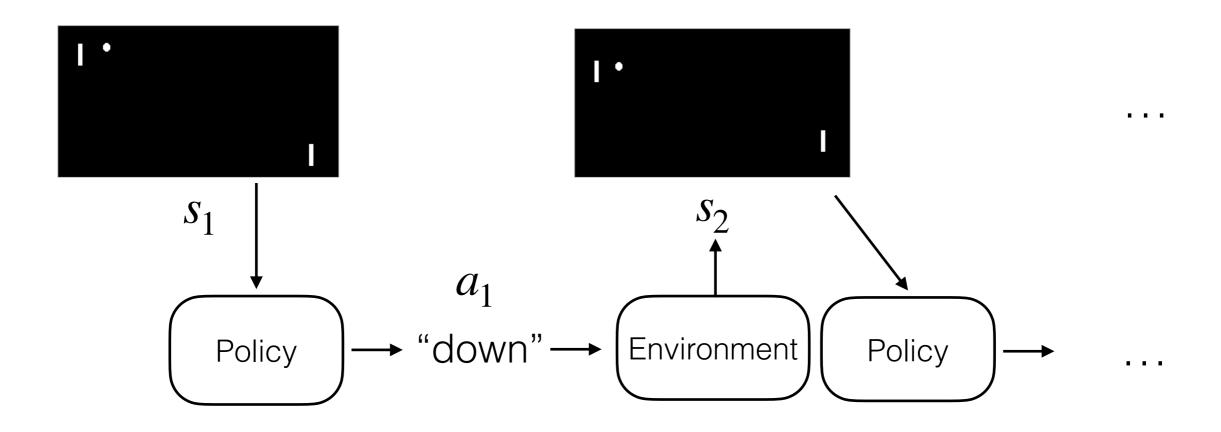
Recap: language models



Recap: RL MDP



Recap: RL MDP



Today: (LM) Agents

- What is an agent?
- Agent environments
- Agent design patterns
- Agents and RL

Code examples: https://github.com/cmu-l3/anlp-fall2025-code/tree/main/17_agents

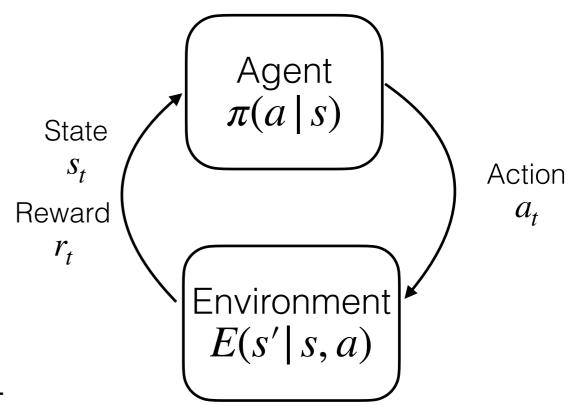
What is an Agent?

An agent is simply a policy!

Input: state

Output: action

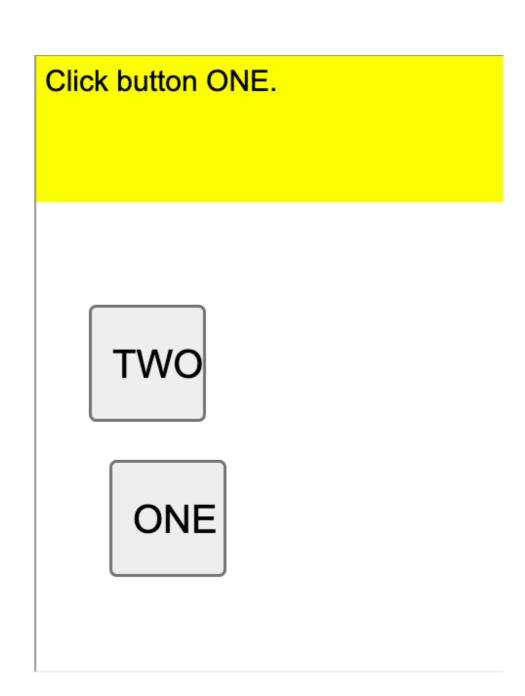
Operates in an environment



Example: simple web environments

Move the cube around so that "5" is the active side facing the user.	Set the sliders to the combination [13,20,13] and submit.	Draw the number "2" in the checkboxes using the example on the right and press Submit when finished.	Select 5 with the spinner and hit Submit.	Keep your mouse inside the circle as it moves around.	Move the cube around so that "4" is the active side facing the user.
Submit Submit	11 6 17 Submit	Submit	Select a value: 0 Submit	*	Submit
Copy the text in the textarea below and paste it into the textbox.	Select 09/23/2016 as the date and hit submit.	Drag all rectangles into the black box.	Enter "qu" into the text field and press Submit.	Select all the shades of blue and press Submit.	Find the 4th word in the paragraph, type that into the textbox and press "Submit".
Enim. Elementu m, tortor ullamcor Submit	Date: Submit	Submit	qu Ţ Submit	Submit	Non arcu ut ultricies est. Gravida gravida. Porta erat nulla eget condimentum posuere a Submit
Enter an item that starts with "Tuni".	Enter "Vb8" into the text field and press Submit.	Focus into the 1st input textbox.	Focus into the textbox.	Move the cube around so that "2" is the active side facing the user.	Select 5Gi and click Submit.
Tags: Submit	Submit			6 13 Submit	Submit

- Environment:
 - Simple webpage
- Task:
 - Click indicated button



```
env = gymnasium.make('miniwob/click-test-2-v1', render_mode='human')
try:
    observation, info = env.reset(seed=41)
    for step in range(100):
        action = agent(observation, env)
        observation, reward, terminated, truncated, info = env.step(action)

if terminated:
    observation, info = env.reset()
```

- Observation/state
 - Image
 - DOM elements

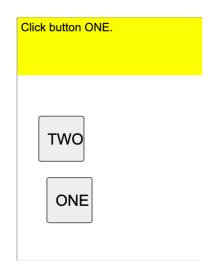
```
TWO

ONE
```

```
'tag': 'div',
 'text': '',
'top': array([50.], dtype=float32),
'value': '',
'width': array([160.], dtype=float32)},
{'bg_color': array([0.9372549, 0.9372549, 0.9372549, 1.
                                                              ], dtype=float32),
'classes': '',
'fg_color': array([0., 0., 0., 1.], dtype=float32),
'flags': array([0, 0, 0, 1], dtype=int8),
'height': array([40.], dtype=float32),
'id': 'subbtn',
'left': array([45.], dtype=float32),
'parent': 3,
'ref': 4,
'tag': 'button',
'text': 'ONE',
'top': array([76.], dtype=float32),
'value': '',
'width': array([40.], dtype=float32)},
{'bg_color': array([0.9372549, 0.9372549, 0.9372549, 1.
                                                              ], dtype=float32),
'classes': '',
'fg_color': array([0., 0., 0., 1.], dtype=float32),
'flags': array([0, 0, 0, 1], dtype=int8),
'height': array([40.], dtype=float32),
'id': 'subbtn2',
'left': array([30.], dtype=float32),
 'parent': 3,
'ref': 5,
'tag': 'button',
'text': 'TW0',
'top': array([109.], dtype=float32),
'value': '',
'width': array([40.], dtype=float32)})
```

Actions

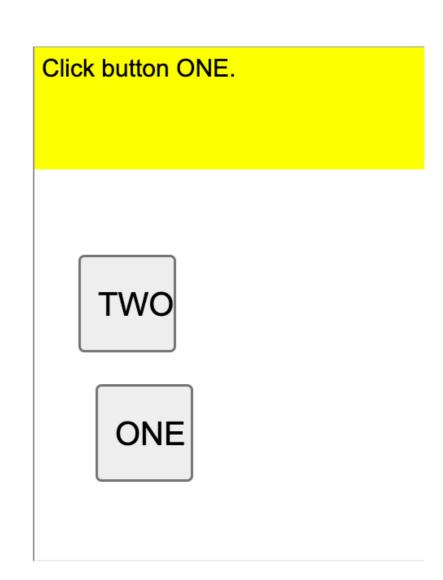
MOVE_COORDS <x, y>,...CLICK_ELEMENT <REF>,



Name	Description
NONE	Do nothing for the current step.
MOVE_COORDS	Move the cursor to the specified coordinates.
CLICK_COORDS	Click on the specified coordinates.
DBLCLICK_COORDS	Double-click on the specified coordinates.
MOUSEDOWN_COORDS	Start dragging on the specified coordinates.
MOUSEUP_COORDS	Stop dragging on the specified coordinates.
SCROLL_UP_COORDS	Scroll up on the mouse wheel at the specified coordinates.
SCROLL_DOWN_COORDS	Scroll down on the mouse wheel at the specified coordinates.
CLICK_ELEMENT	Click on the specified element using JavaScript.
PRESS_KEY	Press the specified key or key combination.
TYPE_TEXT	Type the specified string.
TYPE_FIELD	Type the value of the specified task field.
FOCUS_ELEMENT_AND_TYPE_TEXT	Click on the specified element using JavaScript, and then type the specified string.
FOCUS_ELEMENT_AND_TYPE_FIELD	Click on the specified element using JavaScript, and then type the value of the specified task field.

A simple rule-based agent

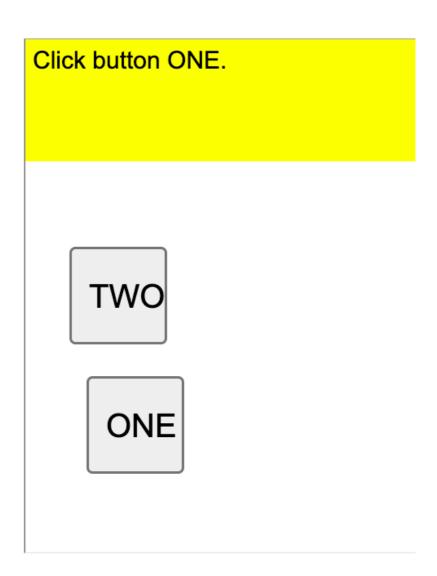
- Write a program that:
 - Finds the DOM element with text 'ONE'
 - Get its REF
 - Output the action CLICK_ELEMENT <REF>



A simple rule-based agent

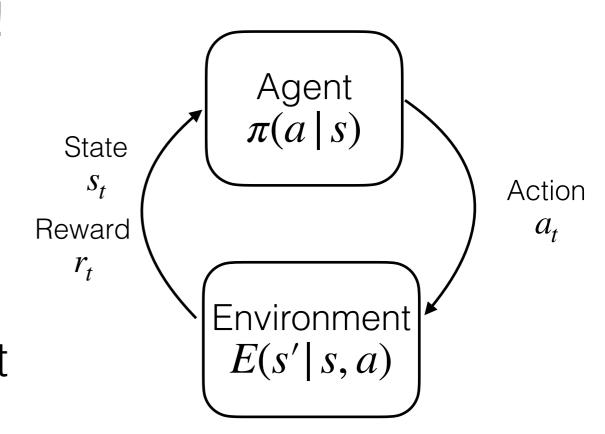
```
def agent(observation, env):
    # Find the HTML element with text "ONE"
    one_element = None
    for element in observation['dom_elements']:
        if element['text'] == "ONE":
            one_element = element
            break

# Click on the element.
        action = env.unwrapped.create_action(
            ActionTypes.CLICK_ELEMENT, ref=one_element["ref"]
        )
        return action
```



What is an Agent?

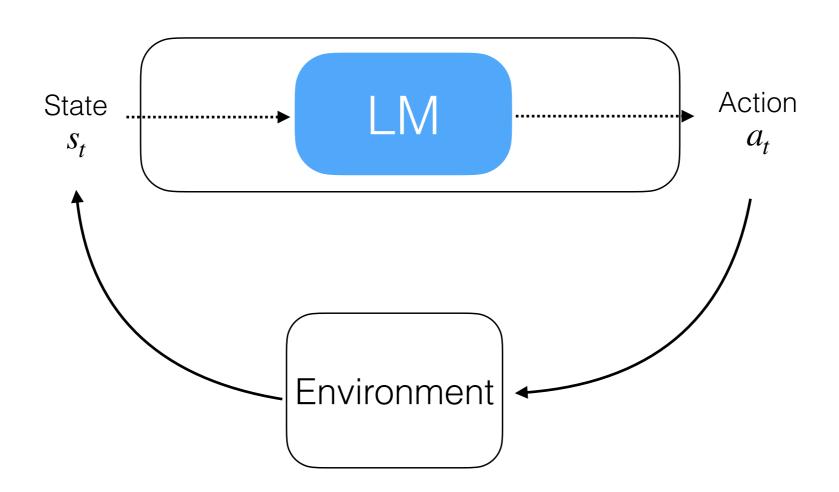
- An agent is simply a policy!
 - Input: state
 - Output: action
- Operates in an environment



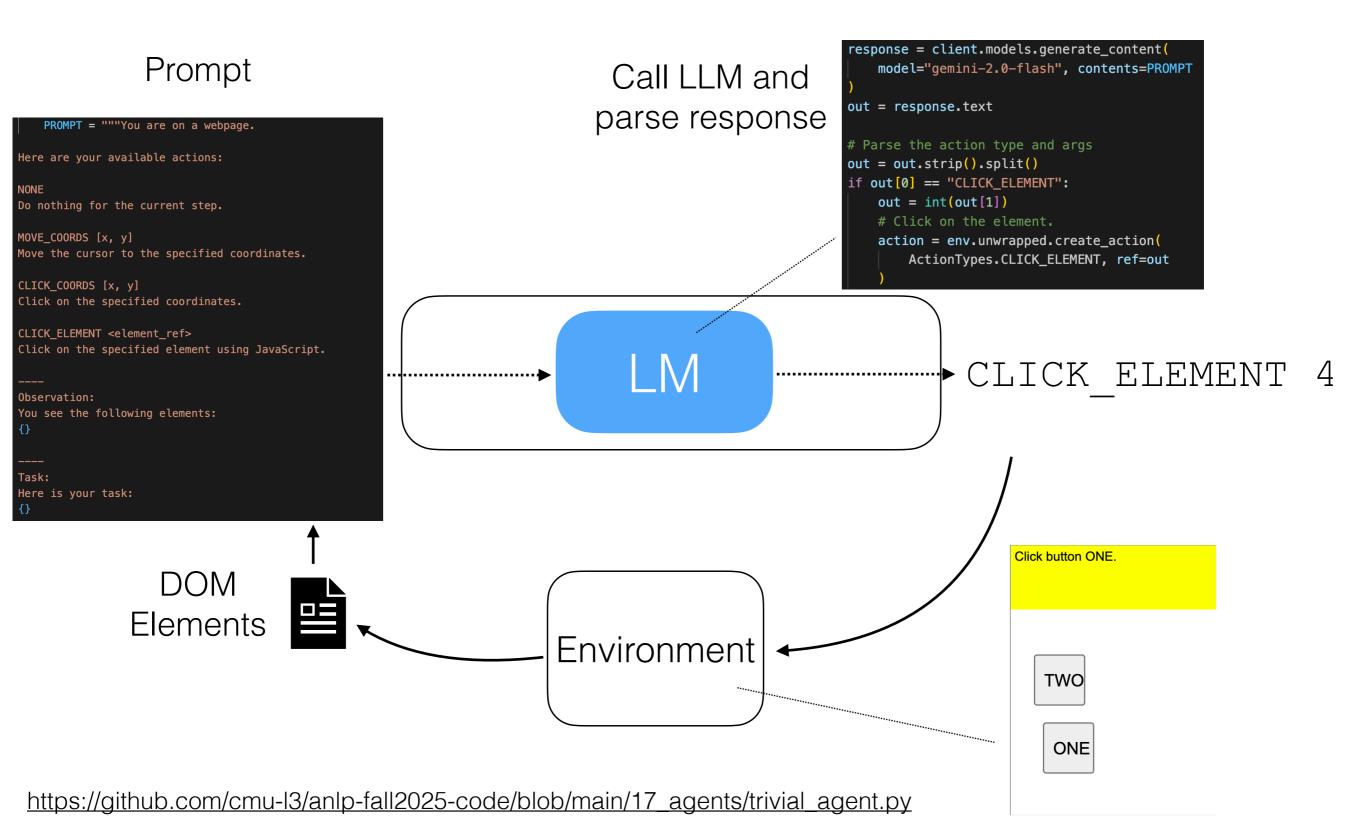
Our rule-based agent: only works for one narrow task

What is a **LLM** Agent?

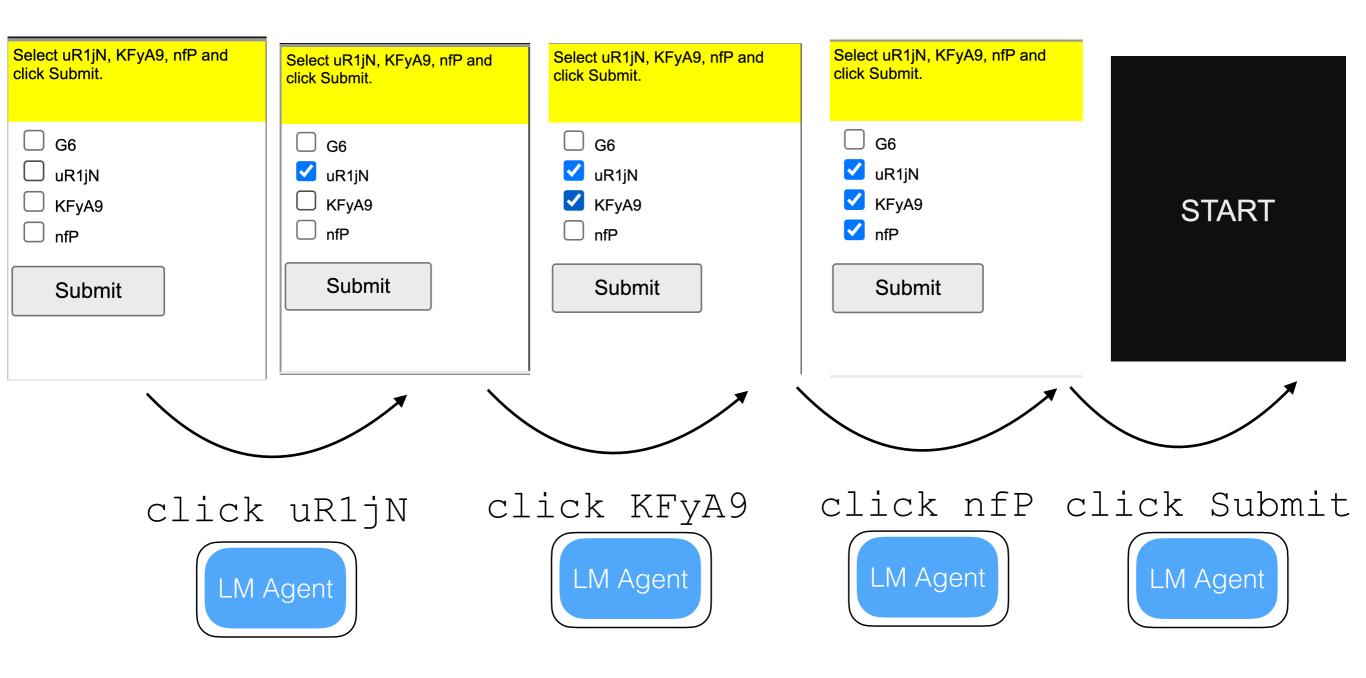
An agent where a language model helps determine actions



Example

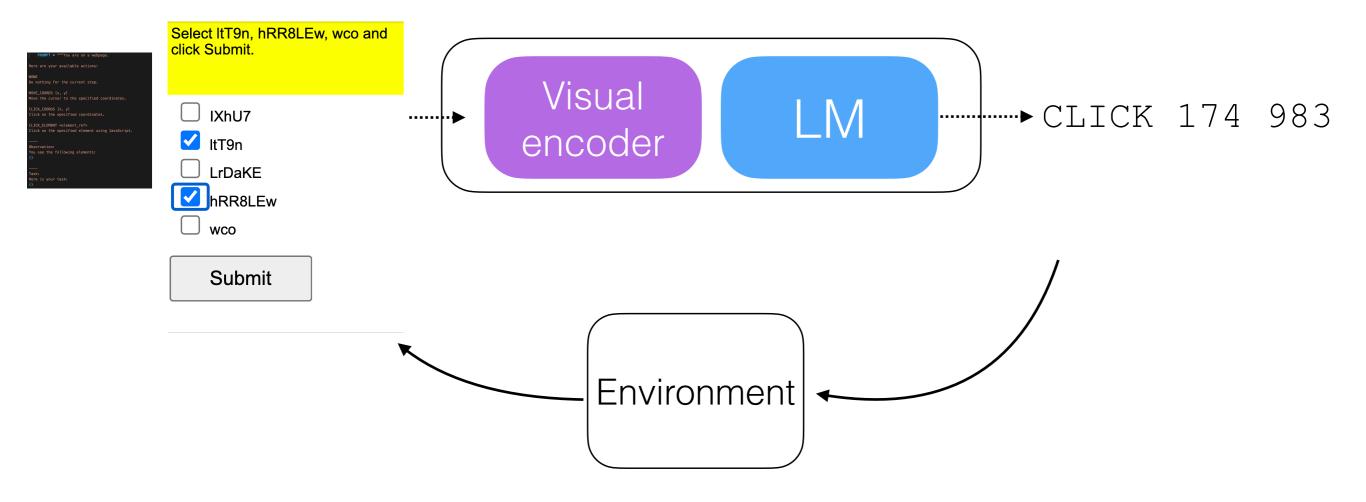


Example



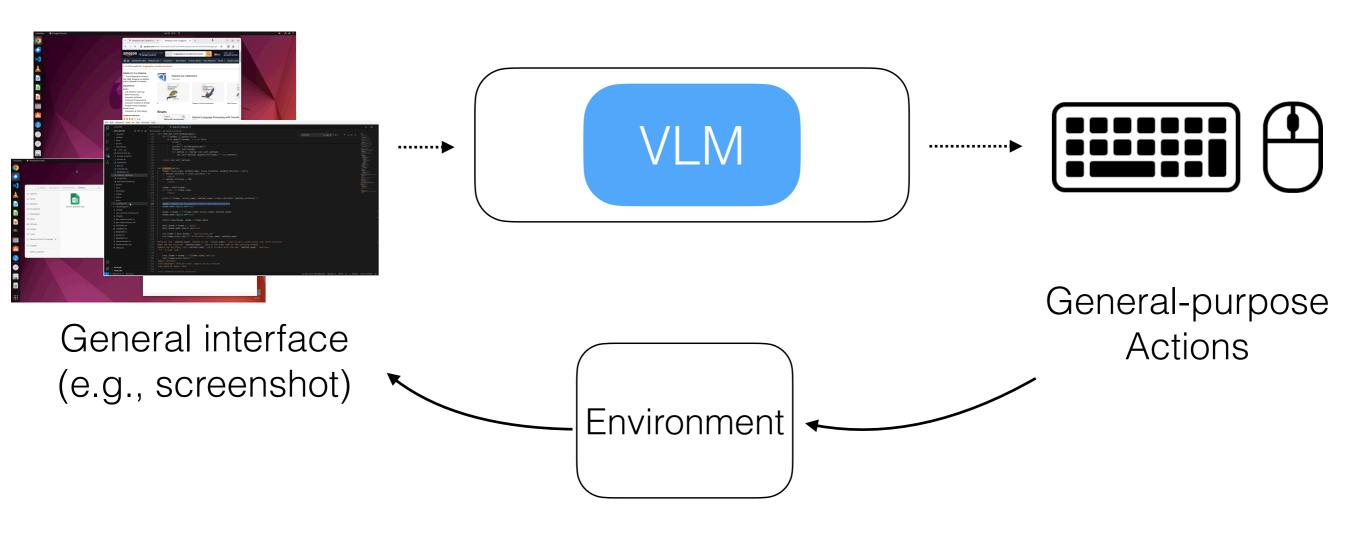
Vision Language Model (VLM) Agent

 An agent where a vision language model helps determine actions



Computer-use agent

 An agent with a general-purpose observation and action space (i.e., works for any computer task)



Recap: what is an agent?

- An agent is a policy $\pi(a \mid s)$ in an environment
- Text agents (e.g., LM agents)
 - Operate with text inputs and outputs
- Visual agents (e.g., VLM agents)
 - Additionally operate with images
 - Example: Computer-use agents
 - A task-agnostic visual agent for computer tasks

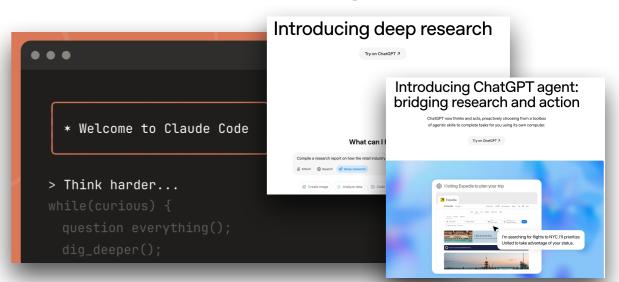
Lots of recent interest in agents

- Software engineering: e.g., "here's a GitHub PR for my repository, can you close it?"
- Research: e.g., "compile a research report on how the Al industry has changed in the last 4 months"
- Web: e.g., "Buy the cheapest color photo printer and send it to Emily's place (address in image)."

Nvidia CEO Says 2025 Is the Year of AI Agents

By Tae Kim Follow

Jan 07, 2025, 5:40 pm EST



Today: (LM) Agents

- What is an agent?
- Agent environments
- Agent design patterns
- Agents and RL

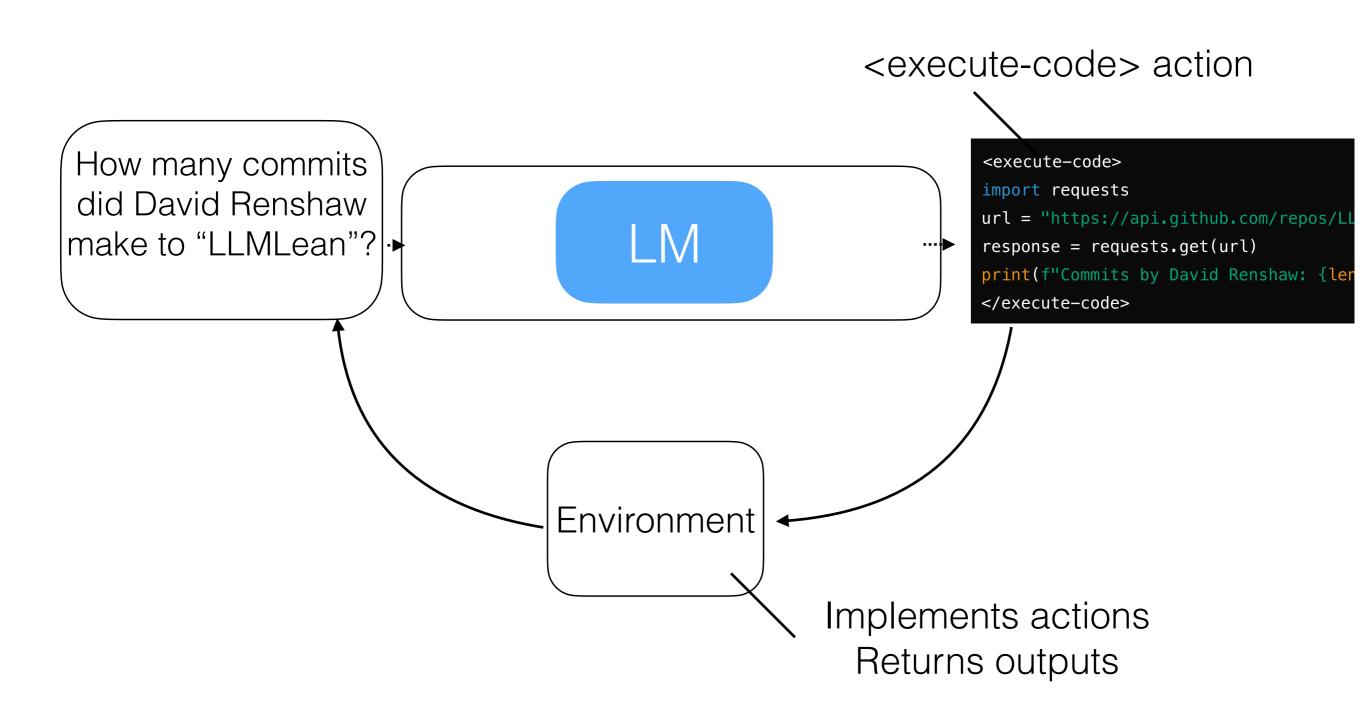
Agent environments

- Text agents
- Visual agents
- Computer-use agents

Domain-specific text environments

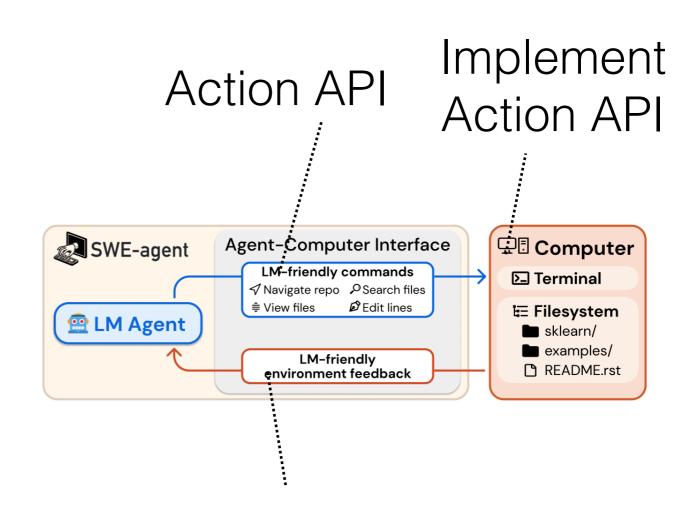
- Basic idea: engineer an environment for a domain
 - Example: software engineering agent
 - Actions: search_repo, execute code, edit file
 - Environment: implements the action API
 - Observations: returns outputs

Example



Example: SWE-agent environment

- Task: close pull requests
- Observations:
 - Action outputs
 - Task instruction
- Actions:
 - Open <path>
 - goto <line_number>
 - search_file <search_term> [<file>]
- Environment: implements action API, formats outputs



Format outputs

psf_requests-2317 **Issue**

We're currently solving the following issue within our repository. Here's the issue text: **ISSUE:**

method = builtin str(method) problem

In requests/sessions.py is a command:

method = builtin_str(method)

Converts method from

b'GET'

"b'GET'"

Which is the literal string, no longer a binary string. When requests tries to use the method "b'GET", it gets a 404 Not Found response.

I am using python3.4 and python-neutronclient (2.3.9) with requests (2.4.3). neutronclient is broken because it uses this "args = utils.safe_encode_list(args)" command which converts all the values to binary string, including method.

I'm not sure if this is a bug with neutronclient or a bug with requests, but I'm starting here. Seems if requests handled the method value being a binary string, we wouldn't have any problem.

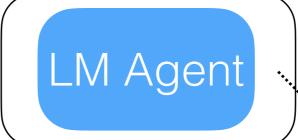
Also, I tried in python 2.6 and this bug doesn't exist there. Some difference between 2.6 and 3.4 makes this not work right.

INSTRUCTIONS: (31 lines)

(Open file: n/a) (Current directory: /psf__requests) bash-\\$

Observation psf__requests-2317

```
Found 2 matches for "sessions.py" in /psf__requests:
/psf__requests/build/lib/requests/sessions.py
/psf__requests/requests/sessions.py
(Open file: n/a)
(Current directory: /psf__requests)
bash-$
```



SWE-agent

psf__requests-2317

To address the issue, we first need to replicate the behavior described. The issue mentions that the 'builtin_str' function in 'requests/sessions.py' is converting a binary string method like 'b'GET' into the literal string "b'GET", which is incorrect. We should start by locating the 'sessions.py' file within the 'requests' library to examine the code around the 'builtin_str' function call. Let's begin by finding the 'sessions.py' file in the current repository.

find_file "sessions.py"

Environment

Basic idea: support new domains/tasks by implementing more action API calls

AutoCodeRover:

API name	Description	Output
search_class (cls)	Search for class cls in the codebase.	Signature of the searched class.
search_class_in_file (cls, f)	Search for class cls in file f.	Signature of the searched class.
search_method (m)	Search for method m in the codebase.	Implementation of the searched method.
search_method_in_class (m, cls)	Search for method m in class cls.	Implementation of the searched method.
search_method_in_file (m, f)	Search for method m in file f.	Implementation of the searched method.
search_code (c)	Search for code snippet c in the codebase.	+/- 3 lines of the searched snippet c.
search_code_in_file (c, f)	Search for code snippet c in file f.	+/- 3 lines of the searched snippet c.

CTF:

Category	Command	Documentation
Static analysis	<pre>decompile <binary_path> [-function_name <function_name>]</function_name></binary_path></pre>	Decompile a binary and prints the decompilation of a given function name, or main by default.
	<pre>disassemble <binary_path> [-function_name <function_name>]</function_name></binary_path></pre>	Disassembles a binary and prints the disassembly of a given function name, or main by default.
Dynamic analysis	debug_start <binary> [<args>]</args></binary>	Starts a debug session with the given binary and optional command- line arguments for the binary.
	<pre>debug_add_breakpoint <breakpoint></breakpoint></pre>	Adds a breakpoint in the debug session.
	<pre>debug_continue debug_step [<number>] debug_exec <command/> debug_stop</number></pre>	Continues the program execution in the debug session. Steps number of instructions (default is 1) in the debug session. Executes arbitrary gdb command in debug session. Stops the current debug session.
Network	<pre>connect_start <server_address> <port></port></server_address></pre>	Starts a new interactive connection to the server address and port.
	<pre>connect_sendline [<line>] connect_exec <command/> connect_stop</line></pre>	Sends a line with Unicode/hexadecimal values to the connection. Sending hexadecimal bytes should be done using \x <hh> where hh stands for the specific byte you want to send. Executes arbitrary connect command in connect session. Stops the current connect session.</hh>
Task	submit ' <flag>' exit forfeit</flag>	Submits your current flag and terminates the session upon correct response, you should be aware to properly escape the flag as this is a bash command and to put your flag in single quotes. Give up on the current challenge and terminate the session.

Example: text-based web browsing

- Task: answer questions
- Observations:
 - Action outputs
 - Question
- Actions:
 - Search <query>
 - Quote <text>
 - End: Answer

. . .

 Environment: implements action API, formats outputs (see Appendix A)

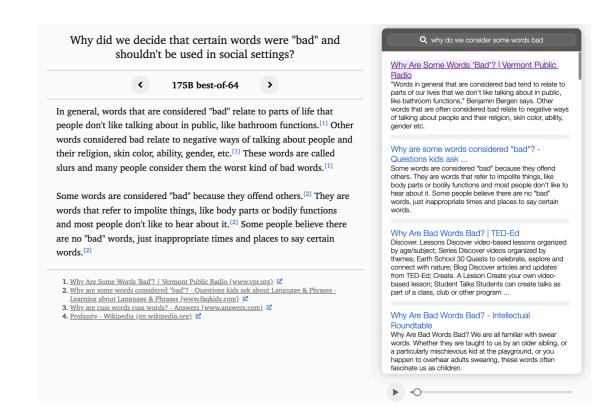
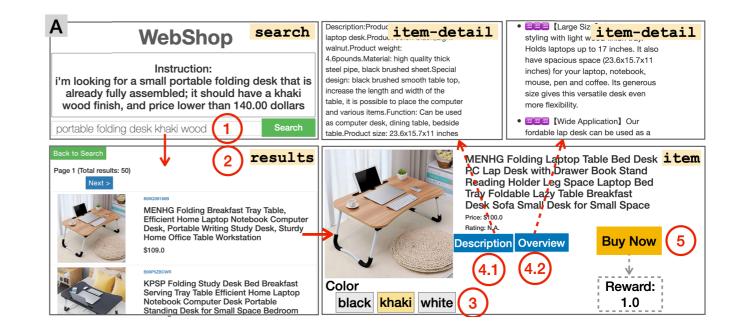


Table 1: Actions the model can take. If a model generates any other text, it is considered to be an invalid action. Invalid actions still count towards the maximum, but are otherwise ignored.

Command	Effect
Search <query></query>	Send <query> to the Bing API and display a search results page</query>
Clicked on link <link id=""/>	Follow the link with the given ID to a new page
Find in page: <text></text>	Find the next occurrence of <text> and scroll to it</text>
Quote: <text></text>	If <text> is found in the current page, add it as a reference</text>
Scrolled down <1, 2, 3>	Scroll down a number of times
Scrolled up <1, 2, 3>	Scroll up a number of times
Тор	Scroll to the top of the page
Back	Go to the previous page
End: Answer	End browsing and move to answering phase
<pre>End: <nonsense, controversial=""></nonsense,></pre>	End browsing and skip answering phase

Example: text-based web shopping

- Task: buy products
- Observations:
 - HTML of web page
 - Instruction
- Environment: simulates an ecommerce website



Type	Argument	State \rightarrow Next State
search	[Query]	Search \rightarrow Results
choose	Back to search	$* \rightarrow Search$
choose	Prev/Next page	Results \rightarrow Results
choose	[Product title]	Results \rightarrow Item
choose	[Option]	Item \rightarrow Item
choose	Desc/Overview	Item \rightarrow Item-Detail
choose	Previous	Item-Detail \rightarrow Item
choose	Buy	Item \rightarrow Episode End

Table 1: Actions in WebShop.

Agent environments

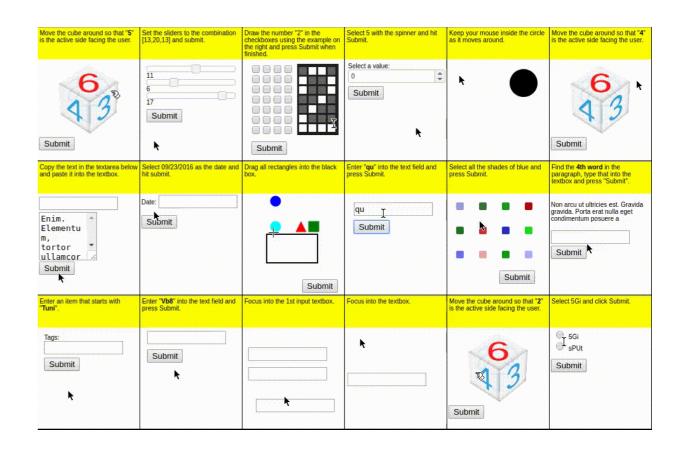
- Text agents
- Visual agents
- Computer-use agents

Visual agent environments

Text + image observations

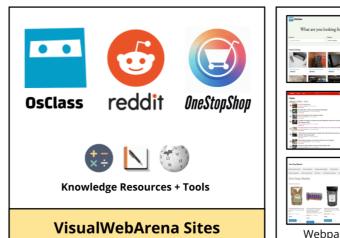
Example: basic web browsing

- Task: described in instruction (e.g., "click button")
- Observations:
 - Image of webpage
- Actions:
 - Move, click, scroll, type, ...
- Environment: implements action API, runs web browser and pages, keeps track of reward



Example: web tasks

- Task: described in instruction
- Observations:
 - Screenshot of webpage
- Actions:
 - Click element, type on element, close tab, go back, ...
- Environment: simulates websites:
 - Web shopping
 - Reddit
 - Classifieds marketplace





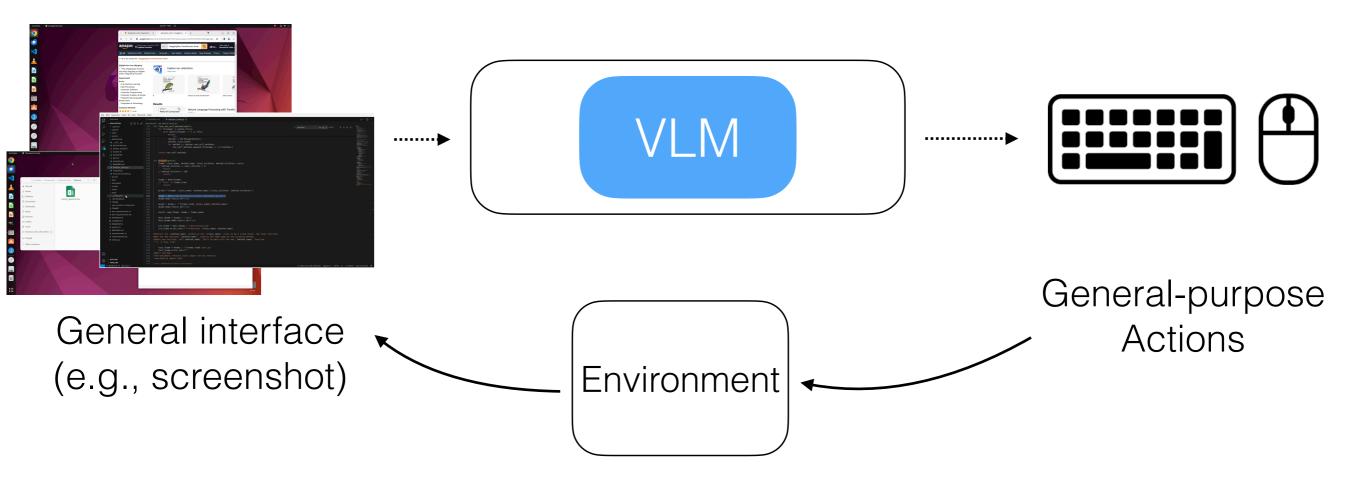
Action Type a	Description
click [elem]	Click on element elem.
hover [elem]	Hover on element elem.
<pre>type [elem] [text]</pre>	Type text on element elem.
<pre>press [key_comb]</pre>	Press a key combination.
new_tab	Open a new tab.
<pre>tab_focus [index]</pre>	Focus on the i-th tab.
tab_close	Close current tab.
goto [url]	Open url.
go_back	Click the back button.
go_forward	Click the forward button.
scroll [up down]	Scroll up or down the page.
stop [answer]	End the task with an output.

Agent environments

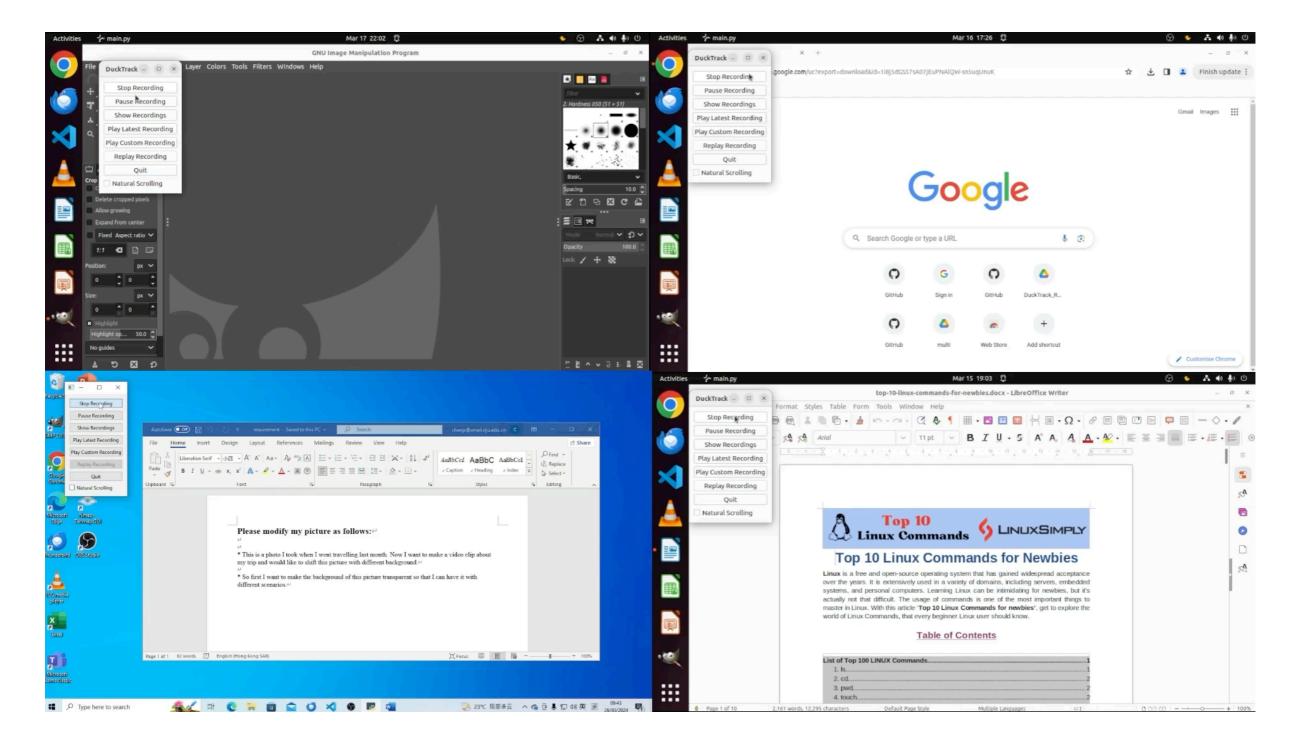
- Text agents
- Visual agents
- Computer-use agents

Computer-use

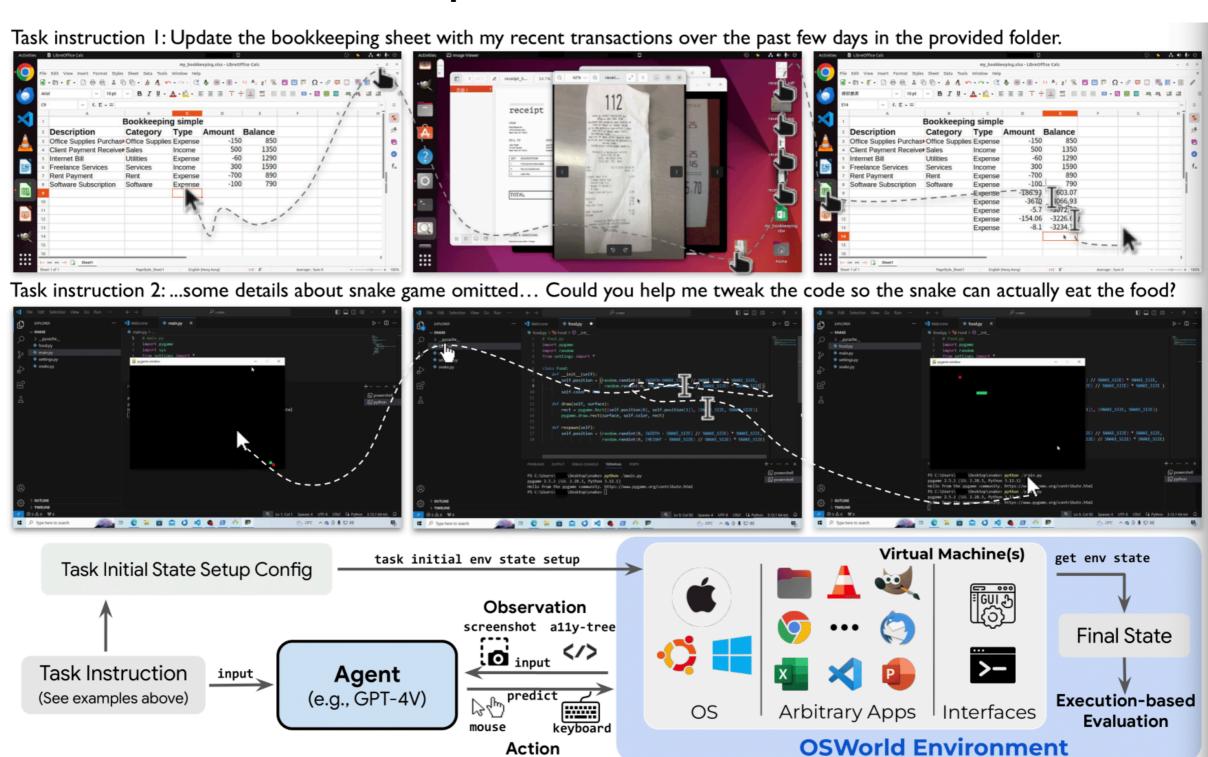
- Basic idea: interact with a computer like a human
 - Expressive: in principle, agents can perform any task a human can on a computer



Example: OS World



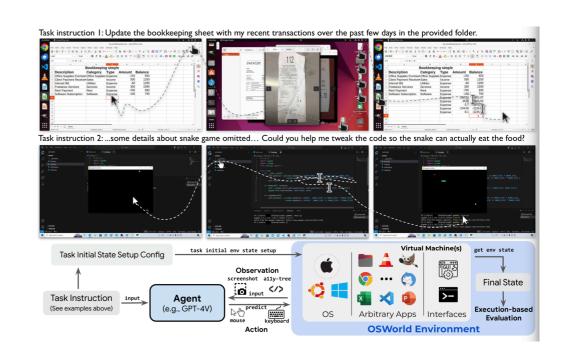
Example: OS World



Action

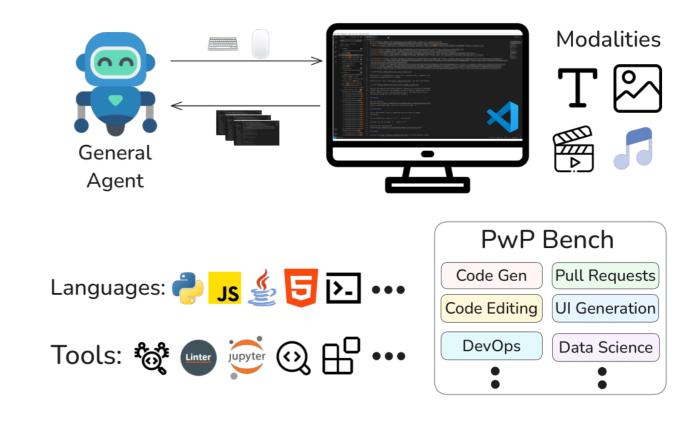
Example: OS World

- Task
 - Any task achievable on an OS
- Observations
 - Screenshot and/or accessibility tree
 - Task instruction
- Actions
 - Mouse, keyboard
- Rewards
 - Execute an evaluation script for the task
- Environment: runs an operating system



Example: Programming with Pixels

- Tasks
 - Any task achievable in an IDE (VS Code)
- Observations
 - Screenshot
 - Task instruction
- Actions
 - Mouse, keyboard
- Rewards
 - Execute an evaluation script for the task
- Environment: runs a modified VS Code



Example: Programming with Pixels

- Tasks
 - Any task achievable in an IDE (VS Code)
- Observations
 - Screenshot
 - Task instruction
- Actions
 - Mouse, keyboard
- Rewards
 - Execute an evaluation script for the task
- Environment: runs a modified VS Code

```
from pwp import PwPBench

bench = PwPBench('design2code')

dataset = bench.get_dataset()
 row = dataset[0]

env = bench.get_env(row)

for step in range(20):
    obs = env.get_observation()
    action = agent.get_action(obs)

env.step(action)

score = bench.get_reward(env, row)
```

Today: (LM) Agents

- What is an agent?
- Agent environments
- Agent design patterns
- Agents and RL

Agent Patterns

- Chain-of-thought
- Scaffolding
- Tool calling and MCP

Chain of thought

- Prompt the model to output a thought prior to the action
- We can view this as:
 - Policy $\pi_{\theta}(a_t | s_t)$ internally generates a thought z_t to determine its distribution over actions
 - The action decomposes into (z_t, a_t)

Response (Thought + Edit Action)

We need to add a check to ensure that all powers are the same. Let's edit the code to include this check.

```
edit 404:407
    elif (total_degree > 3 and
        len(set(k.exp[k.is_Pow])) == 1):
        if all(coeff[k] == 1):
            diop_type = 'general_sum_even'
end_of_edit
```

Scaffolding

- Scaffolding: code that structures how agents interact with the environment
 - Action/observation loop
 - Action definitions and formatting ("tools")
 - State management between turns (e.g., formatting observations, maintaining a history)
 - System prompt: defines agent's role and capabilities

Example: OpenHands

Scaffolding for code agents

```
You are OpenHands agent, a helpful AI assistant that can interact with a computer to solve tasks. <IMPORTANT>

* If user provides a path, you should NOT assume it's relative to the current working directory. Instead, you should explore the file system to find the file before working on it.

* When configuring git credentials, use "openhands" as the user.name and "openhands@all-hands.dev" as the user.email by default, unless explicitly instructed otherwise.

* The assistant MUST NOT include comments in the code unless they are necessary to describe non-obvious behavior.

</IMPORTANT>
```

```
_BASH_DESCRIPTION = """Execute a bash command in the terminal.

* Long running commands: For commands that may run indefinitely, it should be run in the state of the state of
```

```
_BROWSER_TOOL_DESCRIPTION = """

The following 15 functions are available. Nothing else is supported.

goto(url: str)

Description: Navigate to a url.

Examples:
    goto('http://www.example.com')

go_back()

Description: Navigate to the previous page in history.

Examples:
    go_back()

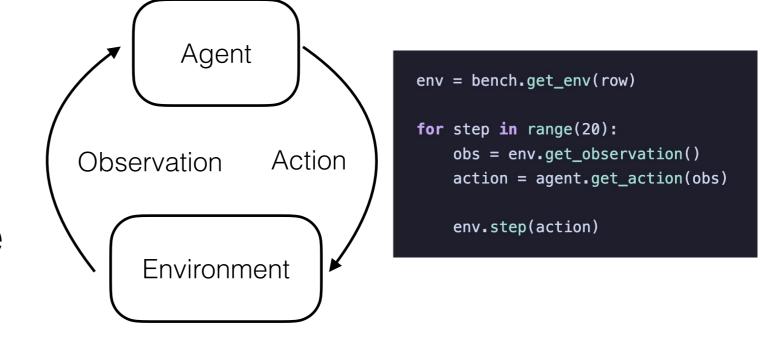
go_forward()

Description: Navigate to the next page in history.

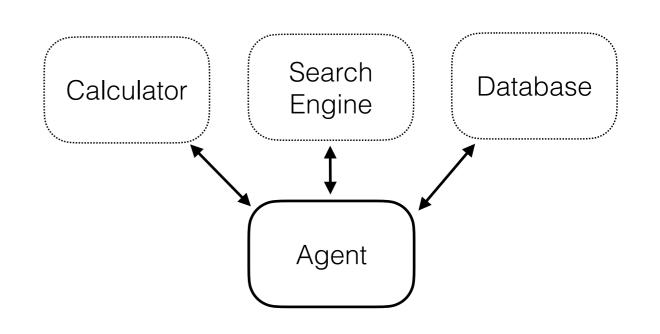
Examples:
    qo forward()
```

Two perspectives on taking actions

 Reinforcement learning: submit an action to the environment, receive a observation

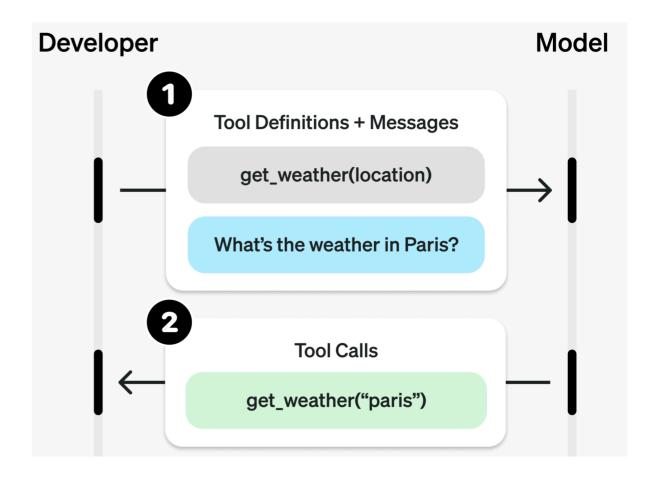


 Tool calling: model calls an external function, the environment returns the output



Tool calling interface

- Tool calling interface: a standardized way to define and invoke tools
 - Developer 1: defines tool as JSON
 - Model: generates tool calls
 - Developer 2: executes tool call, sends output



Tool calling interface

```
Tool Definitions + Messages

get_weather(location)

What's the weather in Paris?

Tool Calls

get_weather("paris")
```

```
"role": "assistant",
  "content": null,
  "tool_calls": [{
      "id": "call_abc123",
      "type": "function",
      "function": {
            "name": "get_weather",
            "arguments": "{\"location\": \"Paris\",
            }
      }]
}
```

Example: Claude computer-use

```
Action = Literal[
    "key",
    "type",
    "mouse_move",
    "left_click",
    "left_click_drag",
    "right_click",
    "middle_click",
    "double_click",
    "screenshot",
    "cursor_position",
```

```
Command = Literal[
    "view",
    "create",
    "str_replace",
    "insert",
    "undo_edit",
]
```

```
A tool that allows the agent to run bash commands.

The tool parameters are defined by Anthropic and are not editable.

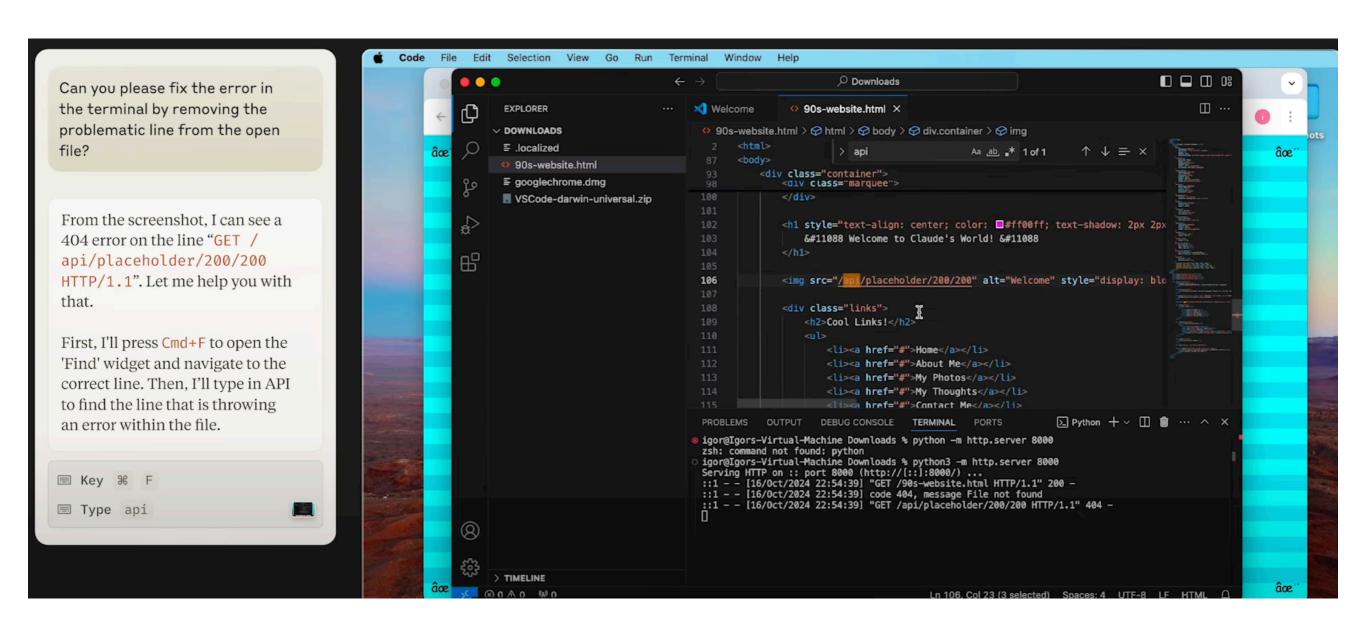
"""

_session: _BashSession | None

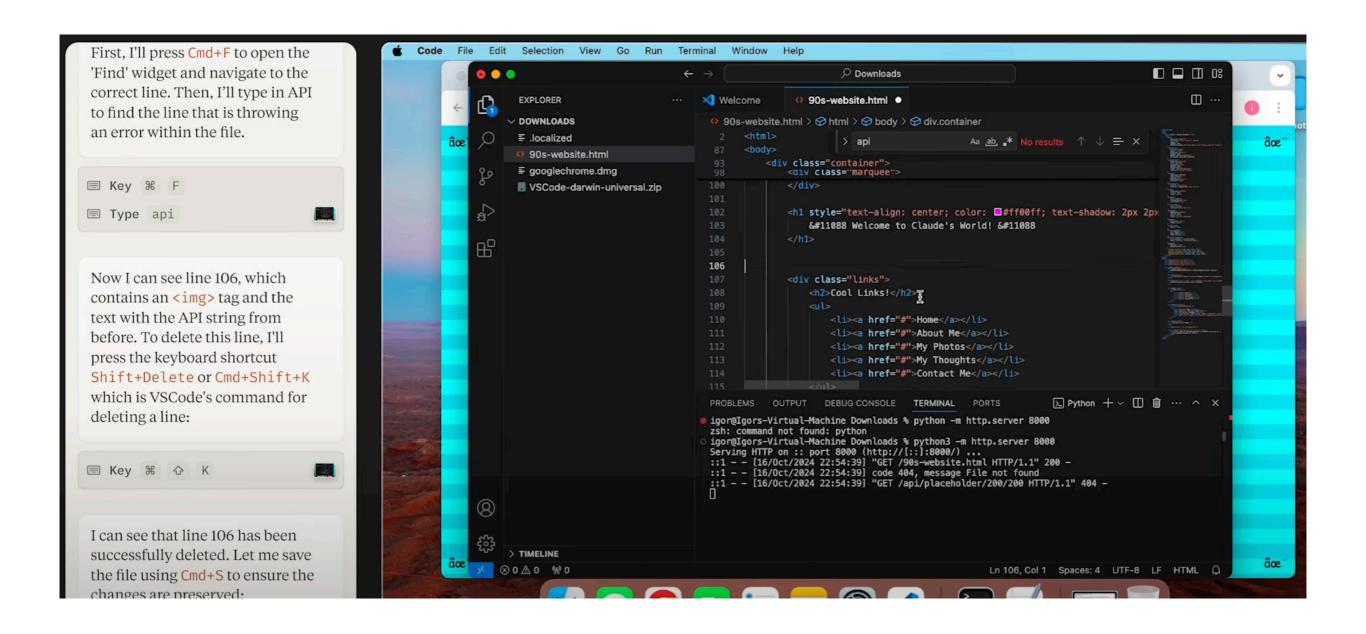
name: ClassVar[Literal["bash"]] = "bash"

api_type: ClassVar[Literal["bash_20241022"]] = "bash_20241022"
```

Example: Claude computer-use



Example: Claude computer-use



Today: (LM) Agents

- What is an agent?
- Agent environments
- Agent design patterns
- Agents and RL

Agents and RL

- Key differences
 - Environment provides new information
 - a_t : "search [...]" $\rightarrow s_{t+1}$: search results
 - a_t : click x y $\rightarrow s_{t+1}$: screenshot after click
 - Multi-step MDP instead of 1-step MDP
 - Environment interaction is often expensive
 - Tasks are often long

Example: RL for computer-use

- $a_t \sim \pi_{\theta}(a_t | x, s_t)$
 - x: task instruction
 - s: screenshot, history of observations and actions
 - a: chain of thought + action
 - => trajectory $\tau = (x, (s_1, a_1), (s_2, a_2), ..., (s_T, a_T))$
- GRPO at a step level

$$\text{ratio}_{t,k} = \frac{\pi_{\theta}(a_{t,k} \,|\, x, s_t, a_{t, < k})}{\pi_{\theta_{old}}(a_{t,k} \,|\, x, s_t, a_{t, < k})}, \text{ single grouped-advantage}$$
 for all steps

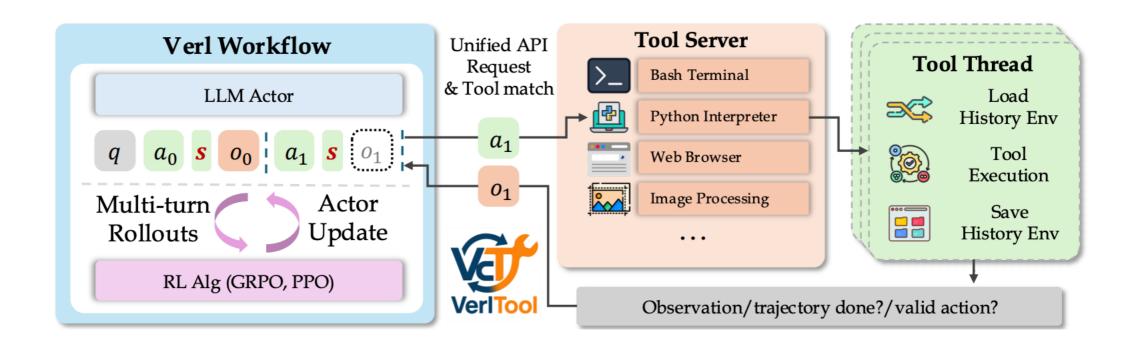
- Innovation: only apply loss on high-entropy states
- Reward: environment provides it by checking the environment state or using a VLM
- CPU Bottleneck: cluster of 180 environment instances

EFFICIENT MULTI-TURN RL FOR GUI AGENTS VIA DE-COUPLED TRAINING AND ADAPTIVE DATA CURATION

Pengxiang Li^{1,2*}, Zechen Hu^{3*}, Zirui Shang^{1,2*}, Jingrong Wu^{3*}, Yang Liu², Hui Liu³, Zhi Gao^{1,2⊠}, Chenrui Shi^{1,2}, Bofei Zhang², Zihao Zhang³, Xiaochuan Shi³, Zedong Yu^{2,4}, Yuwei Wu^{1,5⊠}, Xinxiao Wu^{1,5}, Yunde Jia⁵, Liuyu Xiang⁴, Zhaofeng He⁴, Qing Li^{2™}
¹Beijing Institute of Technology ²State Key Laboratory of General Artificial Intelligence, BIGAI ³DataCanvas ⁴Beijing University of Posts and Telecommunications ⁵Shenzhen MSU-BIT University https://computer-use-agents.github.io/dart-gui

```
System Prompt for GUI Agent
You are a GUI agent. You are given a task and your action history, with screenshots. You need to
perform the next action to complete the task.
## Output Format
Thought: ...
Action: ...
## Action Space
click(start_box='<|box_start|>(x1,y1)<|box_end|>')
left_double(start_box='<|box_start|>(x1,y1)<|box_end|>')
right_single(start_box='<|box_start|>(x1,y1)<|box_end|>')
drag(start_box='<|box_start|>(x1,y1)<|box_end|>',
     end_box='<|box_start|>(x3,y3)<|box_end|>')
type(content=") # If you want to submit, use "\n" at the end
scroll(start_box='<|box_start|>(x1,y1)<|box_end|>',
        direction='down or up or right or left')
wait() # Sleep for 5s and take a screenshot
finished(content='xxx') # Use escape characters \', \", \n
• Use {language} in Thought part.
• Write a small plan and finally summarize your next action (with its target element) in one sentence in
• My computer's password is 'password', feel free to use it when you need sudo rights.
## User Instruction
{instruction}
```

Example: RL for tool-use



VERLTOOL: TOWARDS HOLISTIC AGENTIC REINFORCEMENT LEARNING WITH TOOL USE

- **♣**Dongfu Jiang*†, $^{\Diamond}$ Yi Lu*, $^{\varsigma}$ Zhuofeng Li*, $^{\clubsuit}$ Zhiheng Lyu*, $^{\varrho}$ Ping Nie, $^{\vartheta,\theta}$ Haozhe Wang,
- ♣Alex Su, ^εHui Chen, ^φKai Zou, ♠Chao Du, ♠Tianyu Pang, ♣Wenhu Chen †
- *University of Waterloo, *Sea AI Lab, $^{\diamondsuit}$ University of Toronto, $^{\varsigma}$ Shanghai University, $^{\vartheta}$ HKUST, $^{\vartheta}$ M-A-P, $^{\varepsilon}$ National University of Singapore, $^{\varphi}$ NetMind.AI, $^{\varrho}$ Independent {dongfu.jiang@uwaterloo.ca, wenhuchen@uwaterloo.ca}

Recap: Agent Basics

- What is an agent?
- Agent environments
- Agent design patterns
- Agents and RL

Thank you